KPMG

The future of SaaS

A decision-maker's guide to understanding headless architecture and unbundling the enterprise stack



kpmg.com

AI generated image

For more than three decades, enterprise application buyers made a bargain. In exchange for the certainty of monolithic SaaS, companies effectively outsourced their core infrastructure to vendors who had spent years mastering these disciplines. ERP, CRM, and the workflow platforms that followed were never acquired purely for their features. They were also a deliberate decision to let someone else carry the burden of getting it right. And until very recently, the trade was a good one.

That bargain is now coming undone, but not for the reason most technology decision-makers assume. The emerging narrative—that agentic AI is poised to dismantle SaaS—conflates two distinct shifts. While agentic AI is changing how work gets done, headless architecture and natural language processing are changing where work gets done. For companies, treating these changes as one challenge obscures the real questions facing enterprise leaders: Which risks are they ready to take back, and where is doing so worth it?

But asking that question exposes an even more uncomfortable issue. The teams that own the CRM, ERP, and supply-chain roadmaps are rarely the teams piloting agentic AI. Similarly, the buyer of a headless experience layer is rarely the buyer of the system of record beneath it. And the newest tension, which is whether to adopt the SaaS vendor's own agentic layer or pair an enterprise back end with an independent model from a frontier lab, has no natural owner inside most organizations at all. The decision the market is presenting is not, at root, a technical one. Instead, it's about where accountability now sits, and which executive is willing to sign for it.

The future of enterprise applications is neither a wholesale move to custom agentic stacks nor a passive wait for incumbent vendors to catch up. Instead, it is a selective decoupling: a deliberate, process-by-process judgment about which parts of the stack still earn their keep as vendor-managed risk, which value streams justify owning the experience and orchestration layer, and which warrant building from scratch.

> The winners in this next chapter will not be defined by how aggressively they unbundle, but rather, how precisely they do it.



© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

One enterprise, many answers

Most executives are asking the wrong question. The market frames it as binary: Stay with the SaaS vendor or move to an agentic future. The boardroom version is more nuanced but no less reductive: Pick a lane—vendor-led, hybrid headless, or fully custom—and commit. Neither framing survives contact with how enterprises actually run.

A modern enterprise is not one process but is instead dozens of value streams, each with its own outcome, risk profile, differentiation potential, regulatory exposure, and tolerance for autonomy. The financial close, for example, does not require what the customer-service experience requires. Procurement does not differentiate a company in the market the way that sales does. A pricing engine in algorithmic trading is not a payroll run. Asking a single answer to cover all of them produces either over-investment in places that don't reward it or under-investment in the places that do.

A more useful question is structural. For each value stream, leaders must decide three things in sequence:

First, where does vendor-managed risk still earn its premium?

The thirty-year bargain was never bad economics. For commoditized, compliance-heavy, internally-facing processes, the SaaS vendor's controls, audit posture, and patch cadence remain cheaper and safer than anything an enterprise could rebuild. Staying the course is a correct reading of where the risk-transfer trade still pays, and is likely the most common approach for most organizations.

Next, where does the experience layer itself create competitive advantage?

This is where headless earns its place. When a value stream crosses multiple systems of record, such as in a B2B sales motion pulling from CRM, ERP, and a product catalog, the friction lives between applications, not inside them. Detaching the usage layer from the control layer lets the enterprise build a single surface where work happens, while leaving the systems of record to do what they were engineered to do. The advantage here is the orchestration as opposed to the AI itself.

Finally, where is the process itself the product?

A rarer set of value streams, such as algorithmic trading, specialized logistics routing, or proprietary underwriting models, are differentiated by the workflow itself. For these, even a best-in-class SaaS process is a ceiling, not a floor, and building a custom agentic stack on a modern data platform is the only way to push past it. The upside, when the workflow genuinely is the product, can reset the economics of the business. One supermarket chain that moved procurement onto a fully autonomous agentic system that replaced manual decisions across thousands of SKUs cut waste by 40 percent, reduced stockouts by 80 percent, and doubled its EBITDA margin.

The cost, however, is total risk reassumption: the organization takes on security, integration, governance, model behavior, and hallucination management. The justification is that the workflow is what the customer is paying for.

Sequenced this way, the question stops being which path the enterprise picks and starts being which path each value stream earns. KPMG believes most large enterprises will land on a portfolio of strategies that incorporate vendor-led for the back office, headless across the customer-facing and cross-system value streams where experience is the differentiator, and custom AI-first in the small number of places where the process itself is proprietary. That portfolio is what selective decoupling looks like in practice.

> The question stops being which path the enterprise picks and starts being which path each value stream earns.

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

Two complications make this harder than the framework suggests

The first is organizational. Headless experiences operate across value streams, but most organizations are still structured within them. Until business-process ownership catches up to the architecture, the right answer on paper will struggle to land.

