THE BIG READ

# The fourth meter arrives and a buyer answers with a cap: OpenAI's $500 dot, Microsoft's credits over the seat, 50 million registered agents, and Salesforce reportedly budgeting one technical organisation's employees to $1,000 a month

KEY SIGNALS THIS PERIOD

| 7 | 8 | 6 | 4 |
|---|---|---|---|
| Verticals tracked this period | Vertical packages shipped | Incumbent SaaS responses | Startup signals |

AUTHOR

Brian Letort

BrianLetort.AI

PUBLISHED

November 21, 2026

Issue 31 · 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.

Four meters now sit on enterprise AI price lists where one sat a year ago, and this is the week the fourth became default-on and a buyer-side ceiling was reported in the same seven days. The seat ($30 at Microsoft, $50 at Salesforce's Agentforce Coworker, $20 at Anthropic's Claude for Enterprise, $8 / $40 per million tokens behind it) was joined by the action ($0.10 at Salesforce) and the outcome ($2.00 per resolution at Salesforce, $0.50 at HubSpot); on November 18 OpenAI added the per-agent subscription, turning personal 'dots' on by default in every eligible ChatGPT Enterprise workspace (the help page names Enterprise; the packet summary adds Edu and Healthcare, unconfirmed on the page) at $50, $200 or $500 per dot per month, free until January 13, 2027 and billed from January 14 against existing commitments, with any work the dot delegates to ChatGPT Work or Codex threads metered separately at existing rates. Microsoft did the same thing at Ignite in its own vocabulary: everyday Copilot stays in the $30 seat, while Cowork, Code, the persistent Autopilot agent (with its own Entra Agent ID) and the Work IQ APIs bill through usage-based Copilot Credits, and Agent 365 has registered more than 50 million agents since May (Microsoft's figure). The house arithmetic across price lists: a $500 Expert dot is the price of about seventeen Copilot seats ($30), ten Agentforce Coworker seats ($50), twenty-five Claude seats ($20) or 250 Agentforce resolutions ($2.00) a month, and the buyer has no published measure of what an Expert dot does per month that a Standard one does not. The administrator's only defence is the opt-out, and the default is on.

A buyer answered inside the week, and the answer came from the vendor that sells three of the four meters, with a dispute over its scope. Business Insider and The Information (grade 4) printed a memo from Salesforce's COO setting monthly token budgets after some employees spent thousands a month: $1,000 for technical roles in one organisation per the reports, $500 for other technical delivery roles, $100 for sellers, exceptions by approval, with a pilot team reported running well under budget in month two with no drop in output. Salesforce says it has no companywide individual token budgets and that team leaders manage usage; the house's read is that both can be true and that the memo is a team-level budget, reportedly set in at least one technical organisation, not a corporate

policy. It is not the first per-employee AI spend ceiling at a large company (Business Insider's own article cites Coinbase's $500-$5,000 weekly caps and JPMorgan's $2,000 monthly cap as prior), but it is the first at a vendor that sells the meter, and it arrived in the same week Business Insider reported the company had begun another round of job cuts (scope undisclosed). Set it beside SnapLogic's survey (1,000 IT decision-makers, grade 4): 89% say they have overspent on agents and 42% can forecast token spend. The cap is what a buyer does when the meter is unforecastable, and Salesforce has now done to part of its own payroll what its customers will do to Agentforce.

The remaining events of the week fit the same frame from the other side: who gets paid when the meter runs. Anthropic's final prospectus (grade 5; the weekly owns the offering) fixes the enterprise mix at more than 11,900 customers above $100,000 producing 74% of trailing revenue, every Fortune 10 company a customer, and Claude still under 0.1% of those Fortune 500 customers' expenses, which is the runway the dots and credits are built to fill. Its linked billing (beta, November 17) lets nineteen platforms (Asana, Linear, Ramp, Warp, Factory, Harness among them) run Claude requests that bill directly to the customer's own Anthropic commitment with per-platform caps and no shared keys, moving model cost from the ISV's gross margin to the customer's prepaid budget, which is the mirror image of OpenAI's dots drawing down existing commitments. Intuit showed the incumbent version: QuickBooks takes more than 2 million calls a week from AI tools via MCP connectors inside Claude, ChatGPT, Perplexity, Copilot and Muse, invoicing through the connector runs at a $1 billion-plus annualised payment volume, and the company sold Mailchimp for $3.68 billion to concentrate on being the system behind other people's front doors. Palo Alto Networks bought Applied Compute for $4.5 billion in stock at $100 million of annualised revenue eighteen months after founding, with Prisma AIRS at about $170 million ARR across 1,100-plus customers and the AI Gateway processing up to 3.8 trillion tokens a day in October, which is a security incumbent paying 45x revenue to own the layer where the tokens are counted.

