

The Digital Payroll

Moving AI Out of the IT Budget

July 2026

AUTHORS:



Steve Sarracino

Founder & Partner



Scott Watson

Vice President, Research



Emma Rowand

Research Analyst



About Activant

Activant is a research-led global investment firm that partners with high-growth companies. Since 2015, we have invested in category-defining businesses during their most critical phases of growth, partnering with founders who have won their initial battles and are ready for the next challenge.

Our approach pairs deep, proprietary research with patient, flexible capital and hands-on operational partnership. We work alongside founders and leadership teams to refine strategy, strengthen operations, and accelerate sustainable growth.

Activant Research is dedicated to uncovering the most exciting emerging technologies, sectors, and companies we believe will shape the future. Our research-driven perspective informs everything we do, helping us invest at meaningful inflection points and support founders in building enduring, category-leading businesses.

You can find out more about Activant and our research at <https://activantcapital.com/>.

Salesforce charges per conversation. Workday charges per workflow. Intercom charges per resolution. The new AI-native vendors want a cut of your P&L gains. Per-seat pricing is on its way out. **Software spend is no longer competing with other software. It's starting to compete with payroll.**


In the transition from Software as a Service (SaaS) to Service as Software (SaS), an annual charge of \$1,200 per human seat becomes \$10,000 per automated workflow as vendors start selling completed work over tool access.¹ In 2025, Goldman Sachs Research projected that agents would make up over 60% of the total software market by 2030.² What they didn't account for was what happens when the \$900 billion enterprise software market converges with the \$50 trillion global knowledge worker labor market.³

Pricing is something we've tracked closely at Activant over the years. Our 2024 research on [usage-based billing](#) mapped the infrastructure layer and shaped our conviction in Metronome, an Activant portfolio company recently acquired by Stripe. We've returned to pricing across several other pieces since then. In [The Next Wave of IT](#), we made the case for moving from seat-based to outcomes-based pricing. In [Voice Agents 2.0](#), we argued that usage-based pricing works best for developers getting started since it lowers the barrier to experimentation. Outcome-based pricing, by contrast, tends to work better for end-users because it ties cost directly to the value they receive. [AI Infrastructure: Compute \(4/4\)](#) showed how serverless inference has innovated pricing, charging per token or per second of compute rather than per GPU hour. And more recently, in [AR Operations Today](#), we tested buyer reactions to flat per-agent pricing.

Most vendors are still optimizing around the wrong unit of value, and buyers have already moved past them. PitchBook calls the risk here the "[SaaS J-Curve](#)," a period where declining seat revenue outpaces the ramp of new pricing models. Companies that pick the wrong pricing architecture can watch their economics erode quarter after quarter with no clear way to reverse course.

The key thing to remember is that protecting revenue and creating revenue are two very different games. Incumbents with deep lock-in will need pricing strategies that defend what they already have, while challengers will need inventive models to pry that revenue loose. For now, pricing against labor is the most compelling wedge, but it's not one-size-fits-all. The right approach depends heavily on who you are and what you sell. And to price correctly in the first place, you need to be able to define your value, which means we need the infrastructure to measure it.

What's Broken with the Current Models



Software pricing has always drifted in the same direction: closer to what the customer really values. Perpetual licenses charged for ownership, SaaS subscriptions charged for access, and usage-based models charged for consumption. Each era brought pricing closer to the unit of value, and each was accompanied by infrastructure companies that powered the transition. This is the fourth transition, and unlike the prior three, the industry hasn't converged on a model. The experiments are happening in real time, and they're failing in revealing ways.

Salesforce launched Agentforce at \$2 per conversation, hit customer pushback, and now sells three different pricing paths side by side, with 30% of partners calling the result "confusing" and Salesforce's own SEC filing conceding it has "limited experience with determining optimal pricing."^{4,5} ServiceNow rebuilt its entire tier structure in April 2026. SAP went from a 30% AI premium to freemium to a CEO op-ed announcing consumption pricing, all inside three years. Even the AI-natives have retreated: Cursor moved from flat seats to hybrid usage billing and Intercom cut Fin from \$2 to \$0.99 per resolution and then redefined the unit entirely. SAP's CFO captured the underlying dynamic: buyers "rave about consumption in theory" then ask for "a more predictable model."⁶ When Piper Sandler asked what enterprises prefer, pure consumption showed 16% and outcome-based 7%.⁷ The model vendors keep announcing is the one buyers keep rejecting.