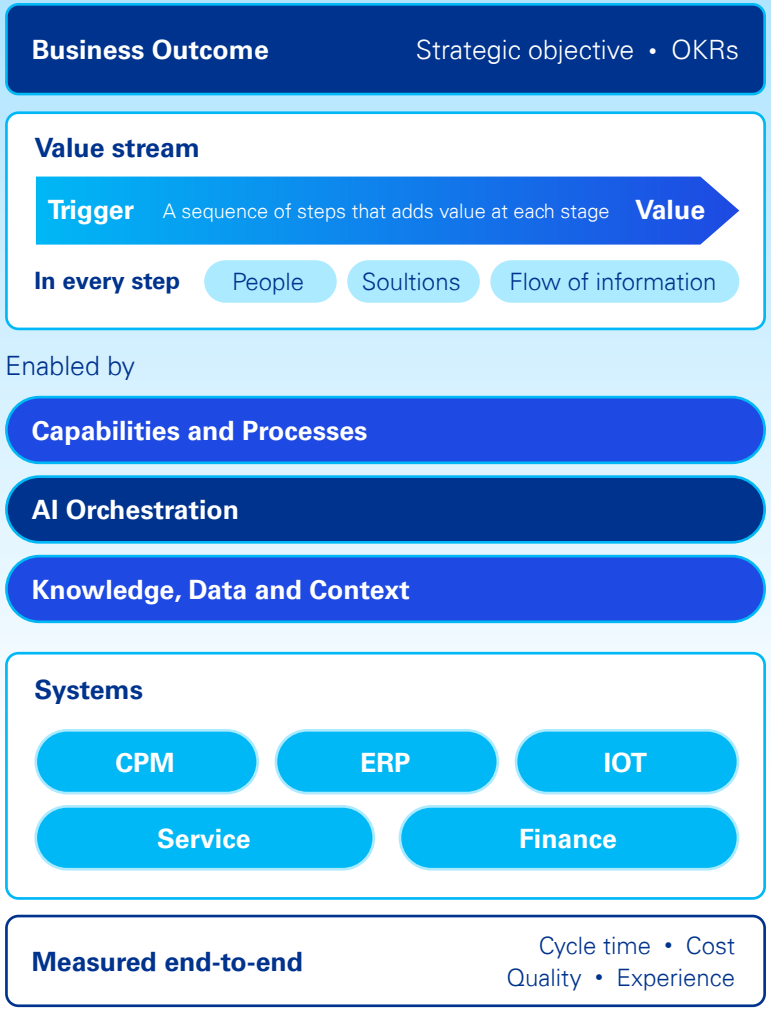
The second is commercial. Even within the hybrid path, a more subtle decision awaits: Which brain will orchestrate the work? The reality is that the best headless architectures will intelligently leverage both. The strategic question is which model will serve as the primary orchestrator—the "harness" that calls upon other specialized agents. The vendors argue, with some justification, that they can deliver a more reliable agentic experience on their own back ends. The frontier labs argue, also with some justification, that their general-purpose models are advancing faster than any application vendor's roadmap. There is no industry-wide right answer yet. There is only a per-value-stream judgment about which harness is more likely to be ahead in 18 months.

The modern value stream

A value stream is the end-to-end sequence of activities an organization performs to deliver value to a stakeholder, whether internal or external, from initial trigger to final outcome. Rather than describing a single function, a value stream represents a strategic capability.

Modern examples include:

- The intelligent lead-to-loyalty journey, which orchestrates marketing, sales, and service interactions across CRM, ERP, and communication platforms
- Predictive field service, which combines IoT data, logistics, and technician dispatch systems to anticipate maintenance needs and proactively resolve issues
- Autonomous financial forecasting, which integrates internal performance data with external market signals to generate continuous, real-time financial models.



© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

What headless changes

Strip away the jargon and the "headless" idea is simple: Separate the interface where work happens from the system that records it. The "head" is the front end, including the screen, the form, the chat window, and the workflow that users, and increasingly AI agents, interact with. The "body" is the back end: the database, the business logic, the controls, the audit trail, and the system of record. Enterprise applications fused the two, but headless pulls them apart.

When the head is detached, the interface relocates. The systems of record still do their job, but they are no longer where humans show up to work.

When the head and body are fused, work has an address. Approving an invoice means logging into the ERP. Updating a customer record means logging into the CRM. Submitting an expense means logging into a different system again. Employees navigate a portfolio of applications, each with its own login, its own conventions, and its own learning curve. The application *is* the workplace.

When the head is detached, the interface relocates. A finance manager can update Workday from inside the Gemini window they already have open.[1] A sales rep can move an opportunity forward from Slack without opening Salesforce.[2] A procurement lead can approve a nonstandard PO from a unified portal that routes the request through SAP, the contract repository, and the policy engine without touching any of them.[3] The systems of record still do their job, but they are no longer where humans show up to work.

The major SaaS providers have read the same signal. In their own ways, they are opening their stacks and acknowledging that some of the work their applications used to host is now going to happen elsewhere. For example, SAP's Business AI Platform integrates with frontier models and non-SAP platforms, making its governed Joule agents available for headless scenarios.[4] Similarly, Oracle has integrated core infrastructure into its Fusion Cloud, including model context protocol and agent-to-agent support, enabling Fusion Agents to be published, discovered, and invoked by external and third-party systems.[5] Salesforce and Anthropic have launched Claudeforce, an expanded strategic partnership that embeds Salesforce's CRM data, workflows, and business logic directly into Anthropic's Claude AI platform.[6] At ServiceNow, Action Fabric opens its full system of action to any external AI agent,[7] and at Workday, Sana AI embeds core HR and finance workflows directly inside Microsoft Copilot to keep users in their primary workspace.[8] Workday calls its version of this shift "breaking the box," and the metaphor travels:
The old model, like Apple's tightly integrated hardware-software-services stack, gave the buyer the whole experience top to bottom.