Decisions. A CIO should treat January 14 as a budget event, not a billing date: inventory every workspace where dots are on, decide the opt-out per workspace before the free window closes, and demand from OpenAI the same thing Salesforce's memo demanded of one organisation's staff, a cap per agent and cost attribution per dot. A CFO should write the per-employee ceiling now, with the reported three tiers (and the Coinbase and JPMorgan caps before them) as the template and SnapLogic's 42% as the reason. An ISV should read linked billing

and dots together as the end of model cost as a COGS line, and reprice before the customer's committed budget becomes the thing both parties draw on. An investor should hold the scorecard's rule: Palo Alto's two same-week security datapoints (Applied Compute and Prisma AIRS) are one week, and the challenger slot waits for the second.

**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.

2026-11-18

PACKAGE 01

OpenAI

LEGAL

FRONTIER LAB SEAT

## Astra for Law (GPT-6.2 Astra configuration)

**Requestable by every ChatGPT Enterprise customer via Trusted Access; full-credit answers on Vals AI Legal Research Bench up from ~46% to 53.8% in under two months**

A frontier lab's legal SKU is now a gated option on every Enterprise contract rather than a design-partner programme, and its one independent number (Vals Legal Research Bench, 53.8%) is below the bar a law firm would accept for unreviewed research but moved eight points in two months, so a general counsel should run it against Harvey and Legora on the firm's own matters now rather than wait, and should note that the seat it ships in is the same seat the dots bill against.

SOURCE: OPENAI

2026-11-18

PACKAGE 02

OpenAI

OPERATIONS

FRONTIER LAB

HYBRID

## ChatGPT Enterprise dots (personal agents, default-on)

**Dots on by default in eligible ChatGPT Enterprise workspaces at $50 / $200 / $500 per dot per month; free to January 13, 2027, billed from January 14 against existing commitments; delegated Work and Codex threads metered separately**

This is the fourth meter: a flat subscription per persistent agent, sold on top of the seat and alongside the token, with the agent's downstream work billed on a second meter, so a finance team now models two lines for one dot; the default-on setting with admin opt-out means the first invoice in January reflects administrator inaction, and a CIO should decide per workspace before the free window closes and ask for per-dot cost attribution across the dot's own subscription and the Work and Codex threads it delegates, which the help page does not describe.

SOURCE: OPENAI HELP CENTRE (PRICING AND BILLING FOR DOTS)

2026-11-17

PACKAGE 03

Microsoft

SALES

INCUMBENT SAAS

HYBRID

## Dynamics 365 in Copilot and role plugins (44 sales, 13 service, 30 finance, 20-plus legal skills)

### Copilot Home and Office in Copilot GA, Autopilot rolling to Frontier customers, 100-plus Dynamics skills inside Copilot; Microsoft's internal sales study claims more than 25% less seller effort and 9% faster deal closure (vendor-run, no baseline published)

Microsoft is selling function-specific skill packs inside the $30 seat and metering the persistent and build modes through Copilot Credits, which puts vertical sales, service, finance and legal workflows on a hybrid price list that competes directly with Salesforce's per-action card and OpenAI's dots; a sales operations leader should price a seller's month on all three (seat plus credits, $2.00 per lead worked, a $200 dot) before renewing any one of them, and should treat the 25% / 9% figures as vendor-run until a customer publishes a baseline.

SOURCE: MICROSOFT DYNAMICS 365 BLOG (IGNITE 2026)

2026-11-17

PACKAGE 04

Salesforce (with Microsoft; Accenture as customer zero)

SALES

INCUMBENT SAAS

HYBRID

## Copilotforce: Salesforce in Microsoft Copilot (public beta, 38-plus sales skills), Sales Cloud in Teams (pilot), zero-copy Data 360 to Fabric (GA)

### Salesforce ships its sales workflows inside Microsoft's Copilot surface with Accenture (740,000 Copilot users, 55,000 Salesforce users) as customer zero; ServiceNow and SAP announced parallel Azure and Copilot agent integrations

The system of record has agreed to be a skill inside the system of action's seat, which is the first time Salesforce has put its sales workflow behind another vendor's default agent surface, and Accenture's ratio (13.5 Copilot users for every Salesforce user) is the reason: the seat with the larger population wins the front door. A buyer with both should expect the Headless Platform Interaction multiplier on Salesforce's Flex Credits card (still 'TBA', carried from W43) to be the price of this traffic, and should not sign the Copilot renewal until that multiplier is published.

SOURCE: SALESFORCE NEWSROOM

2026-11-20

PACKAGE 05

Anthropic

RESEARCH

FRONTIER LAB

HYBRID

## Claude Data (default-on for Enterprise), Computer use (beta, Team and Enterprise), linked billing (beta, 19 platforms)

**Claude Data on by default for Enterprise plans from November 20 (certified queries and live sharing stay off; Sheets and Whiteboards stay off); Computer use beta lets Claude operate allowed desktop apps; linked billing lets Asana, Linear, Ramp, Warp, Factory, Harness and others bill Claude usage to the customer's own Anthropic account**

Three administrator decisions landed in one week on the unmetered seat: a default flip that puts warehouse and CRM data into dashboards unless an owner saved a setting, a desktop-operating agent that a security team must scope by allowed application, and a billing mechanism that routes nineteen ISVs' model spend onto the customer's prepaid commitment with per-platform caps and no shared API keys; an enterprise owner should audit the Claude Data setting this week, scope Computer use to a named application list before enabling, and recognise that linked billing turns the Anthropic commitment into a budget that ISVs can draw on, which is the same shape as OpenAI's dots drawing down existing commitments.