Usage-based Billing: Essential Infrastructure, Wrong Buyer Interface

Usage-based metering is critical, but as infrastructure rather than interface. It's the layer that lets vendors track token consumption, inference costs, GPU seconds, and cost-to-serve per customer. We mapped this ecosystem in our 2024 research on [usage-based billing](#), covering the metering, billing, and revenue recognition

categories that now underpin every AI pricing model. [Metronome](#) (now [Stripe Billing](#)) built this layer for the largest AI companies in the world, and we expect it to remain a permanent, growing part of the AI stack. Every vendor needs usage data to understand its margins.

The problem starts when usage-based billing becomes the buyer-facing model. Four in five IT leaders report unexpected charges from consumption pricing, and 61% have cut projects due to unplanned cost increases.⁸ When CFOs can't predict a line item, they cap it and adoption stalls. The pattern repeats across the market. Salesforce ran through three pricing architectures for Agentforce in 18 months and Deutsche Bank found most customers "still trying to understand what constitutes a unit of consumption."⁹ Pure consumption pricing only works when the metered unit is observable and the buyer already budgets for it. Twilio is a good example where AI simply accelerates an existing per-minute model. Outside those conditions, this approach turns into a procurement fight.

The distinction matters for every AI vendor. Unless your product has a metered unit the buyer already understands and budgets for, usage data belongs in the back office, powering margin management and cost-to-serve intelligence rather than the buyer's invoice.



For many customers, usage-based pricing offers a compelling map of price to value. But if you've ever been on the other side of a 10,000-page AWS bill you know that understanding the value is difficult. To build trust in usage-based pricing you need to give your customers continuous access to their usage and spend data. To support this, your billing system needs to be a data system that users can query, not a black box."

Scott Woody

Co-Founder & CEO, Metronome (acquired by Stripe)

William Blair projects AI-era gross margins will settle at 60-70% versus the 80% SaaS benchmark, and Adobe's own management concedes AI's margin profile is structurally lower because inference costs scale with usage.^{10,11} Take Cursor for example, the AI-native poster child that abandoned its flat \$20 seats for hybrid usage billing once power users started eating frontier-model token costs. Vendors that don't instrument cost-to-serve before pricing will find out the hard way.

Outcome-Based Billing: Elegant in Theory, Messy in Practice

Outcome-based billing works best when the unit of value is unambiguous, measurable without negotiation, and tied to a workflow the vendor controls enough of to deliver.

In our [research on voice agents](#), we argued that cost-per-minute pricing misaligns incentives, rewarding duration over efficiency: a long, unproductive call drives cost without creating value. Outcome-based pricing fixes that. Intercom's Fin charges \$0.99 per resolution, billed only when the customer confirms the issue is solved.^{12,13} Intercom's head of pricing put it plainly: if Fin shrinks a 1,000-person support team to 200, a seat-based model loses revenue exactly when Fin is delivering the most value. In April 2026, HubSpot moved its Customer Agent and Prospecting Agent to outcome-based pricing: \$0.50 per resolved conversation and \$1.00 per qualified lead recommended for outreach.¹⁴

That said, outcome-based pricing introduces real complexity at the infrastructure layer. Metering platforms like [Metronome](#) and [m3ter](#) are built to process events independently, which is what lets them scale. Outcome-based models break that assumption because the result of an event is often only known well after the event itself, and distinct usage events need to be tied together to produce a billable outcome. It gets harder still when the lag between the triggering event and its outcome spans one or more billing periods. Chargebacks are a clear example: a payments vendor pricing on successfully refuted chargebacks may wait weeks to know whether a given dispute counts, and by then the original transaction sits in a closed period.

