THE APPLICATION LAYER ISSUE 30 · WEEK 46 OF 2026

THE BIG READ

Three meters on the record: Salesforce prices an action at $0.10 and a resolution at $2.00, Anthropic puts a work suite inside a $20 seat, Gainsight drops the seat for an outcome

KEY SIGNALS THIS PERIOD

6	8	5	3
Verticals tracked this period	Vertical packages shipped	Incumbent SaaS responses	Startup signals

AUTHOR

Brian Letort

BrianLetort.AI

PUBLISHED

November 14, 2026

Issue 30 · Weekly read

WEB

brianletort.ai/industry/applica

The Application Layer archive

THE BIG READ

The thesis this issue defends.

Independent industry analysis. Compiled from public vendor announcements, SEC filings, earnings releases, and reputable trade press. The body sets the read for the rest of the brief.

Last week the multiplier stayed unpublished; this week three vendors put three different meters on the record, and the buyer's job changed from discovering the quantity to arbitraging the unit. Salesforce's Flex Credits Rate Card, in its revision dated November 3, 2026, converts the $500-per-100,000-credit list into unit prices anyone can read: 20 credits ($0.10) for a standard Agentforce action, 30 ($0.15) for a voice action, 400 ($2.00) for a Help Agent or IT/HR Employee Agent resolution and 400 ($2.00) for a sales 'lead worked', with prompt tiers at 2, 4 and 16 credits and Slackbot actions at 40 and 200. On November 10 the same vendor made AI Skills in Agentforce Coworker generally available on a $50-per-user-per-month seat (185,000-plus users and 21x usage growth in three months, both Salesforce's figures). One incumbent is now running a seat meter, an action meter and an outcome meter side by side on one price list. The twenty-actions-to-one-resolution break-even that the card implies has been worked through by Fidiora, aiagentsinfo.com and calculateforce.com (credited); what this issue adds is the cross-product illustration that a $50 Coworker seat costs the same as 25 Help Agent resolutions a month, and the negotiating rule: the action-versus-outcome ratio is a list-price ratio and says nothing about Salesforce's cost, but an agent that takes more than twenty actions to produce a resolution is cheaper to buy by the resolution, and the vendor's own card tells you where that line is.

The second regime is the seat that absorbs the suite. At a San Francisco event on November 10 Anthropic launched Claude Meetings (beta, Team and Enterprise, Google Meet first), Claude Sheets and Whiteboards (beta, all paid plans), Claude Data (formerly Dashboards, now generally available) and interactive connectors that open the full Adobe, Atlassian, HubSpot, Salesforce, NetSuite and Workday applications inside Claude, all inside the existing $20-per-seat-per-month annual Team and Enterprise price, with meeting notes and transcripts free in beta and in-call documents and whiteboards free until February 10, 2027. Enterprise administrators get per-connector action limits, granular spend limits, audit logs and a Compliance API; Claude Data turns on by default for Enterprise plans on November 20 while Sheets and Whiteboards stay off by default. Microsoft's shares fell between 2% and 3% the day of the launch, in the week of Anthropic's IPO roadshow; the issue records the coincidence, not a cause. The operating

consequence for a buyer is the inverse of Salesforce's: the price per seat did not move while the surface area under it roughly doubled, which is a lab pricing for share of the knowledge-work hour and deferring the meter (the Enterprise plan's usage at API rates, the per-connector action limits) to the administrator's console. A CIO comparing the two this week is comparing a $20 seat with an unmetered suite and a metered perimeter against a $50 seat with a revised per-action card, and the comparison only works at the level of what a resolution costs on each. The S-1/A No. 2's customer disclosures (grade 5 on the house's 1-5 evidence scale, where 5 is a primary filing) say what the seat is worth to the lab: run-rate revenue of about $81B at October 31, more than 11,900 customers above $100K generating 74% of nine-month revenue, more than half the Fortune 500 at $100K-plus, and Anthropic's revenue averaging under 0.1% of those Fortune 500 customers' total expenses. The last figure is the lab's own statement of headroom and the diffuse-return problem in one line: the largest enterprises in America pay it a rounding error. The Weekly owns the IPO range, the $26B placement and the commitments.

The third regime is the seat dropped altogether. Gainsight's chief executive published an open letter on November 12 refounding the company around retention as an outcome: Gainsight CS moves from seat pricing to value-based pricing by ARR and customers managed, every customer gets an Outcome Executive (a human for the largest accounts, an AI agent for the long tail) backed by forward-deployed engineers, MCP (Model Context Protocol) access is not charged, and the Atlas services unit sells gross- and net-retention improvement as a service with more than $75M of customer ARR under management in long-tail renewals (company-reported, October). Investing.com reported about 270 roles eliminated (grade 4; the company confirmed cuts without a number). monday.com's print on November 9 is the hybrid middle: AI ARR (annual recurring revenue) of $36M, agents past 6 million interactions in under two quarters and more than 30% of $50K-plus accounts buying credits, against revenue growth slowing to 19% and a fourth-quarter guide implying about 14%, with management saying the subscription-plus-credit shift is being introduced gradually and may take several quarters into 2027 to stabilise. Wix's Base44 is the version where the meter already pays: nearly $250M of ARR at a 58% non-GAAP gross margin (from near zero in the first quarter), Symphony's multi-agent SMB operator at ARR 'in the millions' three months after launch and inference costs under 30% of second-half bookings (both from the earnings call), a $10M GPU prepayment, and the full-year free-cash-flow margin trimmed to about 16% to fund Base44

marketing. Three incumbents, three pricing regimes, and the one that disclosed a margin on its agent product is the one that let the product cannibalise the core.

