THE APPLICATION LAYER ISSUE 29 · WEEK 45 OF 2026

THE BIG READ

Last week the vendors stopped reporting the seat; this week they stopped publishing the meter, and the first agent-commerce prints said four rides

KEY SIGNALS THIS PERIOD

6	7	5	3
Verticals tracked this period	Vertical packages shipped	Incumbent SaaS responses	Startup signals

AUTHOR

Brian Letort

BrianLetort.AI

PUBLISHED

November 7, 2026

Issue 29 · 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.

The consumption model entered its second phase this week, and the phase is defined by what the buyer cannot see. Microsoft extended 'fair use limits' from image generation to additional Copilot Chat features for every commercial user who lacks a paid Microsoft 365 Copilot licence, rolling from November 2 for tenants with 2,000 or fewer seats and November 9 for larger ones; its support page states that 'because service limits can vary, Microsoft does not publish or share specific numeric limits', administrators cannot view or adjust them, and a user who hits one receives an in-product notice that in most organisations includes an option to request or buy the add-on. OpenAI made 'dots', agents that keep working between conversations, generally available to ChatGPT Enterprise, Edu and Healthcare workspaces with an admin preview to November 18, after which dots are available by default and, in the docs' words, admins can 'review preview access and billing before billing begins'. HubSpot told investors its Phase 2 pricing, a go-to-market seat plus a single credit for all usage-based AI, moves from a Nordics and Benelux pilot to every new EMEA customer in the fourth quarter because pilot cohorts consumed credits at two to three times the rate of other regions while average selling prices fell. Each of the three is a meter whose quantity is set by the vendor after the purchase: an undisclosed cap, a default-on agent, a credit whose consumption rate the vendor learned from the pilot and the buyer will learn from the invoice. UC Today named the seats-to-meters move last week; The Register's reading of the Microsoft change (caps you discover by hitting them) is the sharpest statement of this week's, and this publication's extension is that the unpublished meter is the natural second step, because once revenue is consumption the vendor's interest is in the buyer consuming before the buyer prices.

What a buyer should do differently is set the quantity before the vendor does. For Microsoft that means deciding, before November 9, whether the included Copilot Chat is a product or an on-ramp, because from that date it is a throttled trial whose throttle the administrator cannot read; for OpenAI it means using the fourteen-day preview to decide which members get dots rather than letting the default decide on November 18; for HubSpot it means writing a credit-consumption cap or a published per-unit credit rate into the renewal, since the

vendor has said the ratio (two to three times) and not the amount. The counter-evidence this week came from the buyer's side and from the transaction layer, and it is strong enough to carry in the lead. Goldman Sachs updated internal guidance in October limiting Claude Code and GitHub's coding tools to engineers and quantitative strategists after months of banker and trader use, following what Bloomberg describes as a review of 'productivity and risk management', with a spokesperson saying on the record that after 'a period of experimentation' the tools go to 'those who are best suited to leverage them for the long term'. The report is grade 4 on the 1-5 evidence scale (where 5 is a primary filing and 3 is single-outlet reporting). It reverses the bank's July endorsement of non-technical staff building prototypes for its roughly 12,000 engineers to polish, and it is the first disclosed case in this series of a regulated enterprise narrowing an agent seat after a review; whether the review measured return or managed risk is not stated, and this publication does not assume the former.

The Bank of England's November Monetary Policy Report is the macro context, not the cause: firms on its Decision Maker Panel estimate AI raised their productivity by about 0.5% over three years, roughly 0.1-0.2 percentage points a year, with close to zero aggregate employment effect, while expecting about 0.9 points a year for the next three; Bank staff attribute just under a fifth of 2026 UK growth to AI-related activity and locate it upstream in development and implementation rather than in adoption gains. HubSpot's own print carries the same shape at company scale: net revenue retention of 101%, down a point, full-year net new ARR (annual recurring revenue) guided to decline on what management called budget sensitivity and longer sales cycles, and a $65-75M restructuring charge, while vendor-reported agent metrics rose (Customer Agent past 2 million resolutions in the quarter, 20% of Pro-plus customers using agents, agentic actions up 5x); the two moved in opposite directions in the same quarter, and HubSpot named the causes of the first, not a link between them. Palantir's quarter is the version where the meter wins and the duration loses: US commercial revenue of $984M (up 148%) passed government for the first time, net dollar retention reached 165% and the full-year guide rose to $8.53B, while US commercial contract-value bookings grew 11% against 153% in the prior quarter and remaining performance obligations were flat sequentially, which management attributes to shorter contract durations. That faster-cycling procurement is what a metered product looks like from the vendor's bookings line is this publication's inference, not Palantir's statement.