SOURCE: ANTHROPIC (CLAUDE PLATFORM; CLAUDE HELP CENTRE)

2026-11-19

PACKAGE 06

Intuit

FINANCE

INCUMBENT SAAS

USAGE

## QuickBooks, TurboTax and Credit Karma MCP connectors inside Claude, ChatGPT, Perplexity, Copilot and Muse

**2 million-plus calls a week from AI tools, more than 30% of connected customers engaging weekly, $1 billion-plus annualised invoice payment volume through the connector since June; Mailchimp sold to Bending Spoons for $3.68 billion; Q1 FY27 revenue $4.4 billion (+12%)**

Intuit is the clearest incumbent case for the system of intelligence behind other people's front doors: it is not selling a seat inside the assistants, it is taking the transaction (invoice payment volume, payroll and bill-pay attach) when an agent in someone else's seat calls its connector, and it sold a $3.68 billion marketing suite to concentrate on that position; a finance-software incumbent should count calls per week from assistants as a distribution metric and price the transaction rather than the access, and a CFO should expect the next accounting-system RFP to ask which assistants can reach the ledger.

SOURCE: INTUIT Q1 FY27 PRESS RELEASE AND EARNINGS SCRIPT

2026-11-16

PACKAGE 07

Abridge; Suki; Ambience Healthcare

HEALTHCARE

STARTUP OUTCOME

## Abridge Assistant with Cochrane (10,000 systematic reviews, Wiley licence); Suki OS with the Suki Guarantee; Ambience Rounding Assistant and Discharge Readiness

**Abridge Assistant live at 300-plus health systems supporting 120 million-plus encounters this year and 40,000-plus VA clinician questions last month; Suki OS carries financial guarantees on measurable outcomes; Ambience cites 6.5% higher case mix index at Onvida Health (all company-reported)**

The ambient-scribe vendors repositioned at HLTH as clinical agent platforms, and two of the three moved on price: Suki's financial guarantee is the first outcome-priced contract form in this vertical and Ambience's case-mix-index claim prices the agent on the revenue line the hospital already measures; a health-system CIO should ask every scribe vendor for the contract form Suki published and should treat the Cochrane licence as the content-moat move, since a guideline-grounded answer is now the product and the transcript is the on-ramp. Abridge's two datapoints (NACHC last week, Cochrane this week) are consecutive and are the kind the scorecard counts.

SOURCE: ABRIDGE PRESS RELEASE; SUKI; BUSINESS WIRE (AMBIENCE)

2026-11-18

PACKAGE 08

Palo Alto Networks

SECURITY

INCUMBENT SAAS

HYBRID

## Applied Compute (post-trained open-weight models for enterprises) and Prisma AIRS

**Palo Alto to acquire Applied Compute for $4.5 billion in stock ($100 million-plus annualised revenue 18 months after founding; close expected fiscal Q3 2027; 35-plus specialised security models planned); Prisma AIRS at ~$170 million ARR (+45% q/q) across 1,100-plus customers; AI Gateway at up to 3.8 trillion tokens a day in October**

A security incumbent paid about 45x revenue to own the layer where enterprise tokens are counted and the models that run on open weights behind the firewall, with a $1 billion NGS ARR target from it by fiscal 2029 and 100-plus MW of inference over three to five years at under $3 billion of capex; a CISO should read Prisma AIRS's run-rate and the 9x year-on-year growth in agentic SASE traffic as the demand side of the agent-governance market, and the scorecard holds Codex Security Cloud and AWS Continuum because the row's rule needs a second week, with Palo Alto now carrying the first.

SOURCE: PALO ALTO NETWORKS Q1 FY27 INVESTOR PRESENTATION (APPLIED COMPUTE FIGURES); PRESS RELEASE AT HTTPS://WWW.PALOALTONETWORKS.COM/COMPANY/PRESS/2026/PALO-ALTO-NETWORKS-REPORTS-FISCAL-FIRST-QUARTER-2027-FINANCIAL-RESULTS

INCUMBENT RESPONSES

# How the SaaS estate is answering.

6 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.

2026-11-17

RESPONSE 01

Microsoft

MICROSOFT 365 COPILOT (HOME, OFFICE IN COPILOT, COPILOT MANAGED RUNTIME GA; AUTOPILOT TO FRONTIER; AGENT 365)

### Everyday Copilot stays in the $30 seat; Cowork, Code, Autopilot and Work IQ APIs bill through Copilot Credits; Agent 365 has registered 50 million-plus agents since May and issues Entra Agent IDs to Salesforce, Google, Databricks, Snowflake, Anthropic and OpenAI agents

Microsoft has split its own seat into a flat everyday tier and a metered agent tier, the same seat-plus-credit shape monday.com and HubSpot adopted, and it is registering other vendors' agents in its identity directory, which makes the directory rather than the seat the control point; a CIO should expect the Copilot Credits line to be where the budget variance shows up from Q1 and should put the 50 million registered agents against the number of agents the organisation knows it runs, because the gap between those two numbers is the governance problem Agent 365 is selling into.