The new model concedes that no enterprise runs on one vendor's stack alone, and that the experience layer can live above all of them. The pace and the posture vary, but a year ago, headless was a thesis. Today, it is a product category in which the incumbents compete. That competition is reshaping the buyer's question: Which control layers do we trust, which usage layer do we own, and which agentic capabilities do we buy versus build?

[1] Natalie Morgan, "Workday Pushes Its HR and Finance Agents Inside Google's Gemini Enterprise," Sentinel.ht, June 21, 2026.
[2] Ashley Mao and Phillipe Ozil, "Connect Slack to Salesforce with Hosted MCP Servers," Developer.Salesforce.com, July 15, 2026.
[3] "Flexible Workflows for Procurement in SAP S/4HANA," Community.SAP.com, October 3, 2025.
[4] Phillip Herzig, "SAP and Anthropic Plan to Bring Claude to SAP Business AI Platform," News.SAP.com, May 12, 2026.
[5] "Oracle Advances Enterprise AI with New Agents Across Fusion Applications," Oracle.com, October 15, 2025.
[6] Meredith Anglin, "Claudeforce Explained: Salesforce and Anthropic's New AI-CRM Partnership," SalesforceBreak.com, August 2026.
[7] Nirankush Panchbhai, "ServiceNow opens its full system of action to every AI agent in the enterprise," Newsroom.ServiceNow.com, May 5, 2026.
[8] "Workday Brings Sana Self-Service Agent for HR and Finance Into Microsoft 365 Copilot," Investor.Workday.com, May 13, 2026.

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

The productivity case sharpens once the question is framed that way, because headless unlocks two distinct kinds of gain. The first is that humans do the work they already do, but faster and with less friction. Instead of learning the conventions of half a dozen application UIs, an employee expresses intent in natural language, and the agent navigates the systems on their behalf. The path to the work changes, but not the work itself.

The second gain is more consequential and slower to arrive: Some work that humans used to do is no longer done by humans at all. The agent handles it end-to-end, and the human is involved only by exception. The first gain is available across most of the enterprise within a planning horizon most CFOs would recognize. The second is a longer redesign, governed by calibrated autonomy. Promising the second when only the first is real or settling for the first when the second is achievable, is the most common way headless investments disappoint.

The control plane is the orchestration engine

To understand how headless architecture functions in practice, enterprise leaders must look to the critical layer that makes it possible—the control plane. If the natural language interface is the head and the underlying systems of record (ERP, CRM, SCM) are the body, the control plane is the vital connective tissue between them.

Sitting directly beneath the user experience layer, the control plane provides a system-agnostic orchestration environment. This is where the cross-system APIs, business logic, and unified governance controls now reside. Rather than forcing a user to jump between a Salesforce workflow and a Workday approval, for example, the control plane abstracts that complexity. It translates a single, conversational user intent into a sequence of actions distributed across multiple back-end applications, ultimately delivering a common, frictionless experience.

This system-agnostic layer is what makes modern, cross-functional value streams possible. By centralizing the orchestration, the control plane provides the technical foundation needed to enforce dynamic security, maintain audit trails across disparate applications, and execute the calibrated decision rights required for an AI-enabled workforce.