Outcome pricing struggles in two settings. The first is at the buyer level, where procurement prefers predictability. The second is at the vendor level, once usage scales or outcomes become hard to define. Intercom itself has shifted from "resolutions" to the broader idea of "outcomes" as Fin expands from service into sales and success, an implicit admission that the original unit was too narrow. Where the unit stays clear (a deflected ticket, a qualified lead, an automated case), outcome pricing lands. But paying only for results creates two new problems:

1. **Measurement.** Defining a "successful outcome" is inherently contentious. Outcome-based models introduce significant revenue variability as sales cycles turn consultative and R&D has to build in measurement capabilities.
2. **Accounting.** Under ASC 606, outcome-based revenue is recognized at a point in time upon verification, subject to variable consideration constraints and probability-of-reversal tests. In other words, finance ends up paying for the pricing team's simplicity.


HighRadius recently moved to outcome-based pricing, charging a variable fee tied to a fraction of the savings the client realizes. The model deliberately pushes back against ASC 606, which HighRadius argues is "not designed for innovation" when AI agents deliver instant, quantifiable outcomes rather than mere access.¹⁵

[Parloa](#) takes the opposite view, arguing outcome-based pricing is a principal-agent problem dressed up as alignment: when handle time drops from eight minutes to four, the efficiency gain flows to the vendor while the bill stays flat.¹⁶ They also flag that vendors bake attribution uncertainty into contracts as a permanent risk premium that lingers even after performance stabilizes.

So, when a product lacks a metered unit buyers already understand and budget for, and buyers want predictable costs, is a hybrid model the best of both worlds?

Hybrid Models: The Complexity Trap

Hybrid models are winning the buyer preference battle. Piper Sandler's Q1 2026 partner survey put hybrid at 37% and add-on subscriptions at 29%, followed by pure consumption (16%), higher tiers (12%), and outcome-based pricing (7%).¹⁷



The model works best when it gives buyers a predictable subscription floor with optional consumption upside they can *actually* budget for. It also works for AI-natives, like Cursor, whose unit economics break under flat pricing.

Hybrid makes sense for legacy vendors with large customer bases, who protect their revenue by bolting AI features onto their existing platforms. ServiceNow's tier-upgrade plus Assist Packs structure, Atlassian's seat-plus-credit overages, and Workday's ratable Flex Credits all fit this pattern. And all three are showing strong AI revenue traction without seat compression. RBC Capital Markets calls this the most viable defense for legacy vendors, and they may be right about the defense.¹⁸ The problem is what it does to the buyer.

In practice, hybrid models create multi-variable pricing frameworks that CFOs struggle to model. A mid-market company evaluating Salesforce Agentforce now faces at least four separate pricing variables, with total spend ranging from \$15,000 to \$50,000 per year before implementation. When procurement teams can't estimate cost in one sentence, they delay.¹⁹

KeyBanc captured the underlying buyer friction across the category: customers "have limited context on what a token or credit is worth and how to estimate their consumption."²⁰ That's why Piper Sandler's same survey showed 69% of partners saying AI investments are pushing out non-AI deals and 87% seeing sales cycles extend.²¹ Hybrid solves the philosophical tension between seats and consumption,

but only when the vendor can make the consumption unit legible and the overall package easier to buy than what came before.

The Margin Problem Beneath It All

Whichever model a vendor chooses, AI products carry economics that traditional SaaS never had to deal with. Inference, GPU compute, and token consumption introduce variable COGS that didn't exist in a per-seat world. Traditional SaaS gross margins run around 80%, AI-native products compress that to 50% to 70%, and the gap widens as agents take on more complex work. For every \$1 spent on model usage, enterprises typically incur \$5 to \$10 in hidden infrastructure costs to make those models production-ready.²² Gartner projects that by 2030, the fully loaded cost per AI resolution could exceed \$3, surpassing the cost of many offshore human agents.²³

The squeeze is happening on the demand side too. McKinsey data indicates enterprises use only 20% to 30% of the AI capabilities they've purchased, and Jefferies reports 57% of IT leaders fund AI by reallocating existing budgets rather than adding net-new capital.²⁴ Customers are cannibalizing other software spend to pay for AI, and they aren't fully using what they've already bought. The pricing model is broken at both the point of sale and the point of adoption.

Even the AI-natives have walked back the purest versions of their original pricing. This isn't a transition. It's an open search, and nobody has the answer yet. That's why the infrastructure underneath is just as important as the pricing model on top, laying the foundation for pricing logic. Whatever a vendor lands on, they need three things: visibility into cost-to-serve, accurate and transparent usage data to bill against, and a way to measure value delivered now that the work no longer maps to an employee's 9-to-5.