The third thing the week settled is the size of agent-originated commerce, and the number is small enough to state exactly. Uber's CFO said that in the week of Halloween the company fulfilled about 320 million trips and Eats orders, of which about 500 Eats orders and 'a grand total of four' rides came from third-party AI agents (The Verge drew out the four-rides figure first and is credited); its CEO said eight assistants are negotiating 'terms of trade' and that Uber's own cart assistant produces baskets about twice the average. Shopify, on its November 3 call, reported AI-referred traffic up 2.6x year on year, AI-channel orders up 3.3x, a 64% higher conversion rate for AI shoppers and Catalog-powered searches converting at twice the rate of general AI search, with its president adding that agentic volume 'is still small' and that the quarter was the first with delegated purchasing through Muse, Instinct, Claude Duet and Grok Bot over Shop Pay. The incumbents are pricing access to the catalogue, the wallet and the terms of trade rather than ceding the checkout, which is the commerce scorecard's standing reading and now has two earnings calls behind it. Against those numbers Anthropic's Personal Finance hub, launched November 5 with Plaid, is the week's most carefully bounded vertical package: read-only, up to 24 months of transactions, US Pro, Max and Ultra subscribers only, excluded from Team and Enterprise plans and unavailable to scheduled tasks, Cowork and Claude Code. The data boundary is the product boundary, and it was drawn on the consumer side of the line. The Anthropic S-1/A's updated customer metrics (more than 11,900 customers at $100K-plus, consumption 79% of nine-month revenue) are carried in the pricing section; the Weekly owns the financials.

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.

7 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.

2026-11-02

PACKAGE 01

Microsoft

OPERATIONS

INCUMBENT SAAS

HYBRID

Copilot Chat fair-use limits for unlicensed Microsoft 365 commercial users

Unpublished usage caps extend from image generation to additional Copilot Chat features for work accounts without a paid Copilot add-on; tenants of 2,000 seats or fewer from November 2, larger tenants November 9; admins cannot view or adjust the limits

The included Copilot Chat entitlement in Business Basic, Standard and Premium and in Office 365 and Microsoft 365 E3 and E5 becomes a throttled on-ramp to the paid add-on, one week after Microsoft stopped disclosing paid Copilot seat counts (W44). Microsoft's support documentation is explicit that numeric limits are not published or shared, that they apply across web, desktop and mobile, and that the limit notice includes a licence-request or purchase option where self-service is permitted. A university IT department published a staff notice for the November 9 wave, which is how most organisations will learn the limit exists. The operating consequence is that the free tier is now a meter with a hidden quantity and a buy button at the edge.

SOURCE: MICROSOFT SUPPORT; THE REGISTER; COMPUTERWORLD

2026-11-04

PACKAGE 02

OpenAI

RESEARCH

FRONTIER LAB

HYBRID

Personal dots (persistent agents) in ChatGPT Enterprise, Edu and Healthcare

Agents that keep working between conversations reach GA with an admin preview to November 18, after which dots are available by default and billable; GPT-6.1 Terra is among the models powering them

A new always-on agent surface enters Enterprise credit consumption on an opt-out basis inside two weeks. Admins can control which members get dots, view personal dots and their connections, pause or resume them, review usage, and owners can delete them; the docs say admins can 'review preview access and billing before billing begins on November 18'. Dots support EU and US inference residency, Enterprise Key Management and are HIPAA-eligible in covered workspaces, and are unavailable to UAE data-residency workspaces; dots created before GA in residency workspaces move to the required region within 30 days. Separately, OpenAI is merging the Chat and Work tabs and pooling ChatGPT and Codex usage into one allowance for selected users (TechCrunch, November 2), which makes the dots meter and the chat meter one bucket. Agent Techniques carries what a persistent agent needs from the harness.

SOURCE: OPENAI HELP CENTER (ENTERPRISE AND EDU RELEASE NOTES); CHATGPT DOCS

2026-11-04

PACKAGE 03

HubSpot

SALES

INCUMBENT SAAS

HYBRID

Phase 2 pricing: go-to-market seat plus a single AI credit, expanded to all new EMEA customers

The Nordics and Benelux pilot (hubs removed, Starter-to-Pro step cut, all usage-based AI on one credit) saw ASPs fall, Pro-plus counts rise and credit consumption run 2-3x other regions; it covers every new EMEA customer in Q4