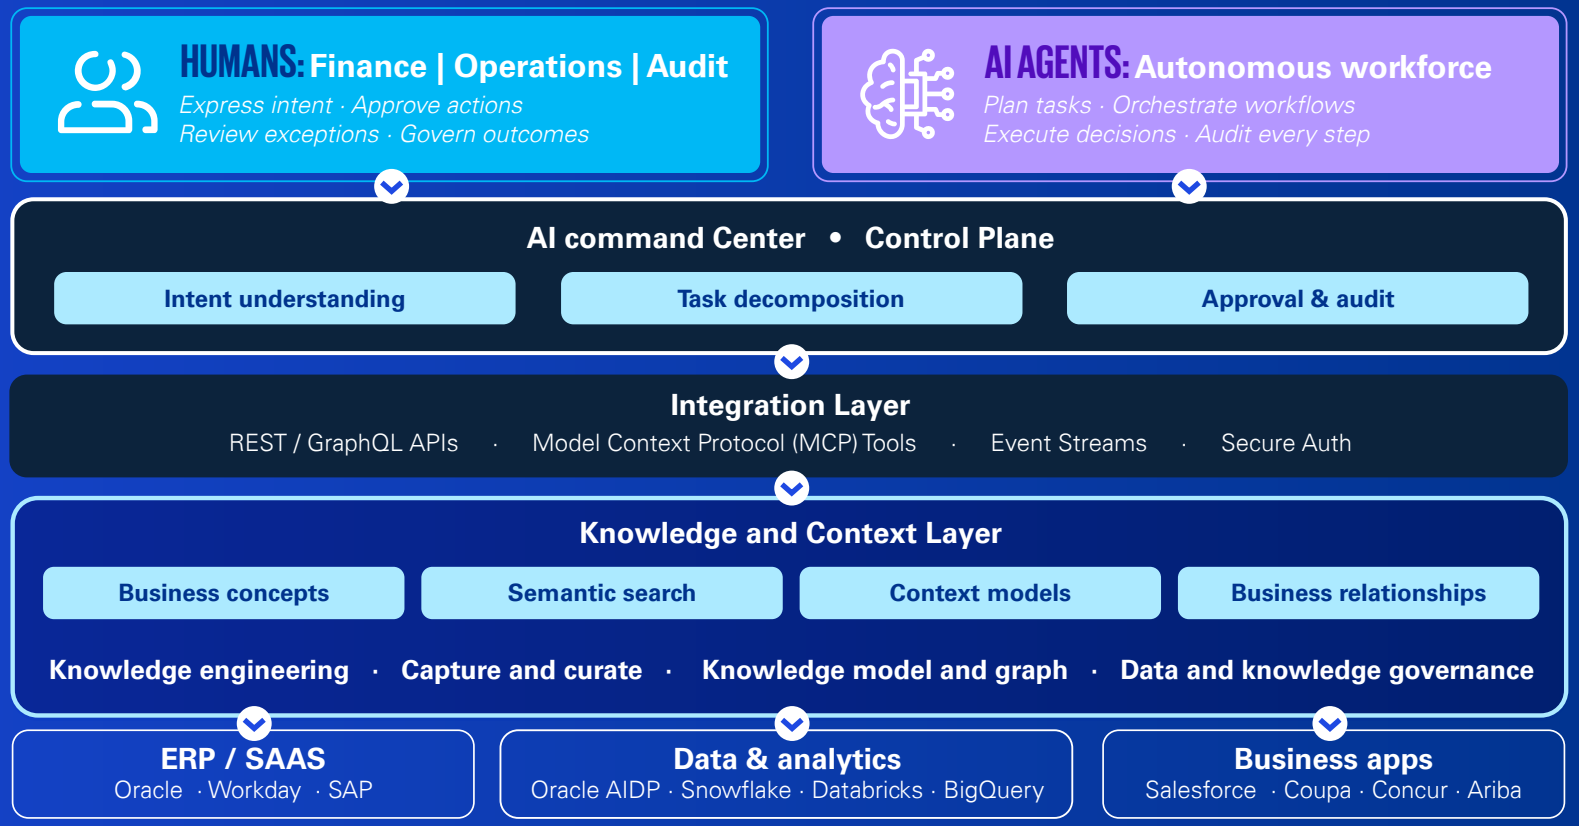
© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

Three implications follow, and they matter more than the architecture itself

1 The interface stops being a tax and starts being a strategy. When every system has its own front end, the interface is something to be trained around. When the interface is decoupled, it becomes a place to design. The way a service rep, a buyer, a planner, or a finance manager experiences their work is now an enterprise design decision instead of a vendor inheritance.

Orchestration becomes a first-class capability. A headless experience that touches three systems becomes a workflow that has to reason, handle exceptions, and be governed. Agent-to-agent coordination—one agent calling another, with the right context, the right permissions, and the right fallback—is harder than the demos suggest. Most enterprises do not yet have a team that owns this.

2 The data foundation either pays off or shows up as a bill. A headless usage layer is only as good as the data it can reach. If enterprise data is searchable, contextualized, governed, and trusted, agents can act on it. If it is not, the headless interface becomes a fast way to discover, in production, how unready the underlying data really is. The companies furthest along on data products, semantic layers, and AI-ready data foundations will move fastest. The companies that skipped that work will find headless does not let them skip it any longer.



3 None of this means the application interface disappears. For many commoditized value streams, the vendor's integrated experience remains the right answer. Where a headless architecture is chosen, it is not a vote against vendor-led AI but is instead an architectural choice to leverage it more effectively. A headless model will still call upon the specialized AI agents embedded within the core systems—for example, using a Workday agent for an HR task or a Salesforce agent for a sales query—but it does so from a unified experience layer. Headless is not a destination every enterprise is heading toward, but a powerful capability for the value streams that require orchestrating intelligence across the stack.

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

Calibrated autonomy: Where human judgment belongs

The conversation about agentic AI tends to collapse into two unhelpful extremes. One camp insists that humans must review every agent action, which preserves trust but eliminates the productivity case. The other promises full automation of entire functions, which makes for compelling demos and unmanageable risk. Neither describes how serious enterprises will operate in reality.

A more useful frame is calibrated autonomy: matching the level of human involvement to the risk, reversibility, and judgment intensity of the work itself. Some decisions warrant a human in every loop. Most do not. A handful can be handed to agents entirely, with humans intervening only when the agent flags an exception. Picking the right setting for the right process is the operating-model question of the next decade.



© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

Three models are emerging in practice

Full human oversight:

Every agent action is reviewed before it takes effect. The agent drafts, recommends, or proposes; the human approves. This is the right setting for high-stakes, low-reversibility, judgment-heavy work, such as credit decisions in regulated lending, clinical recommendations, legal advice, anything that affects a customer's standing or an employee's record in ways that are hard to undo. The productivity gain here is real but bounded; the value is in compressing the cycle time between draft and decision.

Consider Workday's Sana AI. Employees faced productivity loss from switching tools for routine Workday HR/finance tasks. Sana AI embedded into Microsoft Copilot executes tasks within existing approvals and security rules, with managers retaining sign-off. Streamlined ERP access, coupled with human-governed approvals, enables employees to complete HR tasks and transactions directly within Microsoft Copilot, reducing the need to switch between systems.

