

# The Great SaaS Reckoning:

Perspectives From The Frontlines  
of an Agentic Enterprise Future



**GLG**

# Table of contents

<b><u>Introduction</u></b>	<b>3</b>
<b><u>The Virtual Collaborator and the End of the SaaS Interface</u></b>	<b>5</b>
By Mike Sabes	
The headless future · new map for SaaS market · outcome vs. consumption-based pricing · unit economics in agentic future	
<b><u>The Use Case Gap</u></b>	<b>10</b>
By Tom Liu	
Use case crux · purpose-built solutions · buy or build · agentic threat to incumbents · pricing: hybrid beats outcome-based · valuation logic	
<b><u>AI is a Sorting Machine, Not a Harbinger of Extinction</u></b>	<b>16</b>
By Jill Sweeney	
SaaS as sorting mechanism · why incumbents are leading · structural durability · the margin trap · orchestration layer as battleground	
<b><u>Core, Peripheral, and Everything at Stake</u></b>	<b>20</b>
By Andreas Welch	
Assessing SaaS risk: core vs peripheral · advantages of embeddedness · disruption at peripheries · governance gaps · budgeting and implementation for buyers	
<b><u>The Architecture of SaaS Market Share Durability</u></b>	<b>25</b>
By Robert Breynaert	
The three pricing models · making the outcome-based model work · value of incumbency · flaws in inertia thesis · interface as liability · diagnostic for investors	
<b><u>Trusted Autonomy</u></b>	<b>32</b>
By David Betancourt, Ph.D.	
Most exposed SaaS categories · P&L after AI · margin economics for vendors · trust as value	

# Introduction

Over a few days in February 2026, \$800 billion evaporated from one of the stock market's historically strongest categories: software as a service (SaaS).

The sudden sell-off coincided with an apparent realization about how AI changes software's place in the enterprise: The seat-based business model that made the economics of SaaS so reliably lucrative – and attractive to Wall Street – verges on the cusp of obsolescence.

The basic reasoning is hard to dispute. For one, AI means enterprises need fewer seats – licenses for individual workers to use the software, each coming with a price tag. With AI, a single user can streamline their tasks and workflows and therefore drastically multiply what they get out of their seat. In some cases, an agent can even operate a SaaS platform on its own, 24/7 – a virtual user who doesn't pay for a seat but demands much more from the platform than a human would, costing the provider further. Meanwhile, as customers look to reduce the seats they pay for, they're also using AI to engineer their own software solutions in-house at a much greater scale than they ever could before. Ambitious start-ups are doing the same, looking for every opportunity to displace legacy SaaS vendors.

As software giants reported revenue hits this winter and spring and struggled to demonstrate a compelling plan to sustain growth for the next couple years – let alone the next decade – investors scrambled to get ahead of an impending catastrophe. On the trading floor, and eventually in headlines and board rooms, a new phrase caught on: “SaaS-pocalypse.”

While the contractions at times resembled a mass extinction, the reality on the ground – at the enterprise, business unit, and user level – looks a lot more nuanced. To map its true topography and currents, we spoke to six of GLG's most sought-after AI experts, who have not only built accomplished careers in tech development, deployment, and implementation, but are leading the future of the space.

All six experts agree that while a pivotal disruption has arrived, the apocalyptic framing may do a disservice to leaders navigating important decisions about what to buy, where to invest, and how to build. The per-seat model that has driven so much growth in the industry is no doubt coming to an end – but what exactly replaces it isn't so clear, the experts show, nor are the markers of imminent winners and losers or the true speed of change.

## ■ Mike Sabes

Former Head of Corporate Finance & Strategy, Anthropic

## ■ Tom Liu

Managing Partner, Ideate Capital

## ■ Jill Sweeney

Former Strategy Director of High Performance Computing and Artificial Intelligence at Hewlett Packard Enterprise

## ■ Andreas Welsch

Former CMO of Business AI & VP of Product Management and GTM, SAP

## ■ Robert Breynaert

Founder, Cloudscale Partners

## ■ David Betancourt, Ph.D.

Head of AI at DB Macro Research; Former Apple AI Scientist

The following report highlights the nuances shaping debates around such questions, with each expert offering a different vantage point on the market. Here are a few of the key areas of analysis:



#### **How severe is the disruption, and how rapid?**

Welsch and Sweeney read what we're seeing as a consolidation: AI is a sorting mechanism repricing vulnerable platforms while strengthening the embedded ones. Liu and Sabes are more aggressive. They see a five-year threat to category leaders, which may arrive even sooner than the market predicts, especially outside the Fortune 500.



#### **How durable is the incumbent advantage?**

Breynaert makes the strongest case for incumbency. Governance infrastructure, data gravity, and security validation give leaders like Microsoft a position that is very difficult to dislodge, he argues. Liu pushes back. Those moats are real at the top of the enterprise pyramid but thin out in the SMB and mid-market, where adoption barriers are lower and the build-versus-buy calculus is shifting quickly. Sabes goes further, suggesting the "you never get fired for buying IBM" dynamic could invert in some categories as AI-native challengers close the gap faster than larger, slower incumbents can keep up.



#### **Does outcome-based pricing actually work for vendors?**

Breynaert lays out the most detailed case for making it viable. Sabes is measured: it works where outcomes are clearly defined, but faces challenges to scale broadly. Liu is skeptical, arguing the model reflects a fundamental limitation of standardized software rather than solving it. Software cannot customize deeply enough to reliably drive specific KPIs across customers with fundamentally different problems.



#### **Is peripheral versus core the right lens?**

Most of the experts use a version of this binary framing as an analytical starting point. Liu complicates it. Even tools commonly considered core, like CRMs, are vulnerable if they can't demonstrate clear use-case alignment to customers. The question isn't only where the software sits in the stack, but whether customers ever felt they were getting the value they paid for. Liu makes the most sustained case for the SMB opportunity, which gets underweighted compared to standard enterprise-first analyses.

The divergent outlooks don't reflect different levels of information, but different positions in the market, different investment theses, and different time horizons. Rather than arbitrating who is right, readers may find the most value in understanding which lens applies most directly to the companies, decisions, or investments they're evaluating.



# The Virtual Collaborator and the End of the SaaS Interface

By Mike Sabes

Former Head of Corporate Finance & Strategy, Anthropic

**One of the most clarifying moments of my career came in my final months at Twitter.**

It was time to negotiate our Salesforce renewal. Our most senior stakeholder made it clear to me we had no realistic option to switch. Simply, the implementation was too complex to recreate, and the switching costs were too high. We were addicted, so to speak – and Salesforce knew it. Six months later, I left Twitter and joined Salesforce.

The experience solidified in my mind what has made SaaS businesses so durable. In many cases, customers may not love the product, but the surrounding architecture locks them in: complex custom implementations, network effects, proprietary data accumulation, and mission-critical

workflows that cut across domains. Systems of record become the nucleus around which enterprise software gets built. Customers are retained at high rates, and a seemingly bullet-proof business model follows: keep adding new contract value to the bucket, don't let the bucket leak, and you can build a \$100B+ business.

With AI, that paradigm fundamentally changes. My combined experiences working at SaaS leaders like Adobe and Salesforce as well as Anthropic have informed my perspective on this topic.

## ■ The death of the per-seat model

Seat compression is real. More and more work is being done by agents rather than people, and any business measuring its value by the number of seats on its platform is mispricing in the new era. I can't think of a single SaaS business that should remain on pure seat-based licensing indefinitely, especially if it's claiming any meaningful AI-enabled value creation.