The Market and the Opportunity

The obvious play for AI-native challengers is to price against the labor they replace. Legora explicitly positions itself as targeting "tasks traditionally handled by junior associates," Dex charges a 20-30% success fee (replicating the executive search fee model), and Avoca captures value based on "gross job value booked."

The staffing industry is already showing the displacement. Randstad's Q4 2025 permanent placement revenue fell roughly 15% to €90 million.²⁵ Hays reported a 17% drop in permanent net fees for the fiscal year ending June 2025.²⁶ In March 2026, Adecco signed an unlimited Agentforce 360 license with Salesforce through 2027, aiming to power over 50% of its revenue with agentic AI by year-end.²⁷ The firms that profited from placing humans are now buying agents to replace them.

PitchBook models an 8x revenue uplift per unit as vendors shift from \$1,200 per seat to \$10,000 per workflow.²⁸ AI-native app spending grew 108% year-over-year.²⁹ Traditional BPO annual contract value dropped 14% to \$7.3 billion in 2025,

the lowest since 2020.³⁰ BCG projects agentic AI will unlock \$200 billion in net new value within five years.³¹ The BPO contraction and the agentic projection are the same story from opposite ends: spend is moving from human labor to AI-mediated work, at a volume large enough to remake enterprise software underneath it.

But labor-anchored pricing has a shelf life. The anchor breaks the moment agents do work no human could (400 fraud checks in 50 milliseconds) or companies run lean enough that there's no team to benchmark against.

If pricing infrastructure captures 1% to 3% of the billing layer in an enterprise software market, the companies building metering, billing, cost intelligence, pricing experimentation, revenue recognition, and agent-native financial infrastructure are positioned to capture that value no matter which buyer-facing pricing model wins.

Understanding the Pricing Stack

Every pricing model requires infrastructure that most AI vendors don't have in-house. They could build it, but most shouldn't. Every engineer assigned to billing, metering, or revenue recognition is an engineer not shipping the product customers actually buy.


Companies trying to monetize agents face three compounding challenges:

1. **Pricing model mismatch:** seat-based models punish companies whose agents shrink headcount, and usage-based models give CFOs bills they can't forecast.
2. **Invisible value:** agents run in the background, autonomously. Without a layer that tracks and surfaces what they're doing, customers can't see ROI and companies can't justify price increases or expansions.
3. **Margin opacity:** AI models, compute, and API costs vary by customer, feature, and usage pattern. Without real-time cost tracking at the customer and feature level, companies are flying blind on profitability.

Solving these means doing six things at once:

1. Metering what agents consume
2. Tracking the cost of every inference call
3. Billing across hybrid contract structures
4. Experimenting with pricing
5. Recognizing revenue compliantly
6. Proving the outcomes that justify the price


A growing number of companies are building tools for each. We mapped more than 50 of them across nine layers. What follows focuses on the layers that determine whether an AI vendor can price profitably and iterate fast enough to land on a model that works.



What the Buyer Sees

At the top of the stack sits the question every vendor needs to answer: **can you prove your product is worth the price?**

[Paid](#) builds pricing infrastructure for AI-native companies, enabling pricing models that align with how customers experience value from AI agents. Its value dashboards give real-time visibility into agent activity and performance, which matters as agents increasingly run in the background where the link between usage and ROI is otherwise hard to measure. Known customers include Artisan, Hyperbound, and Happy Robot and early adopters reportedly see 20-40% revenue increases within six months, alongside higher retention and faster sales cycles.³² Its margin intelligence tracks cost-per-agent profitability in real time, so the vendor knows whether a specific tier is making or losing money before the customer asks. Billing platforms can generate an invoice based on compute consumed. Paid is built for the harder question: **did the agent create enough value for the buyer to keep paying?**



A billing platform can tell you what an agent consumed. It cannot tell you whether the agent was worth paying for. That gap is the entire problem with monetizing autonomous work, because the agent runs in the background, and the buyer never sees it earn its keep. We built Paid to close that gap: real-time margin per agent for the vendor, and a value receipt the buyer can take to their CFO. You can't charge against a labor budget if you can't prove you replaced the labor."