SOURCE: MICROSOFT COPILOT BLOG; AGENT 365 BLOG (IGNITE 2026)

2026-11-17

RESPONSE 02

Salesforce

INTERNAL TOKEN BUDGETS; COPILOTFORCE

### COO memo, printed by Business Insider, sets monthly token budgets of $1,000 (technical roles in one organisation, per the reports), $500 (other technical delivery) and $100 (sellers) with exceptions by approval, after some employees spent thousands a month; a pilot team ran well under budget in month two with no drop in output; Salesforce says it has no companywide individual token budgets and team leaders manage usage; another round of job cuts began the same week, scope undisclosed (Business Insider, The Information)

The vendor that sells the seat, the action and the outcome has reportedly rationed the token for at least one technical organisation of its own staff, while denying a companywide budget; it is not the first per-employee ceiling (Coinbase and JPMorgan caps are cited in the same Business Insider article) but it is a usable template: tiers by role, approval for exceptions, a pilot that shows output holding. That it arrived in the same week as further job cuts and the Copilotforce partnership says the company is paying for its consumption transition on the payroll line and distributing through Microsoft's seat at the same time. Buyers should expect Agentforce's own per-agent caps to follow.

SOURCE: BUSINESS INSIDER; THE INFORMATION

2026-11-18

RESPONSE 03

Zendesk

OUTCOME-BASED BILLING; OUTCOME BUDGET PLANS

### Of 80,000 customers, about 20,000 use AI and 5,000 are on outcome-based billing; unlimited-seat Outcome Budget plans in early access; metered billing on new AI features starts February 1, 2027

This is the first incumbent to publish the denominator: 25% of its AI customers and 6% of its total base are on the outcome meter (the packet does not date its introduction), which is the adoption curve a board should use when a vendor says outcome pricing is the future; the unlimited-seat plan with an outcome budget is the seat and the outcome collapsing into one line, and the February 1 metering date is the next date on the calendar after OpenAI's January 14.

SOURCE: TECHTARGET

2026-11-19

RESPONSE 04

SAP

COMMITTED CONSUMPTION; JOULE STUDIO FREE THROUGH YEAR-END; BUSINESS AI PLATFORM KPI

### CFO at Morgan Stanley TMT: consumption-based cloud revenue to rise from ~10% to ~30% by 2030 via committed consumption with use-it-or-lose-it annual credit tranches; new Business AI Platform revenue KPI from January; Joule Studio free through year-end, contributing to a 70 bp (0.70 percentage point) year-on-year cloud gross-margin decline in Q3 alongside Business Data Cloud mix

SAP has put a number and a date on the seat-to-consumption transition for a system-of-record incumbent (20 points of cloud revenue over four years) and is accepting gross-margin pressure (introductory AI offers are one of two named causes of the 70 bp decline; 50-70 bp expected through 2027) to seed the build surface, which tells a buyer that the use-it-or-lose-it tranche is the contract form to negotiate now while the vendor needs the KPI to move, and tells an investor that the January KPI is the first number that will show whether the finance leader slot on the scorecard is earned on agents or on renamed licences.

SOURCE: INVESTING.COM TRANSCRIPT (SAP AT MORGAN STANLEY TMT)

2026-11-18

RESPONSE 05

HubSpot

HEADLESS PRICING (SEAT PLUS CREDITS)

### Everyday Claude and ChatGPT connector use included in the seat; bring-your-own agents on included API capacity; governed at-scale agent work on HubSpot Credits

HubSpot has formalised the three-tier answer to agent traffic that Salesforce's card leaves 'TBA': free at conversational volume, included at API volume, metered at scale, which gives a buyer the first published template for what cross-platform agent access costs on a system of record; a RevOps leader should take HubSpot's tiers to the Salesforce renewal as the comparable, and the Phase 2 EMEA cohort (W46 watchlist) is still the place the rate will first be visible.

SOURCE: HUBSPOT COMPANY NEWS

2026-11-19

RESPONSE 06

Elastic

HONEYHIVE ACQUISITION (AGENT OBSERVABILITY)

### Elastic agreed to acquire agent-observability startup HoneyHive alongside Q2 FY27 results; terms undisclosed

The second observability incumbent in three weeks to buy the agent-tracing layer rather than build it, which, set beside Palo Alto's AI Gateway at 3.8 trillion tokens a day and Microsoft's Agent 365 registry, says the incumbents have decided the control plane for agents is where the margin is; a platform team evaluating agent observability should expect the independent vendors to be acquired or priced into suites within the year and should contract accordingly.