The context for all three is the week's adoption evidence, which pointed one way. The Federal Reserve's FEDS 2026-074 (a staff review of the existing evidence rather than a new measurement; grade 5 as a first-party document) finds AI adoption growing rapidly with most firms using it lightly or not at all, labour-market effects 'small and somewhat mixed' and aggregate productivity effects that 'appear to be small so far', with the recent productivity pickup predating generative AI. Genpact's survey of more than 2,000 executives has 79% with agents in production, 20% using them broadly, 15% with AI investment fully tied to P&L outcomes, 49% of initiatives funded in the past 24 months delivering expected value, and 92% having kept or grown headcount. Payscale finds 43% of employees saving three or more hours a week and 19% redirecting the time to higher-value work. PA Consulting's 2,000 UK workers have senior leaders using AI 6.9 times a day, about three times junior staff, and 32% seeing no meaningful benefit. Deming and co-authors' NBER paper has more than 80% of occupations above 20% task-level adoption and 2.8% of tasks above 50%, with person-level differences explaining 44% of adoption against 11% for the task itself. The Weekly carries the thesis consequence (the house's second hypothesis is restated this issue with the Fed's reading as its diffuse-leg ruler); the application-layer consequence is narrower and operational: the vendors have moved to meters because the seat was never where the return was, and the buyer who cannot fully attribute AI spend to a P&L line (85% of Genpact's sample) is the buyer the revised rate card was written for. The commerce layer added its own instance: OpenAI's Businesses in ChatGPT lets companies claim a profile, upload product feeds and deploy branded Business Agents that do not complete checkout and route shoppers to the merchant site, Shopify merchants verifying from their admin, which is the Uber-and-Shopify answer of W45 (price access, keep the checkout) arriving from the assistant's side of the counter. ChatGPT Enterprise's listing in AWS Marketplace, where a customer burns down an AWS commitment to buy a bundle of ChatGPT, Work, Codex and dots, is the same lab learning that the enterprise meter it can most easily attach to is a cloud commitment someone else already sold.

THE BAR

An application-layer move matters when it changes one of three things: which seat or surface a buyer pays for, which vertical's system-of-record is at procurement risk, or

which pricing model (seat / usage / outcome) governs the renewal. Moves that change none of these are noise.

VERTICAL MOVEMENTS

Vertical packages shipped.

8 packages this period across frontier labs, open weights, and the insurgent startup cohort. Each entry names the vertical, origin, and pricing posture — with the decision implication for buyers of the displaced tier.

2026-11-10

PACKAGE 01

Anthropic

OPERATIONS

FRONTIER LAB SEAT

Claude Meetings, Sheets, Whiteboards, Claude Data (GA) and interactive connectors

A knowledge-work suite ships inside the existing $20-per-seat annual Team and Enterprise price: Meetings (beta, Google Meet first), Sheets and Whiteboards (beta, all paid plans), Data GA, and connectors that open full Adobe, Atlassian, HubSpot, Salesforce, NetSuite and Workday apps in Claude; Microsoft shares fall 2-3% on the day

The seat price held while the surface roughly doubled, which is a lab pricing for share of the knowledge-work hour rather than for the feature. Notes and transcripts are free in beta and in-call artifacts are free until February 10, 2027, after which the free period's end is the first repricing date; Claude Data turns on by default for Enterprise on November 20 while Sheets and Whiteboards stay off by default, so the administrator's default settings decide what the seat consumes. The Enterprise controls shipped with it (per-connector action limits, granular spend limits, audit logs, Compliance API) are the metered perimeter under an unmetered seat; Agent Techniques Weekly carries the connector-action controls as a harness pattern. Fortune's figure of more than 9,000 apps available as connectors is Fortune's. For a buyer on Microsoft 365 the comparison is a $20 seat with Meet-first meeting capture against a Copilot add-on whose included tier gained unpublished caps last week.

SOURCE: ANTHROPIC (CLAUDE.COM); TECHCRUNCH; AXIOS; BUSINESS INSIDER

2026-11-10

PACKAGE 02

Salesforce

SUPPORT

INCUMBENT SAAS

HYBRID

AI Skills in Agentforce Coworker (GA) and the revised Flex Credits Rate Card

Coworker's reusable plain-language skills reach GA on a $50-per-user-per-month seat (185,000-plus users, 21x usage in three months, Salesforce's figures) while the public rate card prices a standard action at 20 credits ($0.10), a voice action at $0.15 and a Help Agent or Employee Agent resolution at 400 credits ($2.00)

The first incumbent to run a seat meter, an action meter and an outcome meter side by side on one price list, which converts last week's unpublished-meter problem into an arithmetic one: at $500 per 100,000 credits (no rollover) a resolution is twenty standard actions (the break-even Fidiora, aiagentsinfo.com and calculateforce.com have each worked through; credited), a Coworker seat is 25 resolutions a month (the house's cross-product illustration), and a sandbox action is 16 credits. Prompt tiers (basic 2, standard 4, advanced 16) and Slackbot actions (standard 40, advanced 200) are on the same card, as is 'leads worked' at 400 for the sales row. For a support buyer the negotiating number is the action-to-resolution ratio of the agent being bought: above twenty actions per resolution the outcome price is cheaper, below it the action price is, and the vendor has now told you where the line is. Teams, iOS and desktop apps and a Coworker MCP server are slated for GA in November. The 21x usage figure and the user count are vendor-reported; the card is a primary document.