At the same time, the moats that enterprise SaaS has historically enjoyed aren't evaporating overnight. The cost to develop and implement software in a bespoke, custom way is collapsing, which means switching costs are much lower than they used to be. But other moats – like network effects and proprietary data accumulation – can be kept intact.

**“I can't think of a single SaaS business that should remain on pure seat-based licensing indefinitely, especially if it's claiming any meaningful AI-enabled value creation.”**

## ■ The virtual collaborator

If you imagine a world where much more of the work across all industries gets done by agents next to people, the direction we're headed becomes clearer: Every knowledge worker will be paired with a virtual collaborator. The central AI agent, supplied by a frontier lab, like Anthropic, will have access to all the user's systems: all the relevant context around every email, Slack message, meeting transcript, and dashboard. With that rich context, frontier-level intelligence, and access to whatever reference material the organization wants to guide it toward, the virtual collaborator becomes an extremely capable partner for getting important work done.

The implication is that people will spend their workday interacting primarily with their virtual collaborator, not with twenty different SaaS systems. This is what's forcing the concept of “headlessness” on the industry.

Many SaaS incumbents initially wanted to create walled gardens, retaining the attention and managing the interface where end users interact with their systems in an AI-enabled way. Think of going into Salesforce and using Agentforce. But even Salesforce has now announced their willingness to go completely headless, meaning end users don't even need to know they're using Salesforce. The CRM backend could be Salesforce, or it could be an AI-native competitor doing

the same thing much more efficiently and at a much lower price. This raises a sharp question about pricing power. What is a legacy SaaS business worth in a world where its end users never interact with it directly, and the virtual collaborator of choice is the only interface that matters?

Incumbents do have a massive head start, built on customer relationships, brand equity, and entrenched data. It's their race to lose. But going headless fundamentally changes the dynamic in ways that haven't fully been priced in yet.

### ■ Time horizon for adoption

At a company like Anthropic, this world already exists. Workers of all types are highly enabled to do their daily work by interacting with Claude above any other tool or system.

But the transformation timeline looks different at every organization. While the capability for some version of a virtual collaborator exists today, albeit in a rudimentary form, I expect it will take one to three years before most enterprise organizations embrace this vision. Change management is genuinely hard, especially when the change has visible implications for labor. Technology companies will move faster; industrial and retail organizations will be slower. Different capabilities diffuse across the economy at different speeds. And the organizations that adopt AI the fastest will have a serious advantage in maximizing productivity of their workforces.


### ■ Mapping SaaS today

Investors who are thinking through the strategic positioning of any given company can use a simple framework to categorize software businesses into three buckets.

The first is AI-augmented – what we see most often from legacy enterprise SaaS. These are workflow systems or systems of record that have bolted on AI features like a chatbot and call it a day. The real question for any company in this category is, where is the value actually being created? If it comes from the underlying software rather than from AI, the business may actually be more defensible against disruption. ERPs aren't going away anytime soon. But the value creation is fundamentally different from what AI is making possible.

The second category is AI-enabled – businesses that couldn't exist without AI. The clearest example is modern customer support, which has gone from onshore representatives to offshore to largely automated, and still has further to go. This category is genuinely differentiated and should be valued differently from AI-augmented businesses.

The third is AI-native. It's hard to appreciate the ocean of difference between companies founded before 2022 and those founded after until you see it up close. AI-native organizations are smaller,



more nimble, free of tech debt, and unburdened by the sunk cost fallacy or the innovator's dilemma. They expect more from each individual contributor because everyone uses AI heavily. Strategic and product changes that take months in an organization of thousands happen in weeks or even days in a small team. The AI-native version of any given software category will look profoundly different from its legacy competitor – even in the exact same domain – and will be more successful in the long run. Legacy incumbents have real advantages from the outset, but they will be slower and less empowered by AI than the AI-native businesses emerging to challenge them.

The strength of competitive positioning of AI-augmented, AI-enabled, and AI-native businesses can be evaluated with a handful of factors.

Of course, companies that weren't built AI-native can survive. It usually comes down to moats that go beyond software development. I use six factors to assess how well-positioned any company is:

1. **How hard is it to recreate the system?**
2. **Is there meaningful proprietary data?**
3. **Does the company own the action layer?**
4. **Is there a real-world execution element?**
5. **Are there network effects?**
6. **Who is the buyer?**

Different companies can live on an advantage in any one of these factors indefinitely. The analytical work is figuring out which factor a given company is actually relying on, and whether that factor is genuinely durable in a transformed, AI-driven world.

## ■ **The future of pricing: outcomes vs. consumption**

As the industry still explores alternatives to seat-based pricing, a few viable options are emerging: contracts that measure consumption, outcomes, or a combination of the two.

Outcome-based pricing is the theoretical gold standard: pricing against valuable, measurable results. Some companies are already doing it well, like Sierra and Decagon, and other customer experience platforms have built outcome-based models where the definition of success is clear enough to price against. Sales enablement tools pricing per qualified lead can work similarly well, as long as there's alignment on what "qualified" means.

But the model is hard to scale across the full range of enterprise software, because in many cases you can't get to a clean, unambiguous definition of the outcome you're measuring against. Professional services companies have long dealt with this problem through milestone-based pricing – and they've gotten used to tedious disputes over whether milestones are actually being met. The same dynamic will play out in enterprise software.

The more broadly applicable model ties consumption-based pricing to a metered metric of value creation. Not token count necessarily, but some version of a credit or flex unit that reflects consumption in a legible way. The challenge is building enough trust that the metric is consistent

and that vendors aren't manipulating consumption on the backend. That trust takes time to establish. I wouldn't be surprised if a better pricing scheme emerges in the next six to twelve months. The category moves fast enough that the current models should effectively be treated as provisional.

One underappreciated lever here is model routing. Not every workload requires frontier-level intelligence. Software companies that efficiently route tasks to the right model tier – using the most capable models for complex, high-stakes activities and cheaper smaller models for everything else – will end up with sustainable unit economics.

“The challenge is building enough trust that the metric is consistent and that vendors aren't manipulating consumption on the backend. That trust takes time to establish.”

### ■ The flip-side: unit economics in an agentic future

The biggest structural shift from the SaaS era to the agentic era is in the cost to serve. Traditional SaaS had negligible marginal cost of revenue: build the product, graduate from R&D to sales and marketing, scale hosting costs modestly, grow the top line. That model is changing.

In almost every case, the value being created by software today involves generative AI. And in almost every case, that AI comes from a frontier lab, not a proprietary model built by the software company. The software company's job is to use those models efficiently. As long as the cost of AI as a percentage of revenue is sustainably low, the business is viable. If it isn't, the economics unravel, regardless of how good the product is.

The good news is that the cost of AI and the cost to develop software has fallen dramatically, to as low as a hundredth of what it once was. The barrier to building a software business is lower than ever. I expect the market for software will grow exponentially as development costs drop. We will get to use new software products and services for a wide range of valuable things that were never possible before. But the margin structure of software businesses will look different, and investors who apply SaaS-era margin assumptions to AI-era software companies will make systematic mistakes.



**Mike Sabes**

Former Head of Corporate Finance & Strategy, Anthropic

■ Request engagement



# The Use Case Gap

By Tom Liu

Managing Partner, Ideate Capital

## We're headed toward an agentic future.

The main advantage that agentic systems have over legacy software is that they're not one-size-fits-all. There's no way for legacy software to be 100% personal and customizable, regardless of the customer's size. But agents are now making deep personalization possible.