Shared oversight, or the messy middle:

Humans and agents divide responsibility within a workflow. The agent handles the bulk of routine actions; the human handles the parts that require context, judgment, or relationship. The two pass work back and forth, often within the same case. This is where most enterprises will spend the next several years. A customer-service interaction where the agent gathers information, drafts a response, and routes complex cases to a human; a procurement workflow where the agent negotiates standard terms and escalates nonstandard ones; a financial reconciliation where the agent matches the routine 95 percent and the human investigates the rest. Shared oversight is harder to design than either extreme because the handoffs are where errors hide, but it is also where realistic productivity gains live.

Rimini Street, an enterprise IT services provider, wanted to eliminate the manual grind of prepping for client engagements. To meet the challenge, they deployed a "Deep Research" agentic capability integrated directly across its CRM and ERP systems. The agent automatically synthesizes 360-degree customer profiles from disparate internal and external data sources, helping service teams slash pre-call research time by nearly 40 percent under a co-owned governance framework where humans and AI jointly shape customer strategy.

Exception-only intervention:

The agent operates autonomously within a defined envelope. Humans engage only when the agent encounters something outside its envelope, such as a transaction it cannot reconcile, a policy condition it cannot interpret, or a counterparty it does not recognize. This is the right setting for high-volume, well-bounded, well-instrumented work like standard invoice matching, routine entitlement checks, repetitive data transformations, or automated tier-one support for known issues. Exception-only is the relocation of human judgment to the edges of the process, where it is most valuable, and not the absence of human judgment altogether.

In North Carolina, the City of Raleigh built a virtual assistant named "Ral-E" to streamline answers to routine requests from citizens, routing only complex anomalies to staff. This exception-only model achieved a 98 percent deflection rate, resolving public queries independently and saving city staff the equivalent of a full month of labor.

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

Redefining the target operating model and user roles

As enterprises calibrate the autonomy of their systems, the target operating model must fundamentally transform alongside it. When headless architecture and agentic AI take on the execution of a value stream, the human role shifts from process execution to process governance and exception handling.

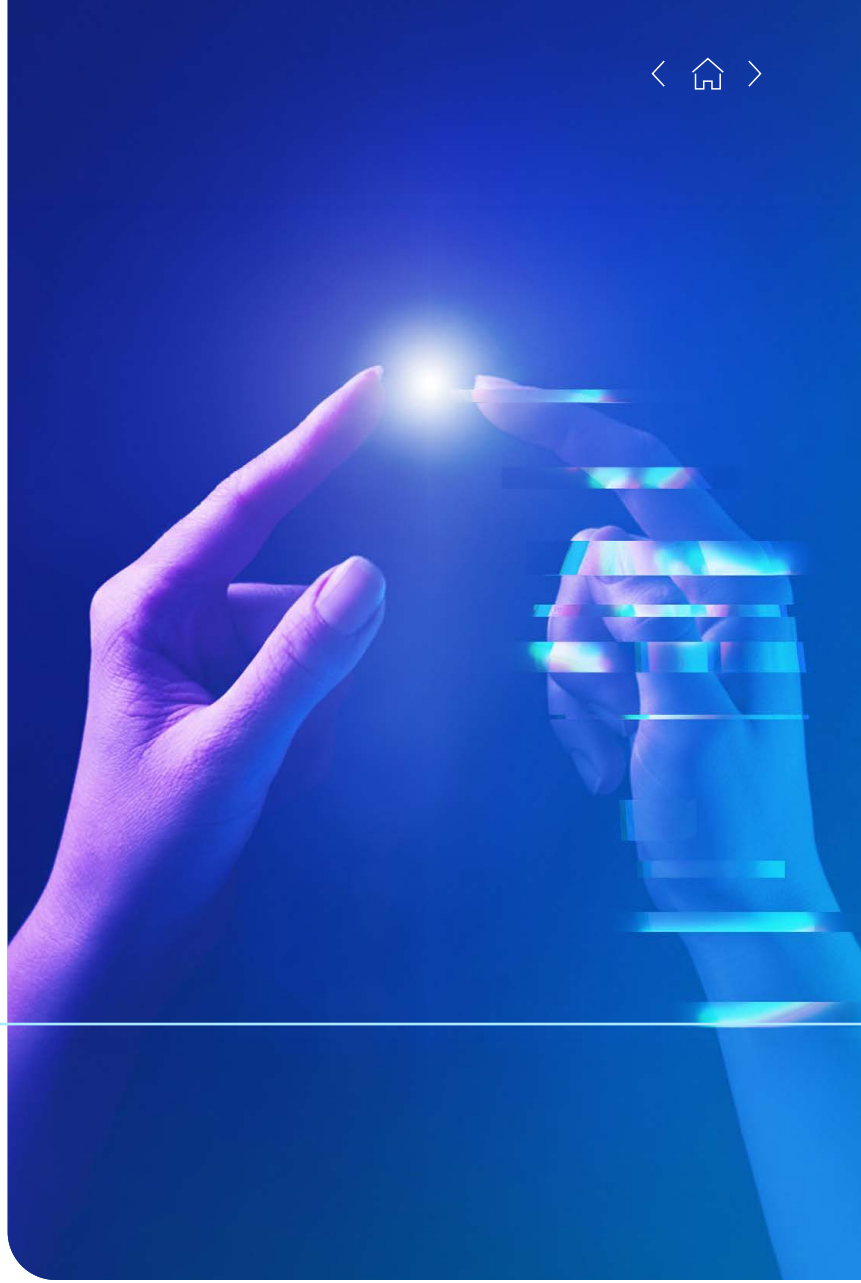
Consider a traditional accounts payable function. In a legacy ERP environment, an "AP Admin" spends their day manually matching invoices, flagging discrepancies, and routing approvals. In a converged, headless model, the AI agent performs the matching and routing autonomously within its calibrated risk envelope. The human role elevates to an overseer who only intervenes when the agent flags a high-risk anomaly or a scenario it hasn't seen before.