SOURCE: SALESFORCE (BLOG; NEWSROOM; FLEX CREDITS RATE CARD PDF, UPDATED NOVEMBER 3, 2026; AGENTFORCE PRICING PAGE)

2026-11-12

PACKAGE 03

Gainsight

SUPPORT

INCUMBENT SAAS

OUTCOME

Gainsight CS value-based pricing, Outcome Executive and Atlas retention-as-a-service

Gainsight 'refounds' around outcomes: CS pricing moves from seats to a price set by ARR and customers managed, every account gets an Outcome Executive (human for the largest, AI agent for the long tail), MCP access is unmetered, and about 270 roles are cut (Investing.com; the company confirmed cuts without a number)

The clearest incumbent case this week of system-of-record economics replaced by outcome pricing plus agents, and the first in this series to drop the seat entirely rather than add a meter beside it. The Atlas services unit sells gross- and net-retention improvement as a service with more than $75M of customer ARR under management in long-tail renewals and more than 7,000 agentic interactions (October disclosure; company-reported), with forward-deployed engineers behind the agent. For a customer-success buyer the shift changes what is being bought from software access to a share of the retention outcome, which means the contract needs a baseline and an attribution rule before it needs a price; for a vendor it is the cost-side companion to last week's headcount pattern (ServiceNow, SUSE, Microsoft), paid for in roughly 270 roles. The headcount figure is grade 4 and does not enter the scorecard; the pricing change is the company's own.

SOURCE: GAINSIGHT (OPEN LETTER; PRICING ANNOUNCEMENT); INVESTING.COM

2026-11-09

PACKAGE 04

monday.com

OPERATIONS

INCUMBENT SAAS

HYBRID

monday.com AI agents and credits (hybrid subscription-plus-credit pricing)

AI ARR of $36M with agents past 6 million interactions in under two quarters and more than 30% of $50K-plus accounts buying credits, against revenue growth slowing to 19% and a Q4 guide implying about 14%; management says the hybrid shift is gradual and may take several quarters into 2027 to stabilise

The hybrid middle of the week's three regimes, and the one that put a timeline on the transition: revenue of $377.5M, net dollar retention of 108% overall and 113% for $100K-plus customers, 2,127 customers above $100K (up 33%) and a largest deployment above 87,000 seats, beside an AI line that is under 3% of annualised revenue and a growth rate slowing to 19% with a fourth-quarter guide implying about 14%. Management's own framing is that credit pricing is being introduced gradually with stabilisation 'several quarters' out, which for a buyer means the credit rate is negotiable now and will not be later. The $36M and the interaction count are vendor-reported; the growth and retention figures are from the results release. HubSpot's Phase 2 (W45) is the same structure one quarter further along.

SOURCE: MONDAY.COM INVESTOR RELATIONS (Q3 2026 RESULTS); Q3 CALL HIGHLIGHTS (YAHOO FINANCE)

2026-11-10

PACKAGE 05

Wix

ENGINEERING

INCUMBENT SAAS

HYBRID

Base44 (AI app builder) and Symphony (multi-agent SMB operator)

Base44 nears $250M ARR at a 58% non-GAAP gross margin (from near zero in Q1), Symphony's ARR is 'in the millions' three months after launch, inference runs under 30% of second-half bookings with a $10M GPU prepayment, and Wix trims full-year FCF margin to about 16% to fund Base44 marketing

The first agent product in this series to disclose a gross margin that cleared 50% inside three quarters, which is the datapoint every seat-to-meter vendor is implicitly promising and none had filed. The 8-K exhibit carries revenue of $570.2M (up 13%), bookings of $573M (up 11%) and a September Base44 cohort with 20% on annual plans against about 5% a year earlier; the Base44 margin and ARR are company-reported in the filing, and the Symphony ARR and the under-30% inference figure come from the earnings call. For an engineering buyer Base44 is the no-code end of the build mode (Agent Techniques Weekly's frame) priced as an application rather than as tokens, and the 58% margin at under 30% inference cost says the unit economics of selling built software are better than the unit economics of selling the tokens that build it. The trimmed free-cash-flow guide is the price of growing it; the core Wix business is what is being cannibalised, and Wix said so by funding the marketing.

SOURCE: SEC EDGAR (WIX 8-K EXHIBIT, Q3 2026 RESULTS); Q3 CALL TRANSCRIPT (MOTLEY FOOL)

2026-11-10

PACKAGE 06

OpenAI

COMMERCE

FRONTIER LAB

UNKNOWN

Businesses in ChatGPT (Business Hub, product feeds, Business Agents) and ChatGPT Enterprise in AWS Marketplace

Companies claim a ChatGPT profile, upload product feeds with pricing and availability, see 28-day mentions and click data, and deploy branded Business Agents (eligible US ecommerce and online services) that route shoppers to the merchant site rather than completing checkout; Shopify merchants verify from their admin; ChatGPT Enterprise becomes purchasable through AWS Marketplace private offers

The assistant's side of the W45 commerce answer: Uber and Shopify priced access and kept the checkout, and OpenAI's launch concedes the checkout (agents do not complete it) and prices the access through ads, since Business Agents can be the destination for ChatGPT Ads. For a merchant the product-feed upload (a merchant report puts the spec at 53 fields; not OpenAI primary) is the structured-catalogue investment Shopify's Catalog made for its merchants, now requested directly by the assistant, and the 1.2 billion weekly users OpenAI cites are the distribution being sold. The AWS Marketplace listing (custom pricing by private offer; same models and release timing as direct; ChatGPT, Work, Codex and dots bundled, Healthcare and Financial Services SKUs available) is a distribution move with a pricing consequence: an enterprise can pay for ChatGPT seats out of an AWS commitment it has already signed, which makes the cloud commitment the meter. Both are OpenAI's own announcements.