SOURCE: ELASTIC INVESTOR RELATIONS

STARTUP SIGNALS

# The insurgent vertical cohort.

4 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.

2026-11-17

SIGNAL 01

Model ML

FINANCE

$250M SERIES B AT $1.5B

## $250 million Series B at $1.5 billion co-led by General Atlantic and QED, a year after a $75 million Series A; agents build financial models, pitch decks and diligence reports inside Excel and PowerPoint; company claims it runs on open-source models at about 2% of frontier-model cost

The round prices a vertical finance agent at $1.5 billion a year after a $75 million Series A (the Series A valuation is not in the packet) and its stated cost structure (open models at 2% of frontier cost, company-claimed) is the margin argument Palo Alto paid $4.5 billion for the same week, so an investor in Hebbia, Rogo or the labs' finance SKUs should assume the cost of the model layer is now a competitive variable rather than a shared input, and a bank buying these tools should ask which model runs each task and at what cost per deliverable.

SOURCE: TECHCRUNCH

2026-11-18

SIGNAL 02

Payhawk

FINANCE

$280M SERIES C AT $2B

## $280 million Series C at $2 billion led by Goldman Sachs Growth for AI-agent-driven enterprise spend management across 33 countries; total funding over $500 million

Spend management is the finance workflow where agents meet the token-budget problem reported at Salesforce this week, and Goldman's growth fund pricing a European spend-management vendor at $2 billion says the capital markets expect the FinOps-for-agents category to be a line item; a CFO should expect its spend-management vendor to offer per-employee AI budgets as a feature within two quarters and should ask for it at renewal.

SOURCE: TECHCRUNCH

2026-11-18

SIGNAL 03

Alecto

LEGAL $20M SEED

## $20 million seed led by Insight Partners for an AI-native UK conveyancing firm at $8 million annualised revenue after eight months on fixed fees; Perry AI raised $2.2 million for private-capital legal agents

A law firm, not a legal-software vendor, is the entity raising venture capital on fixed-fee AI-delivered conveyancing, which is the outcome-priced legal model Harvey and Legora sell tools for rather than practise; a managing partner should read $8 million in eight months on fixed fees as the first disclosed revenue curve for AI-native legal practice, and the legal row's challenger rule now has a startup datapoint beside OpenAI's Astra for Law.

SOURCE: GLOBAL LEGAL POST

2026-11-18

SIGNAL 04

Applied Compute

SECURITY

$4.5B ACQUISITION

### Acquired by Palo Alto Networks for $4.5 billion in stock (including about $500 million of cash and prepaid compute) at $100 million-plus annualised revenue 18 months after founding; post-training and inference platform for open-weight models

The largest exit in this series for a company whose product is post-training open weights for enterprises, at about 45x revenue, which values the capability of running a customer's own model behind the firewall above any single vertical package; a founder building on open weights should note that the acquirer is a security vendor rather than a cloud or a lab, and a buyer should expect 'your model, our gateway' to become a security-suite SKU by fiscal Q3 2027.

SOURCE: PALO ALTO NETWORKS INVESTOR PRESENTATION

PRICING SHIFTS

# Seat to outcome, one move at a time.

5 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.

**2026-11-18**

SHIFT 01

OpenAI

## Seat-based → Hybrid

ChatGPT Enterprise (the help page names Enterprise; the packet summary adds Edu and Healthcare, which the house has not confirmed on the page) adds a per-agent subscription on top of the seat: personal dots at $50 (Standard), $200 (Advanced) or $500 (Expert) per dot per month, on by default unless an admin opts out, free through January 13, 2027 and billed from January 14 against existing usage commitments or credits; work a dot delegates to ChatGPT Work or Codex threads is billed separately at existing rates, so one agent runs two meters.

SOURCE: OPENAI HELP CENTRE

---

**2026-11-17**

SHIFT 02

Microsoft

## Seat-based → Hybrid

Microsoft 365 Copilot formalises a split seat at Ignite: everyday Copilot (Chat and Cowork in Home) remains inside the $30 per user per month seat, while Code, Autopilot (persistent tenant-resident agent), Copilot Managed Runtime workloads and Work IQ APIs bill through usage-based Copilot Credits; Microsoft's business pricing page carries both lines.

SOURCE: MICROSOFT COPILOT BUSINESS PRICING

---

**2026-11-18**

SHIFT 03

Zendesk

## Seat-based → Outcome-based

Unlimited-seat Outcome Budget plans enter early access, collapsing the seat into an outcome allowance; metered billing on new AI features begins February 1, 2027 after a free period. Zendesk discloses the denominator: 5,000 of about 20,000 AI customers (of 80,000 total) are on outcome-based billing.