Take two rival large healthcare companies, for example. They might do nearly the exact same thing and even with similar features of internal and external data, but they still require different use cases and requirements from their CRMs. Their data pipelines, infrastructure, and tech debt are extremely different and cannot be solved with legacy, out-of-the-box software. And even when companies try to extract value from off-the-shelf tools, the results depend heavily on the

users providing high-quality inputs, which makes it hard to get full value without a lot of work. Getting even 50 out of 5,000 employees to fully use a SaaS platform has historically been a real challenge.

Now they can connect agents instead. They probably can't fully replace their past software needs overnight, but they were already building incrementally on top of legacy systems. Agents let them go further in that direction toward personal, customizable, in-house solutions that actually address their use cases head-on.

## ■ Scalability, flexibility = more value

Companies and departments want to address specific use cases. Extracting that value with software tends to be clunky and people-dependent. With agents, even if the result isn't perfect, it can get you most of the way there, but you can aim perfectly at the right target. A scalable legacy system that hits 30-40% of each of seven use cases a company needs is a different proposition from a more limited agentic system that can hit 80-90% of each step because it was built precisely for that purpose. That's the core incremental advantage.

But we're not going to see a head-on software versus agents competition. Where the "SaaS-pocalypse" narrative has real purchase is on the peripheries – the incremental, bundled, upsell-driven revenue that SaaS providers have traditionally relied on. Customers now have more options – agents, in-house build, customization – and many are already choosing alternatives.

## ■ SaaS alignment to use cases will prevail

In the near term, you're not going to see teams completely turn off a software tool.

The real-world decision comes down to one question: how well is a given software tool actually aligned to the use cases that matter? Mar-tech and ad-tech come to mind. If you ran a real blind test, a lot of businesses who know their product well don't see meaningfully higher ROI from those tools. They're essentially paying for peace of mind, and any real value creation is coming from the domain expertise of the users, not from the platform itself.

That pattern shows up across departments. Even with CRM – if you're paying seven figures a year for a HubSpot or SalesLoft, sure there are features you'd miss, but would revenue collapse if you dropped it and cobbled together something in-house? Would you really lose all meaningful customer touch points when your GTM team is already editing and keeping localized versions of their pipeline? Those are questions legacy software had a hard time answering convincingly. And agents are the perfect catalyst to raise them.

“Where the “SaaS-pocalypse” narrative has real purchase is on the peripheries – the incremental, bundled, upsell-driven revenue that SaaS providers have traditionally relied on.”

## ■ Purpose-built is easier than ever to build

It's not winner-take-all. Some tools clearly quantify their value. But agents force the question: can I just build this myself? If I can get 90% of the way there, and in a way that's more personalized to how I actually work, at a fraction of the cost, the ROI payback becomes pretty clear.

Take a CRM we helped build for a bakery chain we've invested in that closely tracks customer trends and flavor preferences – twenty locations, fully digitized, everything running through kiosks. A year ago, what we built wouldn't have been possible. What we ended up with is probably as sophisticated as the CRM a casino would have used 10 years ago to track its top spenders, being able to interact with customers via live intent signals and taste preferences. We built version 1.0 in six weeks.

It's a vivid example of what's changing. A small business can now stand up enterprise-grade customer intelligence, at a fraction of the time and cost, with agents doing the operational work and humans focusing on the decisions.

## ■ To buy or to build?

- 1 Can it speed up solutioning?** The first consideration when deciding whether to buy or build is use-case alignment and iteration speed. Getting a real answer to a complex business question used to take six or seven rounds of follow-up with software providers or data teams across months – sometimes an entire year for CFOs dealing with large data environments. Now those iterations can happen in days. AI doesn't just answer questions faster; it helps users ask better questions, which can be even more valuable than getting a quick one-time answer.
- 2 What is the real cost?** Not just subscription spend, but attention, human labor, and resources. Why maintain an 80-person FP&A team when agents could support a team of five exceptional people? The cost calculation goes beyond just salaries versus SaaS licenses – you need to factor in human time and attention. How many meetings are now obsolete thanks to a five-person team sitting shoulder to shoulder versus the 80-person global team sending invites around all week during international hours? How much faster is the team moving now, even excluding automation and agentic capabilities?
- 3 How error-prone is it?** This is where the bull case for legacy software still has legs: You cannot run a major enterprise of 20,000 employees on something that works 80% of the time. Fortune 1000 companies aren't going to suddenly turn off enterprise software tools – they'll build on top of them until they hit a tipping point. But the threshold for acceptable error rates depends heavily on company size and complexity. A three-person FP&A team can tolerate an 85% accuracy rate and fix the rest manually. A large multinational cannot.

It's not either/or: Agents won't replace software, more likely they'll fill gaps in the use cases that software never fully addressed, running alongside it with time and attention gradually shifting toward the agent-native side.

## ■ How organizations adopt

The conventional wisdom is that security, data access, and permissioning are significant barriers to agent adoption. That's true for large enterprises, but it's much less true for SMBs and mid-market companies. A 30-person team doesn't run into the hard walls of complex permissioning and legacy tech debt. They can get a lot done with agents without hitting those constraints. And the new generation of startups – the 10-person firms generating \$100M in revenue that we're hearing about – doesn't need heavy legacy infrastructure at all.

This is where I'd push back on the view that incumbents are broadly protected by governance and security moats. That argument is most valid at the top of the enterprise pyramid. As you move down market, those moats thin out considerably. HubSpot has already flagged this publicly – SMBs are increasingly asking why they need to pay for software when agents can cover a lot of what they need.

**“The real friction points at large enterprises aren't fundamentally agent problems. They're human and organizational problems.”**

Data access and permissioning are often easier with agents than with legacy SaaS, not harder, at least for those who have the freedom to build the way they want. The real friction points at large enterprises aren't fundamentally agent problems. They're human and organizational problems: incumbent processes, tech debt, institutional inertia. Those take time regardless of how good the technology gets. And for companies that have the motto of “move fast and break things,” they now have an option beyond incumbent software providers and an answer that's satisfactory to outsiders.

“We're just going to use spreadsheets instead of using this software platform” wasn't acceptable for mid-market PE-owned portfolio companies before, but now those same owners applaud and cheer when they hear, “We're just going to build our own Agents instead of using this software platform.”

## ■ Pricing: seats lose, hybrid wins, outcome-based struggles

On the software side, we're seeing a strong push from vendors toward API and usage-based pricing in place of the seat model. It's a response to a double whammy, where customers need fewer seats while simultaneously pushing for lower per-seat prices. Vendors are responding by saying, “okay, your team is smaller, but they'll need more data – here's a credit-based or consumption-based model on top of that.” The Bloomberg and Capital IQ model is a good analogy: Rather than 30 seat logins, customers just take the API and pay per usage.

On the agent build side, the cost management story is more interesting than it might appear. Proof-of-concepts tend to run on premium commercial models, where you're sometimes running parallel agent teams as a check on each other, which doubles token costs. But once the MVP is fully working, you can swap commercial APIs for open-source models. DeepSeek, for example,

runs at a fraction of the cost on standard GPU infrastructure. So, it's a clever path: prove it's possible on commercial credits, then optimize costs with open source once the architecture is validated.