SOURCE: OPENAI (BUSINESSES IN CHATGPT; CHATGPT ENTERPRISE IN AWS MARKETPLACE); REUTERS

2026-11-11

PACKAGE 07

Rox

SALES STARTUP

UNKNOWN

Rox (Global 2000 sales agent)

$100M Series B at a $1.3B valuation co-led by General Catalyst and March Capital, with Sequoia, Factory, Databricks Ventures and Prosperity7; Rox reports ARR up more than 10x in its first year in market and 702% enterprise net revenue retention, with absolute ARR undisclosed

The sales row's best-funded challenger to the incumbent platforms, and a case study in what a startup discloses when the meter is working: a growth multiple and a retention percentage, not a base. TechCrunch's earlier report of about $8M of projected ARR at end-2025 is the only absolute figure in circulation and is not Rox's. The customer outcomes Rox cites ($330M of qualified pipeline in eight weeks for an IT services firm; $24M closed in five months for a medtech company; 24.3x ROI in 90 days for a German engineering firm; 4 million prospects engaged at 20% of traditional outreach cost) are company-reported customer claims and are labelled so; the named customers (MongoDB, Groq, WSP, Puma Energy, Ramp, New Relic, Trafigura) and the plan to grow the US go-to-market team fivefold in 2027 are the firmer facts. For a sales leader the comparison this week is Rox's outcome claims against Salesforce's published $2.00 per 'lead worked', which is the first time the row has had an incumbent unit price to put beside a startup's ROI multiple.

SOURCE: PR NEWSWIRE (ROX); FORTUNE; ROX (DATABRICKS VENTURES INVESTMENT)

2026-11-13

PACKAGE 08

Abridge; OpenEvidence

HEALTHCARE

STARTUP UNKNOWN

Abridge Assistant (local care pathways; free to community health centers) and OpenEvidence Access Intelligence (with ACP licensing)

Abridge's Assistant now cites each health system's own care pathways across more than 300 enterprise systems and 250M-plus patients and is free to 108,000 clinicians at community health centers serving about 52 million Americans (with NACHC); OpenEvidence licenses American College of Physicians and Annals content and launches Access Intelligence, free to all clinicians, covering 2,500-plus patient-support programs and guideline-grounded prior-authorisation letters

Both healthcare leaders chose free point-of-care distribution as the wedge ahead of monetisation in the same week, which is the healthcare row's version of the suite-inside-the-seat move: the clinician-facing surface is given away and the enterprise contract (the health system, the payer, the pharmaceutical programme) is where the price sits. Abridge's local-pathways feature is the enterprise hook (the Assistant cites the buyer's own protocols, which makes it harder to replace); OpenEvidence's 500,000-plus prior-authorisation documents already drafted on the platform and its licensed ACP content are the usage base and the authority it is monetising access to. All counts are company-reported. For a health-system buyer the practical change is that two vendors now compete to put guideline-grounded decision support in front of clinicians at no clinician-side cost, and the procurement question moves to data rights and the pharmaceutical-programme revenue model behind 'free'.

SOURCE: ABRIDGE (PRESS RELEASES, NOVEMBER 12-13); OPENEVIDENCE (ANNOUNCEMENTS, NOVEMBER 9 AND 11)

INCUMBENT RESPONSES

How the SaaS estate is answering.

5 moves from established SaaS vendors reacting to the agentic shift — product launches, repositioning, earnings color, partnerships. The system-of-record incumbents defending turf against systems of action.

2026-11-10

RESPONSE 01

Salesforce

AGENTFORCE COWORKER ($50 SEAT) AND THE FLEX CREDITS RATE CARD

AI Skills in Coworker GA on a $50-per-user-per-month seat; the November 3 rate card published with per-action (20 credits, $0.10), per-voice-action (30, $0.15), per-resolution (400, $2.00) and per-lead-worked (400, $2.00) lines against $500 per 100,000 credits

The incumbent answered last week's unpublished-meter pattern by publishing, and by running three meters at once. A buyer should now negotiate the action-to-resolution ratio rather than the credit price: the card fixes the ratio at twenty to one, so an agent that averages more actions per resolution than that is cheaper on the outcome line. The headless platform interaction multiplier this publication has tracked since W43 is listed on the card as 'TBA', so the line exists and the price does not; confirm it before signing cross-platform traffic. Pricing model defended: seat plus action plus outcome, all published (hybrid).

SOURCE: SALESFORCE (FLEX CREDITS RATE CARD PDF; AGENTFORCE PRICING PAGE; BLOG)

2026-11-12

RESPONSE 02

Gainsight

GAINSIGHT CS (VALUE-BASED PRICING), OUTCOME EXECUTIVE, ATLAS

Seat pricing dropped for a price set by ARR and customers managed; an Outcome Executive (human or AI agent) for every account; retention improvement sold as a service; about 270 roles cut (Investing.com, grade 4)

The first system-of-record incumbent in this series to abandon the seat rather than meter beside it, and to say in its own letter that the software category it led is being replaced by the outcome it was supposed to produce. For a customer-success buyer the contract now needs a baseline retention figure and an attribution rule, because the vendor is pricing on the delta. Pricing model defended: outcome (value-based by ARR and accounts), with agents and forward-deployed engineers as the delivery mechanism.

SOURCE: GAINSIGHT (OPEN LETTER, NOVEMBER 12); INVESTING.COM

2026-11-09

RESPONSE 03

monday.com

AI AGENTS AND CREDITS ON THE WORK-MANAGEMENT PLATFORM

AI ARR $36M; more than 30% of $50K-plus accounts buying credits; agents past 6 million interactions; revenue growth 19% with Q4 guided to about 14%; hybrid pricing 'gradual', stabilisation several quarters into 2027

The seat-plus-credit transition with a disclosed timeline and a disclosed cost in growth: the credit line is under 3% of annualised revenue while the core decelerates, and management has said the structure will not settle until 2027. A buyer renewing in the next two quarters is negotiating with a vendor that has told the market its credit pricing is provisional. Pricing model defended: subscription plus credits (hybrid), introduced gradually.