SOURCE: TECHTARGET; ZENDESK SUPPORT ANNOUNCEMENT

---

**2026-11-19**

SHIFT 04

SAP

## Seat-based → Usage-based

SAP moves cloud contracts toward committed consumption with use-it-or-lose-it annual credit tranches, targeting consumption-based cloud revenue rising from about 10% to about 30% by 2030, with a Business AI Platform revenue KPI reported from January; Joule Studio is free through year-end, with introductory AI offers one of two named causes of a 70 bp year-on-year cloud gross-margin decline (CFO at Morgan Stanley TMT).

SOURCE: INVESTING.COM TRANSCRIPT

---

2026-11-17

SHIFT 05

Anthropic

## Usage-based → Hybrid

Linked billing (beta) on the Claude Platform lets nineteen third-party platforms (Asana, Factory, Linear, Ramp, Harness, Warp among them) run Claude requests that Anthropic bills directly to the customer's own organisation, drawing down existing commitments with per-platform spend caps and no shared API keys; model spend moves from the ISV's cost of goods to the customer's committed Anthropic budget.

SOURCE: ANTHROPIC (CLAUDE PLATFORM)

VERTICAL SCORECARD

# Who leads each vertical.

Leader-vs-challenger by vertical, useful for procurement shortlists when matching workload to vendor cohort. As of Nov 21, 2026.

| VERTICAL | LEADER | CHALLENGER | READ |
|---|---|---|---|
| **Commerce** | **Amazon Seller Assistant** | Sierra-led Poppy standard partners (Meta, Shopify, Stripe, Walmart) | Held; no commerce package or checkout disclosure in the window. The adjacent datapoint is Intuit's $1 billion-plus annualised invoice payment volume through assistant connectors, which is a payments event in the finance row, and Microsoft's Agent 365 issuing identities to Shopify-class partners is a governance event. Neither moves the slots. |
| **Engineering** | **Codex (on GPT-6.2 Sol and Astra from November 17)** | Claude Code (on Opus 5.5 and Fable 5.2) | Held, with the leader's default model updated: GPT-6.2 Sol is now Codex's default and Astra did not move the independent Coding Agent Index (The Model Pulse). Gemini 4 Neon in Antigravity CLI took first on that index at 59 against Claude Code's 58, which is one datapoint for a Google challenger under the row's rule; Cognition's reported $75 billion talks are unconfirmed and do not count. Harness releases are in Agent Techniques Weekly. |
| **Healthcare** | **Abridge** | OpenEvidence | Held, with the leader's second consecutive datapoint: the Cochrane licence (10,000 reviews) and 300-plus health systems / 120 million-plus encounters (company-reported) follow last week's NACHC deal. Suki's financial guarantee is the first outcome-priced contract form in the row and is one datapoint toward a challenger change; Ambience's 6.5% case-mix-index claim is a second vendor's single datapoint. OpenEvidence had no disclosure this week. |
| **Support** | **Sierra** | Salesforce Agentforce | Held. Zendesk published the first denominator on outcome billing (5,000 of 20,000 AI customers) and an unlimited-seat Outcome Budget plan, one datapoint toward a challenger change; Salesforce's $2.00 per resolution stands and its reported internal token budgets are a buyer-side event. Sierra's outcome pricing remains unpublished, so the row still has one published meter and one private one. |
| **Sales** | **ZoomInfo** | Salesforce Agentforce with Docusign IAM | Held on the standing rule. Two incumbent events landed in the row this week (Microsoft's 44 sales skills and Dynamics 365 in Copilot with vendor-run 25% / 9% claims; Copilotforce with Accenture as customer zero) and both are first datapoints for a Microsoft challenger; Rox |

| VERTICAL | LEADER | CHALLENGER | READ |
|---|---|---|---|
| | | | (W46) has one datapoint. A change needs the second week from any one of them. |
| **Security** | **Codex Security Cloud** | AWS Continuum | Held under the rule, with the flag that the rule exists for: Palo Alto Networks produced two datapoints in one week (the $4.5 billion Applied Compute agreement and Prisma AIRS at ~$170 million ARR across 1,100-plus customers), which is one week, not two. A second disclosure next week (close terms, AIRS customer count, AI Gateway volume) moves Palo Alto into the challenger slot. Elastic's HoneyHive purchase is observability, adjacent to the row. |
| **Legal** | **Harvey** | Legora | Held. Astra for Law is now requestable by every ChatGPT Enterprise customer with an independent 53.8% on Vals Legal Research Bench (one datapoint for a frontier-lab challenger); Alecto's $20 million seed at $8 million revenue on fixed fees is one datapoint for an AI-native firm as a different kind of challenger; Microsoft's 20-plus legal skills in Copilot are a third. Neither row holder disclosed in the window. |
| **Research** | **Claude Data (default-on for Enterprise from November 20)** | OpenAI dots (default-on from November 18, $50 / $200 / $500 per dot per month) | Challenger changed from the Gemini agent to OpenAI dots under the row's rule: two consecutive weeks of disclosed datapoints (W46: billing-by-default date and AWS Marketplace availability; W47: the published per-dot price list, the January 14 billing date and the two-meter structure) while the Gemini agent had none. The leader's default flip landed on schedule; Computer use beta extends the leader's reach to the desktop. |
| **Finance** | **SAP Joule** | Workday agentic finance | Held, with the leader's first number: SAP's CFO put consumption at ~10% of cloud revenue rising to ~30% by 2030 and will report a Business AI Platform KPI from January. Intuit's connector volumes (2 million-plus calls a week, $1 billion-plus invoice payment volume) are a small-business finance datapoint and the strongest incumbent-distribution evidence in the row; Model ML's $250 million at $1.5 billion is a banker-workflow startup datapoint. Workday had no disclosure. |