What hasn't worked is outcome-based pricing. Vendors will say, "if we hit these milestones, here's what you pay." The problem is that software, by its nature, can't customize deeply enough to reliably drive specific KPIs across 1,000 different clients with 1,000 different problems. That's exactly the pattern where AI outperforms off-the-shelf software – high customization needs and specific KPI improvement. So, the outcome-based pitch exposes the fundamental limitation of standardized software rather than solving it. Procurement teams are skeptical. I think vendors will have more success transitioning from seat-based to usage and credit-based pricing than they will making outcome-based models stick.

## ■ Valuation

The bull case for legacy software isn't wrong, but it's incomplete. It's right that large enterprises with complex use cases aren't ripping out compliance-grade systems just because agents can handle many of the less sophisticated tasks.

While Fortune 1000 companies aren't churning today, the trends we're seeing should still raise questions. What does the five-year roadmap look like? In three years, are you looking at a cigar butt or a business that has grown alongside AI?

Consider what happens as teams get leaner. Instead of 40 new hires on 40 seats, a company wants two seats and an API. Maybe they want no seats at all. What happens when SMB and mid-market customers make their own data tools through general-purpose AI?



## ■ What is the software provider worth in that world?

The question is nuanced and domain-specific – you have to go software by software, company by company, end market by end market. The variables to watch are already visible – and they're worth taking seriously before the market prices them in. Proprietary data, network effects, dynamic user-driven customizability beyond just chatbots – these are the key differentiators.

Hybrid SaaS companies that allow users to build and create value within the platform would harness the best of both worlds: scalability and personalization. On the other hand, incumbent providers that end up operating solely as a data allocator and governance tool are fated to become more of a nuisance for enterprise and likely a dead-on-arrival option for the SMB and mid-market.

Being able to build and embrace AI capabilities is going to be key in this world. Investors will probably turn their attention from traditional financial metrics – such as ARR, NDR, or Rule of 40 – to time spent per user, volume of data pulled, or the number of Agentic integration touch points that prove mission criticality as a network hub.



**Tom Liu**

Managing Partner, Ideate Capital

■ [Request engagement](#)



# AI is a Sorting Machine, not a Harbinger of Extinction

By Jill Sweeney

Former Strategy Director of High Performance Computing and Artificial Intelligence at Hewlett Packard Enterprise

**There's a popular narrative out there that AI represents the extinction of SaaS. I see that outlook as overblown.**

In reality, public market sentiment moves faster than enterprise behavior. While investors may imagine AI immediately displacing legacy software, in practice enterprise adoption follows a slower, more risk-managed curve. Buyers still prioritize security, integration, governance, procurement simplicity, and continuity of workflow. They are testing AI, but mostly within existing systems and contracts. So, there's a disconnect between valuation-level disruption and operational-level pragmatism.

“The key distinction isn’t SaaS versus non-SaaS. It’s shallow SaaS versus operational SaaS.”

I view AI as more of a sorting mechanism for the SaaS market, compressing value for software that’s interface-driven or reliant on the seat expansion model while strengthening platforms that are deeply embedded in workflows, controls, and governed enterprise data. The key distinction isn’t SaaS versus non-SaaS. It’s shallow SaaS versus operational SaaS. The vendors who can convert embedded

workflows, governed data, and distribution into AI-native operating leverage will win.

### ■ Why incumbents are maintaining an early lead, for now

Signals in adoption and budget trends suggest that enterprises are extending incumbent platforms with AI, defaulting to bolt-on AI solutions from incumbents rather than replacing core workflows with new vendors. Resets may happen at the margins, but most enterprise AI adoption is flowing through platforms that already own the identity, the workflow, and the data.

Incumbents own security governance and data gravity – real challenges for AI integration. A bolt-on from an existing vendor comes on top of a vetted stack with access to operational data. For procurement, it’s easier to extend a current contract or add a new module than to justify a new vendor, a new integration, and a new risk profile. And boards remain nervous about AI governance.

Niche point solutions come in later, around a gap. AI-native tools get adopted more selectively, where an incumbent’s AI is too generic, or where an acute use case justifies a separate system. So the sequence in an enterprise looks more like: turn on AI where I already pay – and then evaluate specialized tools next. That’s what’s happening in the early curve.

### ■ The structural advantages that create durability

More than inertia in the procurement process, the early incumbent advantage also reflects real structural moats.

The workflows customers have gotten used to are often more than just UI. They encode decisions and approvals, compliance permissions, and operational sequencing. AI is valuable only if it can operate inside those constraints. Legacy incumbents already control where AI can act, not just where it can analyze.

“Legacy incumbents already control where AI can act, not just where it can analyze.”

Governed historical enterprise data is also a critical choke point. Enterprise AI systems require trusted permission context, and that data contains not just information, but relationships, roles, and governance rules. AI only creates value when it can operate within that context. That’s extremely difficult for a new entrant to replicate because it’s fragmented, controlled, and deeply embedded in the system of record.

Independent AI tools may succeed at generating insight but most stall when it comes to enabling execution, running into hurdles around security, identity, data permissions, and, most importantly, operational accountability. The enterprise isn't just asking, "is this answer good enough?" They need to know who can act on the information and how it gets tracked.

As a result, durable platforms are separating themselves while the thin point tools and seat-based models are getting structurally repriced as agents replicate or bypass the value their products deliver. That's where AI represents a structural repricing event, not just incremental upside.

## ■ Managing the margin trap

If you're an incumbent vendor, bolting AI onto legacy systems can help you hold onto existing customers, but it introduces a new financial risk: Model inference costs can become a structural margin trap, especially if you can't monetize the value that AI creates. Incumbents should look at managing this through selective usage, creative model routing, and pricing strategies tied to workflow design. If AI increases retention and expansion, the total economic value can still improve even with a higher cost of goods sold.

As an investor, how do you differentiate an incumbent that's successfully folding AI into its moat from one whose AI features are just surface-level? You have to look at operational impact. Does AI materially change workflow completion? Is it improving automation rates and customer retention, or just driving spend expansion? If it's just an assistive UI feature, it won't show up in those metrics. Valuable AI changes how work gets done and how customers spend.

**“How do you differentiate an incumbent that's successfully folding AI into its moat from one whose AI features are just surface-level?”**

## ■ The orchestration layer as the new battleground

Major software platforms have started describing themselves as the operating systems for enterprise AI agents. In practice, that means becoming the control plane where agents get their permissions, context, execution rights, and workflow access. For buyers, this matters because agents need governance as well as intelligence. The platform that can orchestrate those agents becomes central to long-term value capture.

If enterprise users stop logging into front-end UIs because autonomous agents are bypassing them, the nature of lock-in shifts. Where it used to come from the interface and user behavior, it moves to the orchestration layer – permissions, workflow execution, and system integration. The platform that controls how work gets done wins in an agent-driven world.

As enterprise budgets adapt, the orchestration layer and workflow automation are the key areas for expansion. Shrinkage will happen where value is fragmented, i.e. point solutions and tools that offer interface intermediation, not deep operational ownership.

## ■ The bottom line

Ultimately, AI won't wipe out enterprise SaaS as a category. We should view today's trends as a consolidation rather than a wholesale reset. AI threatens weak SaaS products that were essentially just interface wrappers, while strengthening platforms already sitting inside critical workflows and systems of record.

Impressive AI demos and promises of low-friction solutions may earn initial success, but the winners in the long run will be the ones who turn intelligence into trusted execution rights inside the enterprise.