This shift requires a complete reimagining of enterprise controls and segregation of duties (SoD). In this new paradigm, SoD must be redesigned to govern the relationship between human and machine. Controls must focus on who has the authority to change the agent's underlying rules, who monitors its audit logs, and how human oversight is applied to the AI's autonomous decisions.

Exception-only is not fully automated

It is easy to read the third model as the end-state every process is heading toward. It is not. The monthly financial close, often cited as the canonical candidate for full automation, is a useful corrective. Even with mature agentic tooling, the close involves judgment calls that include materiality thresholds, accrual estimates, classification of unusual transactions, and audit-trail explanations that do not survive removal of the human entirely. A well-run close may move from substantial manual effort to exception-only intervention, with the controller's team focused on the items the agents flag. And although that is a significant operating-model shift, it is not the end of the controller's team.

In this new paradigm, SoD must be redesigned to govern the relationship between human and machine.



© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

Calibrated autonomy is a portfolio, not a destination

Just as each value stream earns its own architectural path, each process within a value stream earns its own autonomy setting. A single customer-service value stream might run exception-only for password resets, shared oversight for billing disputes, and full human oversight for fraud claims. The mature operating model is not "we moved to exception-only" but "we know which setting each process is on, and we revisit it as the agents and the data improve."

This is also where the organizational question returns. Picking the right autonomy setting requires a new kind of leader: a Converged Stack Owner. This role's mandate is to own the calibration decision itself by understanding the risk profile of the entire value stream and having the authority to redesign the human role around the new automation envelope. In most enterprises, this leader does not yet exist, because no one owns the end-to-end calibration of risk and autonomy.

This is also where the context layer and the control plane play a decisive role. A well-defined context layer which combines deterministic business rules with the probabilistic outputs of AI models is what enables greater decision autonomy with auditability. The control plane enforces these calibrated settings, ensuring that an agent operating in exception-only mode has a firm guardrail and a clear path for escalation. Until business-process ownership is reorganized around value streams, calibrated autonomy will be set by default rather than by design.

The productivity case for agentic AI is real, but it is unlocked by calibration, not automation. The enterprises that move fastest will be the ones that know, process by process, where the human belongs and when it's time for technology to take the reins.



© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

How to decide which path each value stream earns

Although calibration decides where the human belongs, it does not decide which path the value stream should take. The framework so far has argued that selective decoupling is a portfolio judgment, not an enterprise-wide pick. This section makes that judgment operable. It does not hand the reader a rubric—the rubric belongs in client conversations, not in a published paper—but it equips the executive team to ask the right questions of each value stream in turn, in the right order, with the right people in the room.

The decision turns on five questions, sequenced deliberately. The order matters: Each question filters the value streams that survive into the next.

1. Is the process a source of competitive differentiation or a driver of operational efficiency?

This is the threshold question and the one that reframes the entire portfolio. The distinction is not between important and unimportant. Core back-office functions such as payroll, the monthly close, and statutory reporting, for example, are critical. But they do not typically differentiate an enterprise from its competitors; a process differentiates only when how the work is done affects whether a customer chooses one company over another.

KPMG has observed that this is where the analysis often goes wrong. The conclusion that a process is not a market differentiator is often misinterpreted as a sign to stay the course with a traditional ERP. The better reading is that these processes are prime candidates for AI and automation, but the business case is built on efficiency, speed, and cost reduction rather than top-line growth. The back office will lean heavily into AI, but the question is whether that occurs through the ERP's embedded agentic tools, which prioritize safety and integration, or through a headless or custom build that prioritizes performance and pushes the boundaries of automation.

2. Does the work span multiple systems of record or live inside one?

Headless earns its place when the friction lives between applications. A B2B quote-to-cash motion that crosses CRM, configure-price-quote, ERP, and a product catalog is a different problem from a sales-forecasting workflow that lives entirely inside the CRM. The first is a candidate for a unified usage layer; the second is a candidate for the vendor's embedded agentic capabilities.

The mistake to avoid here is assuming integration complexity equals headless suitability. Some cross-system processes are better solved by the SaaS vendor's own orchestration; for example, where the vendor controls the dominant system of record in the chain. The strategic question is about the degree of ownership the enterprise is prepared to take. Should it design a fully custom, end-to-end orchestration layer that crosses all systems, or should it buy into a primary vendor's ecosystem and use their tools to manage the handoffs? The right answer depends on how much differentiation the orchestration itself is expected to create.