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

## The fourth meter: a flat subscription per persistent agent, sold on top of the seat, with the agent's work metered on a second line

**OpenAI dots: $50 / $200 / $500 per dot per month, default-on, billed from January 14, 2027; delegated Work and Codex threads at existing rates**

**Microsoft Autopilot: persistent tenant-resident agent with its own Entra Agent ID, billed through Copilot Credits over the $30 seat**

**Salesforce Agentforce Coworker: $50 seat alongside $0.10 per action and $2.00 per resolution (W46)**

A year ago enterprise AI had one meter (the seat); W46 recorded three (seat, action, outcome) on one Salesforce card; this week the per-agent subscription became default-on at OpenAI and the persistent agent got its own identity and credit line at Microsoft. The structure matters more than the price: the dot is a subscription whose consumption is billed elsewhere, so the administrator sees a fixed monthly fee and a variable usage line and cannot attribute the second to the first without per-dot cost attribution, which OpenAI's help page does not describe. The house illustration is that a $500 Expert dot equals about seventeen Copilot seats ($30), ten Agentforce Coworker seats ($50), twenty-five Claude seats ($20) or 250 Agentforce resolutions ($2.00) a month with no published measure of Expert's incremental work. The operator's response is the one Salesforce's memo reportedly applied to one of its own organisations: a cap per agent, set before the free window closes.

SOURCE: OPENAI HELP CENTRE; MICROSOFT COPILOT PRICING

PATTERN 02

## The buyer-side cap: per-employee AI spend ceilings reach the vendor that sells the meters

**Salesforce: $1,000 / $500 / $100 per month by role, reported for one technical organisation, exceptions by approval, pilot team under budget with output held; company says no companywide budgets (Business Insider, The Information)**

**SnapLogic survey: 89% have overspent on agents, 42% can forecast token spend, 54% have abandoned agentic projects for unclear value**

**Anthropic linked billing: per-platform spend caps on the customer's commitment**

Caps are what buyers do when meters are unforecastable, and this week's reported per-employee ceiling came from Salesforce's own COO (a memo Business Insider printed, covering at least one technical organisation; Salesforce says there is no companywide budget) in the week the company sold Copilotforce and began another round of cuts. It follows the Coinbase and JPMorgan caps the same article cites. The three-tier design (by role, with approval for exceptions) is the template, and the reported result (a pilot team well under budget in month two with no drop in output) is the argument every CFO will make. Anthropic's linked billing builds the same control into the vendor side with per-platform caps on the customer's commitment, and Payhawk's $2 billion valuation says spend-management vendors expect to sell the feature. The structural read is that the consumption transition now has a ceiling on both ends: the vendor meters, the buyer caps, and the number that clears is negotiated per role rather than per seat.

SOURCE: BUSINESS INSIDER; SNAPLOGIC

PATTERN 03

## The customer's commitment becomes the budget everyone draws on: dots, linked billing and Copilot Credits all bill against prepaid balances

**OpenAI dots draw from existing usage commitments or credits from January 14**

**Anthropic linked billing: nineteen ISVs bill Claude requests to the customer's own Anthropic organisation, no shared API keys**

**Microsoft Copilot Credits over the seat; SAP use-it-or-lose-it annual consumption tranches**

Three frontier and incumbent vendors in one week wired new products to the customer's existing prepaid balance rather than to a new invoice, and SAP described the contract form explicitly (committed consumption, annual tranches that expire). For an ISV the consequence is that model cost stops being a cost of goods and becomes a draw on the customer's lab commitment, which improves the ISV's gross margin and removes its pricing power over the model layer at the same time; for the buyer it means the commitment negotiated for API usage is now also the budget for agents it did not open and ISVs it did not meter. A procurement team should re-read every commitment contract for who may draw on it and ask for per-drawer caps, which Anthropic offers and OpenAI's dots do not yet.