### **Jill Sweeney**

Former Strategy Director of High Performance Computing and Artificial Intelligence at Hewlett Packard Enterprise

■ **Request engagement**



# Core, Peripheral, and Everything at Stake

By Andreas Welsch

Former CMO of Business AI & VP of Product Management and GTM, SAP

**The term “SaaS-pocalypse” helps convey the magnitude of what’s happening to the software market, but it might be a bit extreme. A close look reveals important nuances that investors and business leaders should heed.**

To understand where the market is heading, you need to grasp the convergence of a few factors. One is the ability of organizations to rebuild their software, customize processes, and tailor them to their needs – at the cost of assuming new risks and liabilities they used to defer to vendors of standardized software. The other is estimating the growth trajectory for SaaS vendors, especially as AI is poised to change the composition of how work is done and the systems that support it.

## ■ Not all software carries the same risk

Assessing a SaaS company's exposure starts with understanding where its software sits in a customer's organization. Is it core to running the business, or is it a productivity tool that teams use to work more efficiently?

**“The survivability advantage these providers have boils down to three factors: risk delegation, convenience, and the potential to innovate.”**

If core aspects of a business depend on the app to function – to generate invoices, hire employees, or manage inventory – the company providing it is much more resilient to disintermediation. Think enterprise resource planning (ERP), customer relationship management (CRM), or supply chain software. The survivability advantage these providers have boils down to three factors: risk delegation, convenience, and the potential to innovate.

## ■ Why core SaaS is harder to displace than it looks

Embeddedness is still an important moat. Core SaaS applications can be so deeply embedded in a business's operations that replacing them becomes prohibitively cumbersome and costly. You may be able to save on license costs by switching to an AI or agentic alternative, but the full cost goes far beyond dropping the subscription. Managing the change – not just the technology change but the people change – adds significant financial complexity to the decision.

Then there's risk deferral. Replacing a core SaaS app with something AI-generated, vibe-coded, or fully agentic introduces new risks that IT leaders have traditionally outsourced to vendors who specialize in building such apps – vendors who know the market and the business function inside and out.

If there's an issue with the software – if it's unavailable or doesn't do what it's supposed to – you can call the vendor and appeal to contract guarantees around uptime. You have a support contract and a maintenance contract as part of the subscription. It's about deferring risk from your organization to a vendor who specializes in building this software and making it available in a secure and reliable way.

When organizations start building their own derivatives or in-house multi-agent systems, that responsibility shifts internally. The CIO or CTO becomes accountable for ensuring these systems work as intended – that there's no data leakage, that agents don't go off the rails, that guardrails are in place – on top of the ongoing build, testing, and maintenance burden.

Risk deferral matters especially in regulated industries, where auditability requirements are built into business systems. A third-party vendor can assume that regulatory risk – not only ensuring systems comply with current rules but adapting as new regulations are published. It's a deferral of risk, maintenance, and operational burden from the customer organization to the vendor.

A third factor is the potential for innovation. Vendors of established systems that bring genuine AI innovation to their customers can leverage years of built-up trust to expand their footprint – connecting agentic AI and emerging technologies to core business processes without requiring customers to rebuild everything from scratch or reduce their CRM, HRM, or supply chain system to a mere database. Companies that pay for a product expect performance, reliability, and outcomes better than they could achieve on their own.

“Vendors of established systems that bring genuine AI innovation to their customers can leverage years of built-up trust to expand their footprint.”

### ■ Where disruption is real: the peripheral software layer

Now let's look at the layer on top of core SaaS – peripheral software, made up of apps that aren't necessarily core to the business. Think functional enhancements or project management tools that help teams run workflows more quickly. Now that building an app in-house with AI assistance is increasingly viable, this layer – and the companies that provide it – face much greater risk of disintermediation and slowing growth.

Peripheral apps are simpler, both in how they're built and the workflows they enable. They typically don't require deep domain knowledge to replicate. For large organizations with experienced IT and development teams, using AI to build a comparable product – rather than paying millions in subscription fees – is a much easier lift now. Even if the savings are only a few multiples, the ROI calculation increasingly favors building.

“For buyers, it's a tradeoff: What does it cost to subscribe versus what does it cost to build?”

Peripheral software is also less critical. Interruptions have a much smaller impact on business operations, which means the cost of a homegrown solution not working perfectly right away is lower – and so is the value of the guarantees a legacy provider would have offered.

For buyers, it's a tradeoff: What does it cost to subscribe versus what does it cost to build? Not just the technology itself, but ongoing usage and the maintenance required to ensure compliance and data security.

### ■ The governance gap: organizational maturity determines the calculus

Some have argued the cost of building governance for every deployment flattens the distinction between core and peripheral apps. The truth is it depends on the maturity and size of the organization. It's easier for organizations that have already worked through predictive analytics, machine learning, and generative AI than for those just starting out. Large, multinational organizations with the capital, resources, and experience to build on agentic and generative AI will have an advantage here, whether they build, buy, or partner.

Smaller organizations face a real risk because the gap between what is and what could be is widening quickly. The question returns to basic strategic principles: What are our core competencies? Where should we spend our time and effort? And it's not just about building – you're on the hook to maintain it. These additional costs add up quickly for smaller and mid-size organizations, potentially making an in-house alternative prohibitively expensive. The same dynamic strengthens large incumbent vendors, who have the foundation and scale to deliver these capabilities even to their smaller customers.

## ■ How a buyer budgets and implements solutions of the future

With the market landscape in mind, leaders who decide to replace a peripheral application with an AI-driven alternative – whether built in-house or sourced from a startup – need a new evaluation rubric.

For purchased solutions, IT and procurement will continue to focus on data privacy, protection, and compliance. But now you need to add AI risk management frameworks to the conversation. Is the provider following global standards for AI risk management? Who owns the data? How is it segregated? These questions apply in addition to the usual concerns about a startup vendor's viability, longevity, and runway.

Corporate budgeting cycles also need to evolve. Even with peripheral tools, it's not as simple as swapping them out overnight – and the budget looks different.

Most platform-based offerings are moving to metered billing – like an AI lab charging based on units consumed, with a bill at the end of the month like a utility. Understanding and forecasting that cost, and managing token budgets, becomes critical. We've already seen early lessons from the token-maxing trend that emerged from Silicon Valley, with some companies burning through their 2026 AI budget in a few months because leaders encouraged employees to use AI as aggressively as possible. Companies like AWS, Microsoft, and Uber quickly found that AI costs were growing faster than anticipated and walked back the approach.

**“Most platform-based offerings are moving to metered billing – like an AI lab charging based on units consumed, with a bill at the end of the month like a utility.”**

Finance leaders will need time to establish what an ideal individual usage profile looks like at a company of a given size and industry. Right now, most of those benchmarks are narrowly based on individual company data.

Once costs are understood, leaders need a sound implementation plan.

An initial key to successful AI adoption is creating sandboxes and trial environments for internal users – including enterprise architects – so they can build, test, and experiment with company data in a safe setting. You also need a structured governance process that allows adoption of peripheral AI tools without creating unmanageable shadow AI sprawl.

Ultimately, leaders need to get hands-on. Vendors and consultancies tend to promote their products as low-touch, especially when speaking with C-suite leaders and investors. The reality is more friction and a steeper learning curve than advertised – not just learning the platform, but learning how to build a good agent that actually does what it's supposed to.