SOURCE: MONDAY.COM INVESTOR RELATIONS; Q3 CALL HIGHLIGHTS

2026-11-10

RESPONSE 04

Wix

BASE44 AND SYMPHONY

Base44 near $250M ARR at a 58% non-GAAP gross margin; Symphony ARR 'in the millions' after three months; inference under 30% of second-half bookings; $10M GPU prepayment; full-year FCF margin cut to about 16% to fund Base44 marketing

The incumbent that let the agent product cannibalise the core and filed the margin to prove it works: a 58% gross margin on an AI app builder inside three quarters is the number every hybrid-pricing vendor is implicitly promising. The trimmed free-cash-flow guide is the cost of pressing the advantage. Pricing model defended: application subscription with usage tiers (hybrid), priced as software rather than as tokens.

SOURCE: SEC EDGAR (WIX 8-K EXHIBIT); Q3 CALL TRANSCRIPT

2026-11-09

RESPONSE 05

OpenAI

CHATGPT ENTERPRISE IN AWS MARKETPLACE; BUSINESSES IN CHATGPT

ChatGPT Enterprise, Healthcare and Financial Services SKUs purchasable by AWS Marketplace private offer against existing AWS commitments; Business Hub and Business Agents open to eligible US ecommerce and online-service businesses with no agent checkout

Recorded as an incumbent response because the lab is behaving like one: it attached its enterprise seat to the one meter every large buyer already has (a cloud commitment) and it built a merchant-facing channel that concedes the checkout and sells the access, which is the position Shopify and Uber took from the other side in W45. The dots billing default of November 18 (W45) stands and is on the watchlist. Pricing model defended: seat with a pooled allowance, now billable through a hyperscaler's commitment (hybrid).

SOURCE: OPENAI (CHATGPT ENTERPRISE IN AWS MARKETPLACE); REUTERS

STARTUP SIGNALS

The insurgent vertical cohort.

3 startup signals this period — funding rounds, named customer wins, product launches, and M&A. The cohort between frontier labs moving down the stack and SaaS incumbents defending the record-of-truth.

2026-11-11

SIGNAL 01

Rox

SALES

$100M AT $1.3B

$100M Series B at $1.3B (General Catalyst, March Capital co-lead; Sequoia, Factory, Databricks Ventures, Prosperity7); ARR up more than 10x in year one and 702% enterprise net revenue retention, both company-reported, with no absolute ARR disclosed (TechCrunch's ~$8M end-2025 projection is the only base in circulation)

The sales row's challenger field gains its largest round of the year, and a disclosure pattern to note: multiples without bases. The customer ROI claims are the customers' as relayed by Rox and are labelled; the fivefold go-to-market expansion plan for 2027 is the clearest signal that the company intends to sell against Salesforce's published $2.00 'lead worked' price with outcome claims rather than a unit price of its own.

SOURCE: PR NEWSWIRE; FORTUNE; ROX

2026-11-13

SIGNAL 02

Abridge (with NACHC)

HEALTHCARE

FREE TO CLINICIANS (ENTERPRISE PRICING UNDISCLOSED)

Assistant free to 108,000 clinicians at community health centers serving about 52 million people; local care-pathway citation live across more than 300 enterprise health systems (company-reported)

The healthcare leader converts distribution into a moat by citing the buyer's own protocols and giving the clinician surface away where the payer is public; the enterprise contract carries the price. For the scorecard the row's leader widened its lead on distribution without disclosing a dollar of revenue.

SOURCE: ABRIDGE PRESS RELEASES (NOVEMBER 12-13)

2026-11-11

SIGNAL 03

OpenEvidence

HEALTHCARE

FREE TO CLINICIANS (LICENCE TERMS UNDISCLOSED)

Access Intelligence launched free for all clinicians (2,500-plus patient-support programs; guideline-grounded prior-authorisation letters; 500,000-plus prior-auth documents already drafted on the platform); licensing deal with the American College of Physicians and Annals of Internal Medicine (November 9)

The challenger to the healthcare row's leader is not Suki's ambient-documentation category but the point-of-care reference layer, which this week bought authority (ACP content) and distribution (free access) in the same week; the monetisation is the pharmaceutical programme and the payer, not the clinician. The row's challenger slot is reviewed below.

SOURCE: OPENEVIDENCE ANNOUNCEMENTS (NOVEMBER 9 AND 11)

PRICING SHIFTS

Seat to outcome, one move at a time.

2 public pricing-model shifts this period. The 'data owns the application' thesis predicts a structural move from seat-based to outcome-based pricing across SaaS; the rate of change is itself a market signal.

2026-11-12

SHIFT 01

Gainsight

Seat-based → Outcome-based

Gainsight CS moves from per-seat pricing to value-based pricing set by the customer's ARR and number of customers managed, with MCP access unmetered and retention improvement sold as a service through Atlas (more than $75M of customer ARR under management in long-tail renewals, company-reported). The first seat-to-outcome move by a system-of-record incumbent in this series; about 270 roles were cut in the same restructuring (Investing.com, grade 4).

SOURCE: GAINSIGHT (PRICING ANNOUNCEMENT; OPEN LETTER)

2026-11-09

SHIFT 02

monday.com

Seat-based → Hybrid

Hybrid subscription-plus-credit pricing being introduced gradually across the base, with more than 30% of $50K-plus accounts already buying AI credits and management guiding that the transition may take several quarters into 2027 to stabilise; AI ARR $36M against $377.5M of quarterly revenue. Recorded as a shift on management's own description of the structure and its timeline.