Manny Medina
Co-Founder & CEO, Paid

Price Strategy

Experimentation: Beneath the buyer-facing layer, pricing experimentation determines how quickly a vendor can find the model that works. [Stigg](#) provides no-code plan configuration that lets vendors A/B test agent tiers and change pricing without engineering involvement. Its customer list includes [Miro](#), [Webflow](#), and [PagerDuty](#). [Togai](#) (now Zuora) and [Atlas](#) address the same problem with different approaches. When the unit of value is an agent persona with tiered capabilities, the ability to experiment with packaging and price points without shipping new code is the difference between iterating your way to product-market fit and burning a quarter waiting for an engineering cycle.

Configure Price Quote (CPQ): While most of the CPQ category is being retrofitted onto SaaS, [Nue.io](#) is the Salesforce-native CPQ platform purpose-built for complex agent deals – unifying quoting, billing, and usage metering in one system for vendors that sell through the Salesforce ecosystem.

The Vendors Engine

This is where the back-office infrastructure lives – the systems a vendor needs running continuously to price sustainably.

Billing and Monetization AI is inherently multi-model: a single contract might combine per-agent fees, usage overages, prepaid credit wallets, and enterprise discounts. [Stripe Billing](#), with [Metronome](#) integrated, now offers one of the most complete metering-to-billing pipelines available today, powering the billing stack for OpenAI, Anthropic, [Databricks](#), [Confluent](#), and NVIDIA. [Lago](#) is the leading open-source alternative, with adoption across developer-first AI companies that want to own their billing layer end-to-end.

Observability creates the link between usage and output: [Portkey](#) operates an AI gateway processing over one trillion tokens daily, with budget guardrails that enforce spending limits per agent or per customer tier – a control plane for vendors that need to protect margins in real time. [Braintrust](#) provides tracing and cost-quality monitoring used by [Notion](#), [Ramp](#), and [Cloudflare](#). [Arize](#) is the incumbent leader in AI observability, with a platform that spans evaluation, quality, and safety – the trust layer that underpins outcome-based pricing claims.

Cost intelligence and FinOps turns metering data into pricing decisions. [Vantage](#) integrates with more than 25 providers including Anthropic and OpenAI for token-level spend attribution. This produces cost-per-agent-tier data that determines whether a flat monthly fee is profitable. With organizations spending an average of \$85,521 per month on AI-native apps, [CloudZero](#) shows cost-to-serve per customer and per feature.³³ [Cast AI](#) sits one layer beneath, optimizing Kubernetes costs for AI workloads specifically, which matters as inference moves off shared infrastructure and onto dedicated GPU clusters.

Metering is the foundation that everything else depends on. [Orb](#) is one of the largest independent metering platforms, with a developer-first engine that lets vendors retroactively model pricing architectures before committing. Its customer list reads like the AI-native economy: [Vercel](#), [Replit](#), [Perplexity](#), [Pinecone](#). [Amberflo](#) is the established alternative, with strength in high-throughput event ingestion and credit burn-down management for enterprise AI companies.

[M3ter](#) takes a different angle, targeting the enterprise end of the market where the quote-to-cash stack is more complex and pricing changes must thread through legacy CPQ, ERP, and revenue systems. That positioning was validated in March when Salesforce picked m3ter as its advanced metering and rating partner for Agentforce Revenue Management, and in June announced a definitive agreement to acquire m3ter, folding m3ter's metering and rating capabilities natively into Agentforce.^{34,35}



Usage-based billing, which requires metering and rating data infrastructure, belongs in the back office when the buyer is human. In a Service as Software world the buyer is increasingly an agent, and before it transacts it needs to know its remaining credit balance in milliseconds, not minutes. That pulls the metering and rating engine into the interface layer, even if only through an API. Only the platforms that can do genuine sub-second rating at billions of events will be the ones the agent economy actually runs on.”

John Griffin

Founder & CRO, m3ter (acquired by Salesforce)

The cost intelligence layer ([CloudZero](#), [Flexprice](#)) and the metering layer (Stripe/Metronome, [M3ter](#)) are essential regardless of which buyer-facing pricing model wins.