Getting hands-on doesn't mean becoming a developer. It means seeing firsthand how natural, spoken or written language becomes a final product – even a relatively simple one – and understanding where the process works, where it doesn't, and what data privacy and governance need to look like for agents to function reliably. Reading about this or talking about it on panels is no longer enough.

My recommendation is to start coding with AI. Try turning an idea into an application and go through that full lifecycle. There are things professional developers have internalized over years that only become visible when you try it yourself – where the technology performs well, where it falls short, and where your own proficiency is the limiting factor. Being hands-on is the best way to develop an intuition for where this is all going and how to make it genuinely useful. The time to do that is now.



### **Andreas Welsch**

Former CMO of Business AI & VP of Product Management and GTM, SAP

■ **Request engagement**



# The Architecture of SaaS Market Share Durability

By Robert Breynaert

Founder, Cloudscale Partners

**The “SaaS-pocalypse” poses many complex and interesting questions to investors, vendors, and buyers in the market. But all of them trace back to one phenomenon: AI is making the seat model obsolete.**

Understanding what replaces it – and which vendors survive the transition – requires a close look at two dynamics: the new pricing architectures emerging to fill the vacuum, and the structural forces that will determine which incumbents are genuinely protected and which are vulnerable to slowing growth or extinction.

# “The seat model was built on an assumption that’s now broken: that humans are the primary unit of software consumption.”

The seat model was built on an assumption that’s now broken: that humans are the primary unit of software consumption. For 20 years that was true. More employees meant more seats, which meant more revenue for the provider. The math was clean and predictable.

But AI agents decouple software value from human headcount entirely. When you move from seats to tasks, you’re moving from pricing access to pricing output, and the unit economics flip upside-down.

Three pricing models are replacing the seat model, and I recommend assessing each through three lenses: predictability, visibility, and revenue quality. How predictable is the spend? How much visibility does the person writing the check have? And what’s the revenue quality?

- 1 The first is the bundle model, like the Microsoft E7 approach, which scores high on all three.** You have Microsoft 365, AI built alongside it, Agent 365, and a governance layer, Entra, for identity and access management. They’re still anchoring to a seat as the billing unit, preserving annual recurring revenue, with clean RPO accumulation and the renewal cycle investors want to see. The consumption meter layered on top creates upside without introducing revenue recognition complexity. For a CFO it almost looks just like SaaS. The predictability and visibility are excellent.
- 2 The second is the Palantir outcome model.** Customers pay per business outcome instead of per seat. It’s a more defensible pricing conversation, but the challenge is making that outcome replicable and predictable across the customer base – business results are inherently variable, both within a single customer and across customers. The nuance many people miss is that Palantir doesn’t do pure outcome contracts. There’s a platform commitment floor – a minimum access fee – that provides baseline revenue. So, less predictability, but the gain is revenue quality. Charging against a business outcome embeds you in the customer’s operational P&L. It’s the stickiest possible commercial relationship, where you’re tied into the business outcome – and frankly, into the career of your customer’s CFO and business unit executive.
- 3 The third is the raw token model, like the OpenAI or Anthropic API approach.** It’s the worst of the three from a vendor perspective. There’s no RPO, and while it’s the simplest to understand, visibility beyond 60-90 days is limited. And you can’t build an RPO without a contractual commitment. Volume agreements help, but they leave you exposed to model commoditization – if a competitor releases a model at a fraction of your price, your customer will never meet their volume commitment and you’re in a negotiated true-up situation, which is bad for the relationship.

Of the three models, the bundle companies have the fastest adoption curve with Fortune 500 accounts today because they look most like traditional SaaS on paper – a hybrid that keeps the seat as a governance floor and layers a consumption meter on top, preserving budget

predictability while capturing usage upside. Salesforce and Microsoft are executing this well, using AI to drive premium upsell within existing contract structures. Enterprise buyers who have already committed to the platform are upgrading to the AI tier rather than going through new vendor evaluations. Wall Street favors renewals. That's the stickiness of incumbency.

The outcome model shows interesting signals, but in a narrower band. Palantir's enterprise commitment growth is impressive, but the accounts they deploy – financial services, government defense – are resource-intensive enough that Fortune 500 adoption is highly selective. The bundle model will win on breadth and speed; the outcome model wins on depth and durability, but in a much smaller customer base.



## ■ How does a vendor make the outcome-based model work?

There are a lot of theoretical frameworks for outcome-based pricing. The first crux is defining what an “outcome” actually is. You need to structure your outcome-based contract in a way that doesn't lead to attribution disputes – did the agent actually produce the business outcome? This could be a graveyard for inexperienced companies because it's easy for customers to disaggregate the software company from the decision they said they helped generate.

Strong attribution comes down to four conditions that must be satisfied before the contract is signed:

- **Use the customer's metrics.** You can't bring in your own – it will land like an alien language. If you're arguing with a CFO about whether your proprietary efficiency score is valid, you've already lost. If the customer already believes in their own metric, it's much easier to prove the impact you're having.
- **Establish a baseline before going live, using the customer's historical data.** A deployment bootcamp is fundamentally a baseline-setting exercise. You can't quantify impact without one.
- **Control for external variables.** If you're charging against logistics cost and fuel prices increased 50%, you didn't cause that. Good outcome contracts have adjustment mechanisms where agreed-upon external factors are explicitly excluded from the attribution calculation – similar to how investment managers handle benchmarking performance fees.
- **Keep a commitment floor.** Pure outcome contracts with no minimum are a recipe for disputes and lumpy revenue. A platform access fee – what Palantir does, and what Microsoft approximates with E7 – gives you a commitment floor to build on, with outcome-linked upside above it.

Pricing model dynamics, however, only tell part of the story. Equally important is understanding which vendors are structurally protected as the market shifts, and which are more exposed than the inertia thesis suggests.

## ■ Where incumbency is genuinely valuable

First, let's look at where incumbency is genuinely valuable: authentication and governance infrastructure. Consider what has to happen for an AI agent to do anything useful in an enterprise. It needs to authenticate somewhere – prove it has permission to read data or write to a system. Actions have to be logged and tracked so an auditor can reconstruct what happened. Someone has to be able to pull the plug. That requirement exists regardless of which AI product you're running or which SaaS vendor you use. Every agent from every vendor across every workflow needs to route through the identity and governance layer.

Microsoft's Entra ID is a concrete example. Microsoft spent 30 years building Active Directory and Entra, and it's now the identity backbone for enterprises worldwide – every employee, every device, every SaaS tool federating identity back to it. Now add AI agents into that environment. Each agent needs an identity, scope, and permissions. Unlike a human who logs in during working hours, an agent runs at 3:00 AM and can make 10,000 API calls before anyone notices something has gone wrong. Microsoft's move with Agent 365 and the E7 bundle is a direct extension of that Entra position into this new surface area. They're telling customers, "you already trust us to govern your human identities, now let us govern your agent identities with the same infrastructure." For an enterprise already all-in on Entra, moving to a different agentic governance platform would require a total rebuild.

## ■ Where the inertia thesis falls short

Many also still back the inertia thesis – that incumbents are protected by contract integration, structural inertia, and UX familiarity, creating prohibitively high switching costs.

This argument is valid but incomplete. It's correct on the what, but wrong on the where and when. Contract integration, UX familiarity, organizational muscle memory – these are real switching costs. People have built careers around Salesforce or Microsoft 365 administration, and that human network doesn't evaporate just because a better product arrives. Having sat through countless procurement cycles, mostly at Microsoft, I've seen first-hand that superior technologies lose to entrenched competitors all the time.