SOURCE: MONDAY.COM INVESTOR RELATIONS (Q3 2026 RESULTS); Q3 CALL HIGHLIGHTS

VERTICAL SCORECARD

Who leads each vertical.

Leader-vs-challenger by vertical, useful for procurement shortlists when matching workload to vendor cohort. As of Nov 14, 2026.

VERTICAL	LEADER	CHALLENGER	READ
Commerce	**Amazon Seller Assistant**	Sierra-led Poppy standard partners (Meta, Shopify, Stripe, Walmart)	Held, with the assistant's side of the W45 answer now on record: OpenAI's Business Agents do not complete checkout and route to the merchant, product feeds are uploaded by the merchant, and Shopify merchants verify from their admin, which is the access-priced, checkout-retained shape Uber and Shopify described, conceded by the assistant. Amazon did not update Seller Assistant in the window. The row changes when an assistant completes a checkout at scale or a merchant discloses agent-channel revenue; neither happened.
Engineering	**Codex (on GPT-6.1 Terra and Sol)**	Claude Code (on Opus 5.5 and Fable 5.2)	Held. The week's engineering events were harness defaults and a disclosed CI compromise (Agent Techniques Weekly carries Depi's chain through openai/codex, Claude Code 2.1.325's sandbox prompt and Codex 0.169's MCP lockout) and the model layer (Fable 5.2 at $8 / $40; The Model Pulse). The build-mode application entry is Wix's Base44 at nearly $250M ARR and a 58% gross margin, which is the strongest disclosed economics of any agent application in the series and sits in this row's no-code sub-category rather than against the two coding agents. The challenger gains Fable 5.2 as an option; the leader's model line is unchanged.
Healthcare	**Abridge**	OpenEvidence	Challenger changed from Suki to OpenEvidence under the row's rule (a challenger changes on two disclosed datapoints in consecutive weeks while the incumbent challenger has none): the ACP and Annals licensing deal (November 9) plus Access Intelligence free to all clinicians (November 11) this week, and the 500,000-plus prior-authorisation documents already drafted on the platform as the usage base; Suki had no disclosure in the window or the prior one. The leader widened its distribution (free to 108,000 community-health clinicians; local care-pathway citation across 300-plus systems). Both figures are company-reported; neither vendor disclosed revenue.
Support	**Sierra**	Salesforce Agentforce	Held, with the challenger's price list revised: $2.00 per Help Agent resolution and $0.10 per action on the November 3 Flex Credits card fixes the action-to-resolution ratio at twenty to one, and the $50 Coworker seat is 25 resolutions a month. HubSpot's $0.50 per

VERTICAL	LEADER	CHALLENGER	READ
			resolution (W44-W45) remains the low anchor and Sierra's outcome pricing remains unpublished, so the row has one published incumbent outcome price and one unpublished leader price. Gainsight's seat-to-outcome move enters the field as a customer-success adjacency and is recorded in the pricing table rather than against the row's leader.
Sales	**ZoomInfo**	Salesforce Agentforce with Docusign IAM	Held on the standing rule (a change needs a second datapoint). Rox's $100M Series B at $1.3B with a 702% enterprise net revenue retention claim and more than 10x ARR growth (company-reported, no base) is the first datapoint for a startup challenger; Salesforce's $2.00 per 'lead worked' on the published card is the incumbent's first unit price in this row; HubSpot (W45) carries over. Rox becomes the challenger on a second datapoint (a disclosed ARR base or a named customer's filed figure), the same two-datapoint rule the healthcare row applied this week.
Security	**Codex Security Cloud**	AWS Continuum	Held. The week's security events were harness-side and are in Agent Techniques Weekly: Depi's prompt-injection chain to Codex's npm publish token, two MCP Python SDK advisories, Microsoft's agent-governance controls and Kiro Enterprise Management. None is a packaged security product; all raise the cost of not having one. No disclosure from either row holder in the window.
Legal	**Harvey**	Legora	Held; no legal package or customer disclosure in the window. The nearest adjacent event is OpenEvidence's guideline-grounded prior-authorisation letters, a document-drafting workflow in healthcare rather than law.
Research	**Claude Data (formerly Dashboards)**	Gemini agent	Leader renamed, not changed: Claude Dashboards reached general availability as Claude Data on November 10 and turns on by default for Enterprise plans on November 20, now alongside Sheets, Whiteboards and Meetings inside the same seat. OpenAI's dots (W45) begin billing by default on November 18 and remain the persistent-research entry without a deployment figure. The Gemini agent had no disclosure in the window; Gemini 4 Carbon's GA and Hydrogen's launch are model events in The Model Pulse.

VERTICAL	LEADER	CHALLENGER	READ
Finance	SAP Joule	Workday agentic finance	Held. OpenAI's ChatGPT for Financial Services SKU became purchasable through AWS Marketplace, a distribution event rather than a finance-function package; Anthropic's interactive connectors now open NetSuite and Workday inside Claude, which places the frontier lab's seat on top of both finance incumbents' systems of record without displacing either. No disclosure from SAP or Workday in the window.

ARCHITECTURE WATCH

Patterns to track.

5 cross-vendor patterns this period. Each pattern names the trend, the exemplar moves, and what it changes for procurement, cost, or vendor-risk posture.