Compliance and Enforcement

On the compliance side, the accounting complexity is real. The Financial Accounting Standards Board (FASB) issued an Accounting Standards Update in September 2025 to modernize software cost accounting, and platforms like [BillingPlatform](#) and [Maxio](#) specialize in ASC 606/IFRS 15 compliance – which becomes more demanding as enterprises reclassify AI agent costs from SG&A to COGS.³⁶

Revenue recognition determines whether creative pricing models survive the audit. Under ASC 606, variable pricing creates variable recognition rules. [RightRev](#) is one of the category leaders, automating compliance for consumption-based and outcome-driven models where traditional revenue recognition systems break down. [Leapfin](#) sits next to it, reconciling revenue data across billing, payments, and CRM into a single source of truth. This becomes essential when a single customer

contract spans per-agent fees, usage overages, and outcome-based bonuses, each with different recognition treatment.

Entitlements and Access Management enforce what was sold. When a vendor offers three agent tiers at different price points, something needs to control which features, usage limits, and capabilities each tier receives. [Schematic](#) builds a pricing and packaging control plane on top of Stripe that decouples entitlement logic from application code – a cleaner architecture than bundling it into billing. [Wingback](#) addresses the same problem with a different approach focused on real-time usage enforcement. Beyond these two, the layer is largely empty. Most entitlement functionality is still bundled into billing platforms rather than existing as purpose-built infrastructure. That gap is the clearest white space in the pricing stack.

With the infrastructure layer increasingly accounted for, the harder problem is the one sitting on top of it: choosing a pricing model that holds up with buyers and knowing when to change it.

A Framework for Navigating the Transition

We're not proposing a universal billing metric. The right model depends on the product, the buyer, and how much of the workflow the vendor truly controls. What every model needs, though, is the same underlying architecture: usage-based metering and cost intelligence operating behind the scenes to manage the vendor's margins, and a buyer-facing price that a CFO can explain in one sentence and finance can recognize cleanly under ASC 606.

The vendors winning this transition share a recognizable pattern. They've stopped treating pricing as a packaging decision made once at launch and started treating it as infrastructure that evolves as autonomy improves, costs fall, and the unit of value shifts. Four principles separate the companies building durable AI businesses from the ones stuck renegotiating every renewal. These aren't prescriptions for a single pricing model. They're the constraints any model must satisfy.

The Four Imperatives:

For AI-native challengers, pricing is the moat, and four imperatives drawn from our internal research define what that moat has to look like.

1. Price against the buyer's largest budget line, not their smallest

Software is now competing with payroll instead of competing with other software. Evercore ISI estimates a portion of \$60 trillion in global operating expenses is being reallocated into AI to generate permanent labor efficiencies.³⁷ The vendors capturing this are anchoring to the cost of the work, whether that's a junior associate's hourly rate, an executive search fee, or a fully loaded support agent. The anchor will eventually break once agents do work no human could do, but

until then, labor is the largest budget on the customer's P&L and the easiest reference point for a CFO to defend.

2. Make the unit of value legible in one sentence.

Pure consumption only works where the metered unit is something the buyer already budgets for. Outcome pricing only works where the outcome is unambiguous and the vendor controls enough of the workflow to deliver it. Hybrid only works when the consumption layer is genuinely optional rather than required to make the product usable. If procurement can't model annual spend without a usage forecast, the deal slows down regardless of how elegant the model looks on a pricing page.

3. Instrument cost-to-serve before you instrument the invoice.

The vendors who survive margin compression are the ones who built metering, observability, and cost intelligence into the back office before they needed it for billing. The buyer never sees this layer, but it's what determines whether a given price point is profitable, which customers are subsidizing which tiers, and how much room exists to discount in a competitive deal. Founders who defer this work until finance asks for it are heading for a rude awakening.

4. Build for the model to change.

Salesforce, ServiceNow, SAP, Intercom, and HubSpot have all repriced inside the last 24 months. The companies handling this well aren't the ones who guessed right on day one. They're the ones who can run a pricing experiment without an engineering cycle, migrate a customer from one architecture to another as autonomy improves, and recognize revenue compliantly across multiple contract structures at once. The pricing strategy will keep moving. The ability to move with it is the actual moat.

The vendors getting these four right are the ones worth backing through the transition, regardless of which pricing model they're running this quarter.