**“Blowing up the incumbent is not how enterprise transitions occur in real life.”**

What the inertia thesis gets wrong is it assumes disruption has to be a frontal attack. Blowing up the incumbent is not how enterprise transitions occur in real life. They happen at the edges – when a new division gets stood up, when an acquisition brings in a tech stack that gets evaluated at integration, or when a shadow organization brings in a new tool at renewal time.

Now, while outcome-driven architectures may be structurally confined to the edges today, structurally confined is not the same as permanently confined.

Core systems of record – SAP, Salesforce, Oracle – serve a function that has nothing to do with software efficiency. They transfer operational risk. When a CFO in a bank has SOX risk, they transfer it to SAP. When a hospital uses Epic for clinical documentation, they’re creating a legal record of patient care that governs liability under HIPAA. An outcome-based vendor can’t just walk in and say “we’re replacing your workflow.” They have to replace your risk. That’s a structural wall – not technical, but institutional and regulatory.

But that wall can erode if specific conditions change. If the regulatory barrier shifts, it can be weakened over time as AI systems develop their own audit trail history and start to look more credible to regulators. Nibbling at the edges, getting into specific business units, can lead to the core over a three-to-five-year period. We’re already seeing early versions of this in financial services, where AI-driven trading systems have established enough of a risk mitigation track record to displace existing systems – and taking on that risk comes with a pricing premium, increasing revenue and margins over time.

**“Outcome-based architectures will make inroads at the edge, then over three to five years begin to penetrate the core.”**

The key signal to watch is what I’d call process-agnostic regulation. Most core system protection derives from regulations that are process-prescriptive – they don’t just require a result, they require a specific documented process. The moment a regulatory body shifts from “you must follow this process” to “you must achieve this outcome with a documented audit trail,” the incumbent process system becomes vulnerable, and an outcome-based vendor like Palantir can come in and compete. That shift will happen incrementally, not

uniformly. Outcome-based architectures will make inroads at the edge, then over three to five years begin to penetrate the core. Some domains – healthcare, certain financial and government compliance functions – may never make that shift. But companies like Palantir are going to take a strong run at it.

### ■ **When the interface becomes a liability**

There’s another key point the inertia thesis misses: It tends to conflate the UI layer with the whole product. Yes, incumbents protected on their data and compliance core will be difficult to displace. But if it’s UI and workflow sitting on a common data warehouse that’s already being bypassed by aggressive AI programs, that incumbent is in trouble. And they’re often participating in their own disaggregation – selling the AI layer themselves, disintermediating their own UI to charge more for orchestration. The question is whether they execute that transition before a startup does it for them.

The compound interest analogy only holds as long as the knowledge built around the incumbent’s platform remains valuable. If an AI agent becomes the primary executor of a workflow, the

Salesforce admin’s expertise, for example, becomes less critical. When that expertise depreciates enough, the social infrastructure backing the incumbent weakens. We’re probably three to five years from that manifesting as a material factor.

In some cases, a familiar front-end interface can actually become a commercial liability when autonomous agents start executing workflows. Companies like Adobe and Salesforce built moats around rich, complex UIs. It took years to build, users had to be trained, and that familiarity created switching costs. But when the primary user of your software shifts from a human to an AI agent calling your API, UI becomes irrelevant. If an agent is executing a workflow in Salesforce – pulling contact data, updating pipeline stages – it’s hitting APIs directly. That UI Salesforce spent billions building goes out the window, not to mention investments in sales reps, trainings, and demos. Even worse, a complex UI now signals to an enterprise procurement team that the product’s entire value is built around human interaction. It feels like a legacy constraint.

**“The companies at risk are those who conflate interface complexity with structural defensibility.”**

It also reveals whether data is trapped behind the interface – not programmable, not open access, not extensible. The incumbents getting this right, like Salesforce and Microsoft 365, are exposing their underlying data through clean RESTful APIs and treating the UI as one option among many. The companies at risk are those who conflate interface complexity with structural defensibility.

Inertia may correctly describe the next year or two for core systems. It poorly describes the next five to ten. Holding these two timeframes separately is how you get an accurate view, rather than collapsing them into one that says incumbents are either safe or doomed.

## ■ A practical diagnostic for investors

If I were an institutional asset manager reallocating capital, I’d run five tests to assess a software vendor’s vulnerability. Score each one, weight by importance to your thesis, and you get your evaluation matrix:

- 1 Where does the data live?** If the vendor disappeared tomorrow, does the customer’s data stay with the customer or go with the vendor? If the underlying data lives in a customer-controlled warehouse – Snowflake, Databricks, an ERP like SAP – and the vendor is providing a UI and workflow layer on top, they’re extremely vulnerable. The customer can swap that out without losing data. The moat only exists if the vendor has proprietary schema and its own data – think Microsoft, Salesforce, Palantir – where migrating off is a multi-year project.
- 2 Is the core function regulatory or operational?** Ask whether the enterprise customer uses the product to satisfy a compliance or audit obligation – SOX reporting, HIPAA data handling, financial close attestation for the SEC. Anything where a regulator or auditor accepts the vendor’s output as authoritative. If so, switching isn’t just a software change: You need to replace an auditor-accepted process. High switching costs. If the vendor isn’t embedded in regulated workflows and everything they do is easy to rip and replace, they’re prone to displacement from an AI native.

- 3 **Can an API call replace the UI?** This is the most direct technical test. Does the vendor expose a complete, documented API that covers the same functions as the UI? If so, the UI is aesthetic preference, not structural dependency. Vendors who deliberately constrain their API relative to their UI aren't protecting themselves – they're advertising their exposure.
- 4 **What does non-recurring revenue look like and where is it coming from?** A vendor whose net revenue retention is driven by seat expansion is pricing against headcount growth, which is structurally declining. A vendor whose expansion revenue is driven by consumption and deeper workflow usage is pricing against something that grows with AI adoption rather than against it.
- 5 **What does the ecosystem look like?** Check developer activity, ISV ecosystem health, and growth in marketplace transactions. If the ecosystem is building deeper integrations that make the vendor more central – Agentforce being a good example – that's a reinforcing moat. If developers are building integrations that export data out into more open architectures, that's a warning sign: There isn't ecosystem lock-in.

## ■ The durability test: final verdict

The disruption is here – the question for investors is which vendors have built moats at the right layer. The evidence points to a clear hierarchy of durability. Vendors anchored in identity, governance, and regulatory infrastructure – Microsoft, SAP, Epic – hold structural positions that cannot be dislodged by a better product alone. Vendors executing the bundle model with consumption upside, like Microsoft and Salesforce, are winning in the near term by making AI adoption look like SaaS on paper. And outcome-based vendors like Palantir are establishing footholds at the edges of regulated industries that will compound into core penetration over a three-to-five-year horizon.

The vendors at risk are those whose moat was always the interface. When agents replace human users as the primary software consumer, UI complexity becomes a liability, not an asset.

The recommendation: allocate toward vendors who own the data layer, control the identity plane, or price against outcomes embedded in the customer's P&L. Avoid vendors whose defense is a switching cost story built on user familiarity. Run the five diagnostic tests before allocating capital, and hold the two timeframes separately: Inertia may correctly protect incumbents over the next one to two years, but not the next five to ten.