PATTERN 01

Three meters on one price list: seat, action and outcome now coexist in public, and the ratio between them is the negotiable number

`Salesforce (Coworker $50 seat; $0.10 action; $2.00 resolution; $2.00 lead worked)`

`Gainsight (seat dropped; price by ARR and customers managed)`

`monday.com (subscription plus credits, 'gradual', stabilising in 2027)`

`Anthropic (seat held at $20 with the suite inside and a metered connector perimeter)`

Last week's pattern was the meter the vendor would not publish; this week three vendors put three different meters on the record, and the practical consequence is that the buyer can compute cost per resolution across regimes from primary documents. Salesforce's card makes a resolution twenty standard actions (a break-even Fidiora, aiagentsinfo.com and calculateforce.com have each published; credited) and, on the house's cross-product illustration, a seat twenty-five resolutions a month; Gainsight prices the outcome directly and charges nothing for the MCP access that lets other agents reach it; monday.com sells the seat and meters the credit with a stated timeline for settling the rate; Anthropic holds the seat and pushes the meter to the administrator's connector limits. What a buyer does differently is stop negotiating the credit price and start negotiating the ratio: for any agent whose actions-per-resolution exceeds the vendor's implied ratio, the outcome line is cheaper, and for any suite-inside-the-seat product the administrator's default settings (Claude Data on by default November 20; Sheets and Whiteboards off) are the real price control. The vendor's interest, as W45 said, is in the buyer consuming before the buyer prices; the published card is the vendor conceding that buyers now ask.

SOURCE: SALESFORCE FLEX CREDITS RATE CARD; GAINSIGHT PRICING ANNOUNCEMENT; MONDAY.COM Q3 CALL; ANTHROPIC TEAM AND ENTERPRISE PRICING

PATTERN 02

The suite inside the seat: a frontier lab bundles office surfaces into a flat seat and prices for share of the hour

- Anthropic (Meetings, Sheets, Whiteboards, Data GA, interactive connectors at $20 per seat; Microsoft -2-3% on the day)
- Interactive connectors opening full Adobe, Atlassian, HubSpot, Salesforce, NetSuite and Workday apps inside Claude
- Free in-call artifacts until February 10, 2027 as the deferred repricing date

The move is the inverse of the metered incumbent's: hold the price, double the surface, and let the connector perimeter (per-connector action limits, spend limits, audit logs, a Compliance API) carry the metering that the seat does not. The connectors are the strategic part, because a seat that opens the full Salesforce, Workday and NetSuite applications inside the lab's surface puts the lab between the user and every system of record the incumbents sell, without displacing any of them, which is the position Microsoft's Copilot was built to hold, and the position its shares fell 2-3% on the day the announcement landed, during the lab's own roadshow (coincidence recorded, not cause). For an architect the pattern says the knowledge-work seat is becoming the harness (Agent Techniques Weekly's frame) and the systems of record are becoming tools under it; for a CFO the February 10 end of the free period and the Enterprise plan's usage-at-API-rates clause are where the meter will appear. The S-1/A's statement that Anthropic's revenue averages under 0.1% of its Fortune 500 customers' expenses is the lab's own measure of how much of the hour is left to take.

SOURCE: ANTHROPIC (CLAUDE.COM, NOVEMBER 10); BUSINESS INSIDER; SEC EDGAR (S-1/A NO. 2)

PATTERN 03

Free at the point of care: healthcare vertical AI gives the clinician surface away and prices the enterprise, the payer and the programme

- Abridge Assistant free to 108,000 community-health clinicians (NACHC); local care-pathway citation across 300-plus systems
- OpenEvidence Access Intelligence free to all clinicians; 2,500-plus patient-support programs; ACP and Annals content licensed
- 500,000-plus prior-authorisation documents drafted on OpenEvidence (company-reported)

Both healthcare leaders made the same move in the same week from different categories, and the shape is the vertical's version of the suite-inside-the-seat: the user-facing product is free, the lock-in is the buyer's own content (Abridge citing a health system's protocols; OpenEvidence licensing the College's guidance), and the revenue sits one layer up with the health system, the payer and the pharmaceutical patient-support programme. For a health-system buyer the procurement question moves from licence price to data rights and from the clinician's seat to the programme's economics behind 'free'; for the scorecard the row's challenger changes to OpenEvidence because the point-of-care reference layer, not ambient documentation, is where the second datapoint landed. The adoption figures are the vendors'; the distribution agreements (NACHC, ACP) are public.

SOURCE: ABRIDGE PRESS RELEASES; OPENEVIDENCE ANNOUNCEMENTS

PATTERN 04

The assistant concedes the checkout and sells the access: agent commerce arrives from the assistant's side in the same shape the marketplaces chose

- OpenAI Business Agents (no checkout; route to merchant site; ads destination)
- Merchant product feeds uploaded to the Business Hub (a merchant report puts the spec at 53 fields)
- Shopify merchant verification from the Shopify admin
- Uber and Shopify pricing access, keeping checkout (W45)

W45 had two marketplaces answering disintermediation by pricing access and keeping the transaction; W46 has the largest assistant conceding the transaction and pricing the access through ads and distribution (1.2 billion weekly users, OpenAI's figure). The two sides agree on the shape, which makes agent commerce for now a referral and merchandising channel priced like one rather than a checkout to be defended; the structured catalogue is the asset on both sides, requested by the assistant and built by the marketplace. For a merchant the operating decision is to treat the product feed as a distribution investment with a measurable 28-day return in the Business Hub, and to watch whether 'agents do not complete checkout' survives the first quarter in which the ads destination earns more than the referral. The Poppy standard remains the attempt to make the access a protocol.