SOURCE: ANTHROPIC (CLAUDE PLATFORM); OPENAI HELP CENTRE

PATTERN 04

## The system of record becomes a skill inside the system of action's seat, and the system of intelligence takes the transaction behind both

**Copilotforce: Salesforce sales workflows as skills in Microsoft Copilot; Accenture 740,000 Copilot users to 55,000 Salesforce users**

**Intuit: QuickBooks connectors in Claude, ChatGPT, Perplexity, Copilot and Muse at 2 million-plus calls a week, $1 billion-plus invoice payment volume**

**Agent 365: Entra Agent IDs issued to Salesforce, Google, Databricks, Snowflake, Anthropic and OpenAI agents**

The 'system of record versus system of action' fight moved inside Microsoft's surface this week, with Salesforce accepting the skill position because the population ratio at its customer zero is 13.5 to one, and Intuit showed the third position: a system of intelligence that neither owns the seat nor the skill but takes the payment when either calls its connector. The architecture that emerges has three layers, the seat (Microsoft, OpenAI, Anthropic), the skill (Salesforce, Dynamics, HubSpot) and the transaction (Intuit, payments), with the identity directory (Agent 365) cutting across all three. Incumbents should decide which layer they are pricing, because Salesforce's Headless Platform Interaction multiplier ('TBA' since W43) is the price of being a skill, and Intuit's connector volumes are the revenue of being the transaction.

SOURCE: SALESFORCE NEWSROOM; INTUIT EARNINGS SCRIPT

PATTERN 05

## Adoption measured, again: one in three firms sees P&L results, one in five AI chiefs controls the budget, one in two has abandoned an agent project

**Bain (n=200, August survey): 33% report measurable revenue or cost results (+10 points YoY); software-development headcount expectation swung from -2% to +10%**

**IBM IBV (n=1,667): 18% of AI leaders have full control of AI investment; high-impact CAIOs report 63% of projects with measurable outcomes vs 43%**

**SnapLogic (n=1,000): 54% abandoned AI projects for unclear value; 48% cannot measure agentic ROI**

Three surveys in one week give the same shape from three angles: results exist and are concentrated (Bain's one in three, with workflow redesigners twice as likely to exceed goals), authority over the spend is rare (IBM's 18%), and measurement is the failure mode (SnapLogic's 48% and 42%). Bain's headcount reversal in software development (from expected cuts to expected 10% growth) is the single most useful number for a board, because it says the firms that got results are hiring engineers to get more rather than cutting them, which runs against the severance-funded transition recorded at Salesforce. The surveys are vendor-adjacent (two from consultancies with AI practices, one from an integration vendor) and are labelled as such; their agreement is the signal, not any one percentage.

SOURCE: BAIN; IBM; SNAPLOGIC

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.

BEFORE JANUARY 13, 2027

WATCH 01

## Dots opt-out rates and the first per-dot usage report from OpenAI

The default-on free window ends January 13 and billing starts January 14 against existing commitments; the number to look for is any disclosure of what share of Enterprise workspaces disabled dots, and whether OpenAI publishes per-dot consumption before the first invoice, which is the control a CIO needs and the help centre does not yet describe.

NEXT WEEK

WATCH 02

## Palo Alto Networks' second security datapoint: Applied Compute close terms, Prisma AIRS customer count or AI Gateway volume

The scorecard's rule needs two consecutive weeks; Palo Alto produced two datapoints in one week, and a second disclosure next week moves it into the security challenger slot over AWS Continuum.

Q4 PRINTS (DECEMBER) AND SALESFORCE'S NEXT CARD REVISION

WATCH 03

## Salesforce Headless Platform Interaction multiplier published, and whether Agentforce gets per-agent caps

Copilotforce makes cross-platform agent traffic the main event for Salesforce's sales workflows, and the multiplier (TBA since W43) is its price; the reported token budgets inside its own technical organisation make per-agent caps on Agentforce the natural next product, which would be the vendor-side answer to its own buyer-side move.

FEBRUARY 1, 2027

WATCH 04

## Zendesk metered billing on new AI features begins; Outcome Budget plans move from early access

Zendesk published the denominator (5,000 of 20,000 AI customers on outcome billing); the February 1 metering date is the first chance to see whether unlimited-seat outcome budgets move that share, and it is the next pricing date after OpenAI's January 14.

JANUARY 2027

WATCH 05

## SAP's first Business AI Platform revenue KPI and the consumption share of cloud revenue

The finance row's leader has put a number (~10% to ~30% by 2030) and a reporting date on the seat-to-consumption transition; the January KPI is the first datapoint on whether Joule's free-through-year-end seeding converts to committed consumption.

BY FISCAL Q3 2027 (PALO ALTO) AND ONGOING

WATCH 06

### Linked billing adoption: how many of the nineteen platforms route material Claude spend to customer commitments, and whether OpenAI matches it

Linked billing moves model cost from the ISV's COGS to the customer's prepaid budget with per-platform caps; a disclosed share of ISV traffic on linked billing, or an OpenAI equivalent for dots-delegated Codex work, would confirm the commitment-as-shared-budget architecture.

HLTH FOLLOW-THROUGH, DECEMBER

WATCH 07

### A health system publishes the contract form behind the Suki Guarantee or a second Abridge outcome figure

Suki's financial guarantee is the first outcome-priced form in healthcare and Ambience's case-mix claim is vendor-reported; a customer-published guarantee term or an independent outcome figure is the datapoint that moves the healthcare challenger slot.

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 31**

Week 47 of 2026 · November 21, 2026

WEB

**brianletort.ai/industry/applica**

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