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

3. Is there a converged stack owner with authority across the value stream?

This is the question most enterprises skip. The honest answer for most large enterprises is that the owner does not exist yet for most value streams.

A unified usage layer needs a single accountable executive with authority to govern the entire converged stack. This leader must have the budget to fund the orchestration layer, the cross-functional standing to arbitrate when the CRM, ERP, and AI teams disagree, and the mandate to make the final call on the levels of automation and human intervention. If that owner does not exist, the value stream is not ready for a headless strategy, regardless of the technical justification.

4. Is the data foundation ready to be acted on?

Agents are only as capable as the data they can reach. A value stream where the underlying data is well-modeled, searchable, contextualized, and governed is a value stream where headless or custom-agentic capabilities will deliver. A value stream where the data is duplicated across systems, locked in unstructured fields, or governed by policies no one has refreshed in five years is one where the architecture will expose the data problem faster than it solves the experience problem.

This question often reorders the rest. A value stream with high differentiation potential, clear ownership, and an underdeveloped data foundation is not a candidate for headless this year. It is a candidate for a data-foundation investment that makes it a candidate next year. The architecture must follow the data, not lead it.

5. What is the risk envelope the enterprise is prepared to reassume?

Every value stream that moves away from vendor-led AI reassumes some portion of the risk inventory. For some value streams, like high-velocity, well-instrumented, low-blast-radius work, the reassumption is manageable. For others, such as regulated processes, high-stakes customer-facing decisions, or work with material legal or financial exposure, the reassumption may be wider than the differentiation justifies.

This question is also where the harness decision lives. Within a value stream that earns the hybrid path, the choice between the SaaS vendor's agentic layer and an independent frontier-lab model is a risk-envelope decision as much as a capability decision. The vendor's agentic layer carries more of the integration and governance risk on the vendor's side; the independent harness carries more capability ceiling on the enterprise's side. Neither is universally right. Both are answerable per value stream.

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

Assuming the risk and forging the path forward

The strategic upside of selective decoupling is real, but every value stream that earns a path away from vendor-led AI also earns a corresponding obligation. Before committing to a path, enterprise leaders must be clear-eyed about the risks they are choosing to reassume. These risks look very different depending on whether the chosen path is headless or fully custom.

Functional duplication

Headless architectures tempt enterprises to rebuild what the underlying system already does well. Validation rules, approval routing, role-based access, and exception handling represent years of accumulated edge-case logic. A custom orchestration layer must reimplement and maintain them. The discipline is knowing which functions to call from the body and which to replace, while resisting the urge to redesign the plumbing underneath.

Agent orchestration complexity

Coordinating one agent against one system is a tractable engineering problem. Coordinating dozens of agents across multiple systems is not, and agent-to-agent handoffs are where most pilot programs stall. Questions around agent authority are complex, making orchestration ownership a precondition for going headless at scale.

Value-stream ownership

Headless experiences operate across value streams, but KPMG has observed that most organizations are still structured within them. When the experience layer spans CRM, ERP, and beyond, the question of who owns the redesigned end-to-end workflow has no default answer. Without budget, authority, and accountability that crosses application boundaries, headless investments risk being designed by committee and executed by no one.

Data foundation exposure

A headless or custom-agentic experience makes problems with the state of the enterprise's data foundation immediately visible. If data is searchable, contextualized, governed, and trusted, agents act on it effectively. If it is duplicated across systems, locked in unstructured fields, or governed by policies no one has refreshed since the last platform migration, the agents will fail in ways the previous application UI used to hide.

Economic and commercial volatility

The shift from seat-based licensing to consumption-based pricing accelerates as work moves outside the application. For some workflows, consumption pricing is materially cheaper than the seat license it replaces. For others, such as agentic workflows that call reasoning models unnecessarily at every step, or headless interfaces that fan out across dozens of back-end APIs without coordinated rate management, the consumption bill can exceed the cost it was meant to displace. Workflow design determines the answer: Retrieval-based execution can run several times cheaper than reasoning-based execution for the same outcome. The cost differential is not a rounding error, and the architecture decisions made in the first quarter will determine whether the commercial model in the fourth quarter looks like a win or a write-down.

The early cloud migration is the right historical analogue, but the lesson is more specific than the bill outpacing the value. For the enterprises that instrumented cost early, cloud was genuinely cheaper—for most workloads. The ones that did not discovered that elastic infrastructure is also elastic in the wrong direction.

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

Key recommendations

With the risk inventory in view, four actions follow.

1

Prioritize the context layer before the experience layer

Data foundation exposure is the risk most likely to derail headless investments that are technically sound on paper. Before committing to a new experience layer, validate that your underlying data is ready to be acted on by investing in data products, semantic layers, and governance that create a single, reliable context layer agents can query with confidence. A value stream with strong differentiation potential but an underdeveloped data foundation is not a candidate for headless this year; it is a candidate for the data investment that makes it a candidate next year. The architecture must follow the data, not lead it.

2

Name a converged stack owner before committing to a path

Value-stream ownership gaps and agent orchestration complexity share a common root: No single leader holds end-to-end accountability for the combined stack. Before a headless or custom-agentic program moves beyond pilot, a single executive must be named with authority to govern the full portfolio of vendor-led, headless, and custom solutions; arbitrate between SaaS and AI vendors; and make the final call on automation calibration for each value stream. Without this role in place, the organizational question will reliably outlast the technical one.