SOURCE: OPENAI (BUSINESSES IN CHATGPT); REUTERS

PATTERN 05

Adoption is measured in usage and return is measured in the P&L, and the central bank's first reading of the gap is 'small so far'

- Federal Reserve FEDS 2026-074 (aggregate productivity effects 'small so far'; labour effects 'small and somewhat mixed')
- Genpact (79% agents in production; 15% fully tied to P&L; 49% of funded initiatives delivering expected value)
- Payscale (43% save 3+ hours a week; 19% redirect it to higher-value work)
- PA Consulting (senior leaders 6.9 uses a day, ~3x junior; 32% see no meaningful benefit; 37% spend as long reviewing as producing)
- NBER / Deming (80%+ of occupations above 20% task adoption; 2.8% of tasks above 50%; person explains 44%, task 11%)

W45's version of this pattern was one buyer (Goldman Sachs) narrowing a seat and one central bank (the Bank of England) putting realised productivity at 0.5% over three years; W46's is the Federal Reserve's staff saying the aggregate effect is small so far, four surveys agreeing that use is broad and shallow, and a frontier lab's own prospectus saying its revenue is under 0.1% of its largest customers' expenses. None of this contradicts the vendors' usage metrics (Coworker's 21x, monday.com's 6 million interactions, Rox's 702%), because usage and return are different lines and the week's evidence is that they diverge for longer than a procurement cycle. For an operator the task is narrow: pick the one resolution-type your team produces most, price it on each vendor's card, and report that number beside the usage metrics the vendors publish. The wider conclusion is the one Genpact's 15% implies: an AI budget that cannot be fully attributed to a P&L line is the budget the published rate cards were written to capture, and the vendor that publishes a per-resolution price has handed the buyer the unit to tie it with. The Weekly carries the thesis consequence (the second hypothesis is restated this issue with the Fed's reading as its ruler); the application-layer consequence is that the seat was never where the return was, and the vendors moved first.

SOURCE: FEDERAL RESERVE (FEDS 2026-074); GENPACT; PAYSCALE; PA CONSULTING; HBS DIGITAL DATA DESIGN INSTITUTE (NBER PAPER)

WATCHLIST

On the radar next.

7 catalysts in the next 7–30 days that would change the read materially — earnings prints, conferences, expected launches, regulatory decisions, and competitive responses.

NOVEMBER 18

WATCH 01

OpenAI dots billing begins by default in Enterprise, Edu and Healthcare workspaces (carried from W45)

The end of the admin preview starts a persistent agent drawing on Enterprise allowances for every member whose admin did not restrict access; with ChatGPT Enterprise now sold through AWS Marketplace, the first month's dots usage is also the first test of a cloud commitment being drawn down by an agent the buyer did not open.

NOVEMBER 18-19

WATCH 02

Anthropic prices; the 424(b) prospectus may update the customer metrics

The S-1/A No. 2 figures (run-rate ~$81B at October 31; 11,900-plus $100K customers; under 0.1% of Fortune 500 customers' expenses) are the baseline for the suite-inside-the-seat read; any pricing-day update to the customer mix or the consumption share is the first post-launch datapoint. The Weekly owns the IPO itself.

NOVEMBER 20

WATCH 03

Claude Data turns on by default for Enterprise plans; Coworker Teams, iOS, desktop and MCP server GA slated for November

The default flip is the administrator-side price control for the unmetered seat (Sheets and Whiteboards stay off); Salesforce's Coworker surfaces reaching GA puts the $50 seat on the same channels the $20 one just took.

DECEMBER AND Q4 PRINTS (FEBRUARY)

WATCH 04

monday.com hybrid pricing rate and HubSpot EMEA Phase 2 cohort

Both vendors have said the seat-plus-credit structure is provisional; the figures to look for are credit consumption per customer in currency and whether net revenue retention (108% at monday.com; 101% at HubSpot) moves on the agent cohort.

BEFORE CROSS-PLATFORM TRAFFIC IS METERED

WATCH 05

Salesforce Headless Platform Interaction multiplier on the Flex Credits card

Carried from W43-W45: the November 3 card lists the headless platform interaction multiplier as 'TBA', so cross-platform agent traffic has a line and no price; the next card revision is the place to see at what multiple of the 20-credit action.

Q4 2026 AND Q1 2027

WATCH 06

Gainsight's first value-based renewals and Atlas retention figures

The seat-to-outcome move is the first by a system-of-record incumbent in this series; a disclosed retention delta under the new contracts, or a customer naming its attribution rule, is the datapoint that turns the pricing model from an announcement into a result.

FEBRUARY 10, 2027

WATCH 07

Claude Meetings in-call documents and whiteboards free period ends

The deferred repricing date for the suite inside the seat: whether Anthropic meters the artifacts, raises the seat or keeps them free is the first price move the $20 seat will have made since the suite doubled under it.

METHODOLOGY AND AUTHORSHIP

How this brief is built.

Compiled from public vendor announcements, SEC filings, earnings releases, conference coverage, and reputable trade press. The publication sits alongside The AI Stack Weekly (cross-stack flywheel) and The Model Pulse (model layer), each available at brianletort.ai/industry.

WHAT EACH SECTION IS FOR

VERTICAL MOVEMENTS	Vertical-specific packages shipped from frontier labs, open-weights labs, or vertical-AI startups. Each card cites a single primary source and names the vertical, origin, and pricing posture.
INCUMBENT RESPONSES	Established SaaS vendors reacting to the agentic shift: product launches, earnings color, partnerships, repositioning, restructuring.
STARTUP SIGNALS	The vertical-AI insurgent cohort: funding rounds, customer wins, GA announcements, M&A. Tracked between frontier labs (top-down) and SaaS incumbents (defending).
PRICING SHIFTS	Public moves from seat-based to outcome-based pricing. Empty in quiet weeks — that itself is a signal of pace.
VERTICAL SCOREBOARD	Leader vs challenger by vertical. A snapshot for procurement shortlists; refreshed every issue.

AUTHORSHIP

Brian Letort

BrianLetort.AI · Independent analysis. All sources public. Not investment guidance.

THIS ISSUE

Issue 30

Week 46 of 2026 · November 14, 2026

WEB

brianletort.ai/industry/applica

The Application Layer archive, plus the sibling publications: The AI Stack Weekly and The Model Pulse.