**Robert Breynaert**

Founder, Cloudscale Partners

■ **Request engagement**



# Trusted Autonomy

By David Betancourt, Ph.D.

Head of AI at DB Macro Research; Former Apple AI Scientist

**From an economic standpoint, SaaS's core vulnerability is that it has sold generic tools that require humans to perform their workflows.**

Historically, software pricing mapped to the number of human workers using the system. As the customer grows in headcount, so does the vendor's recurring revenue.

With autonomous AI agents, that predictable mapping breaks down. A small team assisted by agents can now execute the same volume of work that used to require many licensed users. Instead of the seat, the better unit of value becomes the workload, the action, the resolution, or the outcome delivered.

A simple example would be a monitoring and incident management software. In the old model, the software detects a problem, alerts a human, and then the human investigates and fixes it. But if AI agents can not only observe the incident, but also diagnose, escalate, roll back, or remediate it under governance, then the enterprise buyer starts caring less about how many people log into the tool and more about whether the system can complete the operational workflow safely. That's a structural challenge to seat-based SaaS. The unit of value shifts from human usage of a generic tool to trusted AI autonomy – governed, observable, end-to-end work.

Software isn't dead. But that seat-heavy, workflow-light SaaS that requires many UIs, dashboards, and human involvement will get drastically repriced when AI agents can complete accountable business workflows.

## ■ SaaS subcategories facing the highest risk

Value compression isn't distributed equally across SaaS. It falls hardest on seat-heavy, UI-heavy, workflow-light applications – and less so on systems that own data, permissions, auditability, and high-stakes execution.

The most exposed categories are those where the product is essentially a human-facing interface wrapped around repeatable workflows: customer support seats, sales engagement tools, basic productivity and collaboration layers, lightweight developer tools, approval flows, and dashboards that require humans to interact with them. Agents can either reduce the number of human users needed or bypass the UI altogether.

We should be more careful about writing off systems of record. CRMs, ERPs, identity platforms, security and compliance infrastructure, and core data platforms hold permissions, business logic, regulatory records, and workflow history in ways that aren't easily replicated. Some of those platforms may actually become more valuable if they evolve into the trusted control plane for agent activity.

The build-versus-buy calculus is also shifting. AI-assisted code generation makes it much easier for engineering teams to build fit-for-purpose internal tools that connect directly to their own data

and workflows without paying a SaaS premium for something generic. That accelerates erosion most in categories where the vendor's product is essentially a thin workflow layer or dashboard with limited differentiation.

**“Enterprises will build where the workload is strategic, highly customized, or significantly overpriced.”**

But internal development comes with its own costs, including around maintenance, security, compliance, governance, and long-term ownership. Enterprises will build where the workload is strategic, highly customized, or significantly overpriced. They'll continue

to buy where the vendor offers deep systems-of-record functionality, regulated workflow depth, enterprise-grade reliability, and trust. The pressure is intense and urgent for vendors whose

moat is UI templates and seat-based distribution. For vendors whose moat is data governance, compliance, and reliable execution at scale, their position remains strong.

## ■ How AI changes the enterprise P&L equation for buyers

From an enterprise corporate P&L perspective, the trust layer is not optimal compared to the cost savings of replacing human labor. In consumer use cases, an AI error may be annoying but reversible. In enterprise workflows, an agent can change customer records, payments, alter access rights, or affect production systems. That means reliability, permissions, observability, escalation, rollback, and human-in-the-loop controls need to be designed into the system. In P&L terms, that upfront investment can be significant. You need data infrastructure, integration, work evaluation, infrastructure monitoring, security review governance, and process redesign.

“Agentic consumption cuts across workflows, departments, tool calls, model usage, human reviews and exception handling that makes finance teams ask, what business outcome did this produce?”

Enterprise buyers still face procurement realities and budgeting constraints. They will pay for AI, but they won't accept unlimited costs just because vendors now have higher inference and infrastructure costs. Procurement teams already face pressures to reduce SaaS sprawl, rationalize overlapping tools, scrutinize seat utilization, and push back on renewal inflation.

The bigger issue is that AI costs are harder to budget than seats. While seat license maps cleanly to headcount, agentic consumption cuts across workflows, departments, tool calls, model usage, human reviews and exception handling that makes finance teams ask, what business outcome did this produce? What labor or risk did it remove, and why should we pay this vendor instead of using internal tools

or a cheaper model layer? So vendors can pass through AI costs only when they attach them to measurable trusted outcomes, they cannot simply repackage structural compute COGS as an enterprise premium. In my view, the buyer pays for trusted AI autonomy, governed, auditable, reliable work, not raw inference.

AI costs shouldn't only be compared to labor savings. They should be compared to the value of safely automating high-value workloads, faster customer resolution, fewer operational exceptions, lower error rates, better compliance and higher throughput. The companies that win won't simply replace people with agents. They will industrialize trusted autonomy, so the marginal cost of each additional workflow goes down over time.

## ■ Margin economics for vendors

With AI rendering the traditionally attractive margin economics of the SaaS business model untenable, vendors face a critical question: Can they price AI to perform work above its true

cost to serve? The winners will be vendors who manage model mix, routing, caching, smaller specialized models, evaluation loops, and workflow design, so that they deliver reliable outcomes without letting costs consume gross margins.

I would be cautious about making any generalization about the nascent market, especially at the large vendor level. We haven't seen a simple universal story where AI immediately destroys

**“We haven't seen a simple universal story where AI immediately destroys software margins.”**

software margins. Some vendors are offsetting AI costs through premium pricing, bundling usage limits, or attaching AI to a higher value enterprise contract.

But the latent risk is real, even where the average margin may look fine. While the AI feature-level economics are still unclear, the important question is not only corporate level non-GAAP gross margin – it's the gross margin of the AI workload itself. I expect the market to further separate vendors that can price trusted

AI work profitably from vendors that subsidize AI usage to defend legacy seat revenue. That distinction is central to the SaaS repricing thesis.

Ultimately, I don't think SaaS disappears. What we're seeing is the old seat-based, application-centric model being repriced. Traditional SaaS sold generic tools and monetized human usage. Agentic software will monetize trusted work – governed, observable, auditable workflows that produce accountable outcomes. That makes me a selective bear on seat-heavy, UI-heavy SaaS, but a bull on enterprise software that can become the control layer for trusted AI autonomy. The winners will own data permissions, governance, evaluation, observability, workflow depth, and the unit economics of reliable AI performed work.

Customers won't pay simply because an agent exists – they'll pay when the agent can safely complete meaningful work under governance.



**David Betancourt, Ph.D.**

Head of AI at DB Macro Research; Former Apple AI Scientist

■ [Request engagement](#)

## Tap into the world's leading platform for trusted human expertise

The perspectives in this report come from some of the most accomplished AI leaders in the world — and yet they represent a fraction of what GLG's network has to offer. Every day, GLG connects clients with senior operators who have deployed AI at enterprise scale, builders who understand where infrastructure holds and where it doesn't, and governance leaders who have navigated ethical and regulatory complexity in high-stakes environments.

GLG is ready to connect your team to leading AI experts, in the format that fits how you work, including one-on-one expert calls, structured roundtables, trainings and in-person events, research engagements, and advisory roles for key decision making and implementation projects. In addition to direct access to experts, GLG helps clients produce in-depth reports on trends impacting enterprise businesses today – such as in our previous report, [The Agentic AI Playbook](#).

■ [Learn more](#)