3

Establish a tokenomics function to govern AI spend

AI is creating a new enterprise resource: intelligence. Organizations need more than token counting and monitoring and cost management. AI activity doesn't necessarily mean value. They need an AI Economics capability that measures not just consumption, but attribution, accepted work, value realization, resilience, and governance. An AI Economics Ledger should connect intelligence consumed to business outcomes, while an AI Value Office provides accountability for proving value, prioritizing investments, and making allocation decisions as AI scales across the enterprise.

4

Build governance at machine pace

Risk reassumption, functional duplication, and orchestration complexity all demand a governance model that can keep up. The control regime designed for quarterly application releases will not survive an environment where agents and models evolve continuously. Enterprises must invest in continuous monitoring of agent behavior, automated policy enforcement at the orchestration layer, and audit trails designed for nondeterministic systems—creating a feedback loop from production back into policy. Organizations that build this capability early will move faster and absorb less risk than those retrofitting old governance models onto a faster substrate.

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

The future of ERP is a series of enterprise decisions

The temptation is to declare a single winner in the shift to AI, but neither incumbent vendors nor frontier labs offer a complete answer for the modern enterprise. The reality is more demanding.

The future of enterprise applications is not a single destination but a series of deliberate decisions, taken one value stream at a time. It requires mastering the discipline of selective decoupling—knowing where vendor-managed risk still earns its premium, which experience layers are worth owning, and which proprietary workflows justify a custom build. It also requires understanding that the goal is not total automation but calibrated autonomy—thoughtfully deciding where human judgment remains indispensable. And the winners in this next chapter will be the companies whose executives understood that the question is how deliberately to unbundle the stack.

How KPMG can help

At KPMG, we understand that business leaders look to technology to enhance operational efficiency, enable data-driven decision-making, improve customer experiences, foster agile workforces, optimize supply chains, and strengthen resilience. We know it's crucial for organizations to align technology initiatives with overall business goals to truly succeed in today's fast-paced environment.

For more than a dozen years, we've built a leading technology organization, powered by more than 15,000 dedicated technology leaders and professionals, specifically designed to help both information technology and business leaders thrive at the accelerated pace business now demands. Unlike business-only consultancies, we bring battle tested tools, deep engineering skills, and strategic alliances with leading technology companies. And unlike technology-only firms, we offer robust business credentials and sector experience, committed to delivering measurable business results, not just technical solutions.

Whether we're helping you implement AI, improve data and analytics, enhance cybersecurity, leverage cloud technologies, deploy a new technology, or outsource challenges with our managed services, you can count on us to help deliver solutions quickly and effectively.

That's speed to modern technology.

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

Authors



Matteo Colombo

Principal, KPMG Global Leader for Cloud, Data, AI, KPMG LLP

E: matteocolombo@kpmg.com

Matteo serves as the global leader for Digital Technologies at KPMG. In this role, he spearheads the growth, transformation, and innovation across the KPMG Digital Foundation and the global Centers of Excellence for Cloud, Data, and AI. Matteo helps shape the use of technology in the firm's collective strategy and enables its effective implementation in local markets. Additionally, he manages relationships with key strategic technology partners. Matteo's experience encompasses both strategy and technology, with a particular emphasis on emerging trends that drive growth and innovation for clients and partners. He is a staunch advocate for the effective and ethical use of AI. With over 20 years of experience, Matteo has adeptly guided large enterprises in harnessing leading technologies to achieve significant, large-scale transformations.



Brian Anderson

Principal, Advisory, Enterprise Solutions, KPMG LLP

E: bcanderson@kpmg.com

With over 30 years of experience advising clients in the financial services industry and implementing strategies to improve visibility into business performance, Brian is a specialist in advancing change the bank programs within the finance function including large ERP transformations. Brian has a successful track record leading and managing numerous global, multimillion-dollar initiatives to transform finance via implementation of process improvements, new technology, and cost-efficient organizational structures. He helps optimize financial business processes and design front-to-back business architectures in banking and capital markets. Brian is an accomplished executive with strong leadership skills, deep relationships with technology vendors, and ability to drive consensus across cross-functional teams and geographies.

Brian currently leads the KPMG US Workday practice across HR, Financials, Procurement and EPM.

Contributors:

The authors thank Lisa Bigelow, Daniel Cohen, Mary Lewis, Michael Ogden, Tim Nexon, Rich Schulte, Vijay Subramanyam, and Lara Volpe for their contributions to this paper.

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

Related thought leadership



Agentic AI untangled: Navigating the build, buy, or borrow decision



Accelerating AI with managed services



Agents of change: New organizational roles in the age of AI

Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.

Please visit us: **kpmg.com** **Subscribe**

The information contained herein is of a general nature and is not intended to address the circumstances of any particular individual or entity. Although we endeavor to provide accurate and timely information, there can be no guarantee that such information is accurate as of the date it is received or that it will continue to be accurate in the future.
No one should act upon such information without appropriate professional advice after a thorough examination of the particular situation.

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

The KPMG name and logo are trademarks used under license by the independent member firms of the KPMG global organization. DASD-2026-20636