Pricing the Next Trillion

The industry doesn't need another debate about whether per-seat pricing is dying. It needs a framework for what replaces it, and the answer won't be a single model. Agent persona pricing anchored to labor economics is the best wedge we've found for AI-native challengers today. Incumbents defending installed bases will price differently than challengers prying revenue loose. What stays constant is the infrastructure underneath: metering in the back office, cost intelligence shaping margins, and a buyer-facing price a CFO can explain in one sentence.

The historical pattern holds. On-premise gave way when cloud was priced against data center costs, not server hardware. SaaS won when it was priced against total cost of ownership, not perpetual licenses. AI agents will be adopted at scale when

priced against the cost of human labor, not against the software they sit next to in a procurement review.

This is also where investors need to update their playbook. The SaaS-era rules don't translate cleanly. Gross margins will look worse, often by 10 to 20 points, because inference, compute, and tokens are real variable costs that didn't exist in a per-seat world. That isn't a sign the business is broken. An AI company running at 65% gross margin with healthy per-agent unit economics and a path to expanding wallet share inside a customer's payroll budget is a fundamentally different animal than a SaaS company at 65% margin fighting for the next seat. Reading the second as a degraded version of the first is how you miss the category. The right questions are what the company is pricing against, whether it can see its own cost-to-serve in real time, and whether the pricing model can move as autonomy improves. Those are the signals that matter now.

The consolidation has already started. Stripe bought Metronome. OpenAI bought Statsig. Kong bought OpenMeter. Fourteen acquisitions in roughly twelve months, concentrated in metering, pricing experimentation, observability and revenue recognition. The acquirers aren't buying experiments. They're buying the plumbing every AI company will need regardless of which buyer-facing model wins. Pricing infrastructure is the spine of the AI era. The companies that own each layer will define how the next trillion dollars of enterprise software gets priced, and the pricing model itself will keep changing as autonomy improves and costs fall. The pricing strategy is the variable. The infrastructure underneath is the constant.

-
- ¹ Pitchbook, [Through the Looking Glass: The Race to Build Enterprise AI](#), 2026
 - ² Goldman Sachs, [AI Agents to Boost Productivity and Size of Software Market](#), 2025
 - ³ Pitchbook, [Through the Looking Glass: The Race to Build Enterprise AI](#), 2026
 - ⁴ Jefferies Research, Partners See an In-Line Q, Stay Patient for 2H Reaccel, 2026
 - ⁵ Salesforce Inc, FY26 Annual Report
 - ⁶ Bernstein Pan European Strategic Decision Premium Review Conference, 2025
 - ⁷ PiperSandler Research, 1Q26 App Software Survey | Steady Demand; AI Investments Extending Non-AI Deals, 2026
 - ⁸ Zylo, [AI Is Driving a New Era of SaaS Cost Volatility](#), 2026
 - ⁹ Deutsche Bank Research, Salesforce: Takeaways from NYC Agentforce World Tour, 2026
 - ¹⁰ William Blair Research, The State of Agentic AI: The Four Key Questions Defining the Enterprise AI Race, 2026
 - ¹¹ Mizuho Securities USA, Where Do We Go From Here? Downgrade to Neutral, 2026
 - ¹² Fin AI, [Pricing](#), 2026
 - ¹³ Stripe, [Intercom on the evolution of value-based pricing for AI agents like Fin](#), 2025
 - ¹⁴ BTIG, Highlights From Partner Convos, Spring Spotlight, & Thoughts Heading Into 1Q, 2026
 - ¹⁵ HighRadius, [HighRadius launches \\$0 Implementation fee. \\$0 Subscription fee via Outcome Based Pricing for CFO software](#), 2026
 - ¹⁶ Forbes, [Outcome-Based Pricing: The Most Expensive Myth In Enterprise AI](#), 2026
 - ¹⁷ PiperSandler Research, 1Q26 App Software Survey | Steady Demand; AI Investments Extending Non-AI Deals, 2026
 - ¹⁸ RBC Capital Market Research, Software: Thoughts from European Marketing; Spring 2026, 2026
 - ¹⁹ Clientell AI, [Agentforce Pricing Explained: Real Costs & Hidden Fees](#), 2026
 - ²⁰ KeyBanc, Next Week in Software Volume CI—NOW and SAP Previews: Preparing for Divergent Tones, 2026
 - ²¹ PiperSandler Research, 1Q26 App Software Survey | Steady Demand; AI Investments Extending Non-AI Deals, 2026
 - ²² PYMNTS, [Businesses Adapt AI Budgets to Match Usage-Based Pricing](#), 2026
 - ²³ Gartner, [Gartner Predicts GenAI Cost Per Resolution for Customer Service Will Exceed Offshore Human Agent Costs by 2030](#), 2026
 - ²⁴ McKinsey Technology, [McKinsey Global Tech Agenda 2026 February 9, 2026 | Survey](#), 2026
 - ²⁵ Randstad, Q4 results: [Staffing Industry Analysts coverage](#), 2025
 - ²⁶ HAYS, [Annual Report & Accounts 2025](#), 2025

-
- ²⁷ The Adecco Group, [The Adecco Group to scale agentic AI at speed with unlimited Agentforce license agreement](#), 2026
- ²⁸ Pitchbook, [SaaS Is Dead, Long Live SaaS](#), 2026
- ²⁹ Zylo, [AI Is Driving a New Era of SaaS Cost Volatility](#), 2026
- ³⁰ ISG, [Index Insider: Will AI Help BPO Bounce Back in 2026?](#), 2026
- ³¹ BCG, [The \\$200 Billion Agentic AI Opportunity for Tech Service Providers](#), 2026
- ³² Paid, [Paid raises \\$21 million seed to help SaaS companies break free from the seat-based trap](#), 2025
- ³³ CloudZero, [AI Cost Management: How To Track, Allocate And Optimize AI Spend](#), 2026
- ³⁴ M3ter Blog, [m3ter Expands Integration With Salesforce to Support Advanced Usage Monetization in Agentforce Revenue Management](#), 2026
- ³⁵ Salesforce, [Salesforce Signs Definitive Agreement to Acquire m3ter](#), 2026
- ³⁶ FASB (Financial Accounting Standards Board), [FASB Issues Standard That Makes Targeted Improvements to Internal-Use Software Guidance](#), 2025
- ³⁷ Evercore ISI Research, Semis-on-Fire 8-May-2026: CPU Renaissance Thesis Update: ARM, INTC, AMD

Disclaimer: The information contained herein is provided for informational purposes only and should not be construed as investment advice. The opinions, views, forecasts, performance, estimates, etc. expressed herein are subject to change without notice. Certain statements contained herein reflect the subjective views and opinions of Activant. Past performance is not indicative of future results. No representation is made that any investment will or is likely to achieve its objectives. All investments involve risk and may result in loss. This newsletter does not constitute an offer to sell or a solicitation of an offer to buy any security. Activant does not provide tax or legal advice and you are encouraged to seek the advice of a tax or legal professional regarding your individual circumstances.

This content may not under any circumstances be relied upon when making a decision to invest in any fund or investment, including those managed by Activant. Certain information contained in here has been obtained from third-party sources, including from portfolio companies of funds managed by Activant. While taken from sources believed to be reliable, Activant has not independently verified such information and makes no representations about the current or enduring accuracy of the information or its appropriateness for a given situation.

Activant does not solicit or make its services available to the public. The content provided herein may include information regarding past and/or present portfolio companies or investments managed by Activant, its affiliates and/or personnel. References to specific companies are for illustrative purposes only and do not necessarily reflect Activant investments. It should not be assumed that investments made in the future will have similar characteristics. Please see "full list of investments" at activantcapital.com/companies/ for a full list of investments. Any portfolio companies discussed herein should not be assumed to have been profitable. Certain information herein constitutes "forward-looking statements." All forward-looking statements represent only the intent and belief of Activant as of the date such statements were made. None of Activant or any of its affiliates (i) assumes any responsibility for the accuracy and completeness of any forward-looking statements or (ii) undertakes any obligation to disseminate any updates or revisions to any forward-looking statement contained herein to reflect any change in their expectation with regard thereto or any change in events, conditions or circumstances on which any such statement is based. Due to various risks and uncertainties, actual events or results may differ materially from those reflected or contemplated in such forward-looking statements.