The clearest disclosed experiment in seat-plus-meter pricing: a lower entry price, more customers on the higher tier, and a credit meter running at two to three times the consumption of customers on the old structure, all HubSpot-reported. The adoption figures underneath are vendor metrics (monthly credit consumption up about 2.5x since January, 20% of Pro-plus customers using agents, up 11 points year to date, monthly agentic actions up 5x, Customer Agent past 2 million resolutions in the quarter, prospecting-agent users reporting 76% more qualified leads, Breeze Assistant used by 61% of Pro-plus customers, active core seats up 34% after seat minimums were removed). The company metrics are the check: revenue $947.6M (up 17%), net revenue retention 101% (down one point), full-year net new ARR guided to decline on budget sensitivity and longer sales cycles, Q4 guided to about 12% growth, and a $65-75M severance charge for an October reorganisation. The meter is growing faster than the account.

SOURCE: HUBSPOT INVESTOR RELATIONS; HUBSPOT PREPARED REMARKS; Q3 CALL TRANSCRIPT

2026-11-03

PACKAGE 04

Shopify

COMMERCE

INCUMBENT SAAS

UNKNOWN

Agentic commerce on Shop Pay (Muse, Instinct, Claude Duet, Grok Bot) and Shopify Catalog

AI-channel orders up 3.3x year on year and AI shoppers converting 64% more often, Catalog searches converting at 2x general AI search, and the first quarter of delegated purchasing through four assistants; the president says agentic volume 'is still small'

The first earnings-call figures on third-party agents as a transaction channel, all Shopify's: AI-referred traffic up 2.6x, AI channels sending shoppers to product pages at three times the rate of organic search, about three-quarters of AI-attributed purchases from categories outside Shopify's top 100, Catalog structuring more than a billion product pages, and 'we're not seeing that prices are being lowered' by agent comparison shopping. Revenue was $3.742B (up 32%) on GMV of $119.1B (up 29.5%) with a 22% free-cash-flow margin, so the agent channel is a rounding error inside a strong quarter. The pricing posture is the item: Shopify sells access to Catalog and Shop Pay to the assistants rather than ceding the checkout to them, which is the commerce row's thesis. Uber's figures on the same question are in the incumbent section.

SOURCE: SHOPIFY INVESTOR RELATIONS; Q3 CALL TRANSCRIPT; CNBC

2026-11-05

PACKAGE 05

Anthropic

FINANCE

FRONTIER LAB SEAT

Personal Finance in Claude (beta) with Plaid

Plaid-linked, read-only view of bank, card, loan and investment accounts with up to 24 months of transactions, US Pro, Max and Ultra only; excluded from Free, Team and Enterprise and from scheduled tasks, Cowork and Claude Code

A consumer finance vertical built with the enterprise agent surfaces walled off by design: Claude sees institution and account names, last four digits, balances, credit limits and transactions, 'can't move, manage, or invest your money', and hub conversations are not used for training; linked accounts cannot be reached from the agentic surfaces at all. It follows OpenAI's Plaid-connected ChatGPT feature from May. The finance scorecard row is about finance-function agents in the enterprise, which this is not; it is recorded because the exclusion list is the first case in this series of a frontier lab shipping a data connector and naming the agent surfaces it will not serve, a pattern the architecture section reads. Plaid's survey figure (69% of US consumers willing to share personal data with AI in financial services given clear benefit) is Plaid's.

SOURCE: ANTHROPIC; CLAUDE HELP CENTER; PLAID

2026-11-04

PACKAGE 06

Cognition

ENGINEERING

STARTUP USAGE

Devin Review 2.0 and SWE-2 Flash in Devin Desktop and CLI

Per-PR review cost down 82% since August to a typical $0.25-0.75 with a configurable per-PR spend limit, and a new in-house model free in Devin until December 4, then $0.10 / $0.50 per million tokens with a 25% enterprise discount to November 18

Agentic code review priced in cents with a hard cap as a feature: Cognition says about 80% of PRs now cost under $1 including all pushes, bugs found per PR rose from 1.22 to 1.32, median latency fell from 139 to 94 seconds, and the cost per bug fixed is about $2.25, with incremental re-review of changed hunks the default and Review 2.0 the default reviewer; every figure is Cognition's. The Devin plan mechanics for SWE-2 Flash (free on self-serve to December 4 at 18:00 UTC, then $0.10 input / $0.01 cache read / $0.125 cache write / $0.50 output for all three effort variants, Medium available on the Free plan, enterprise ACU plans 25% off to November 18) live here; the model, its withheld weights and its benchmark table live in The Model Pulse. The engineering row's price field now runs from $0.25 per PR review to Codex's and Claude Code's plan allowances.

SOURCE: DEVIN BLOG (REVIEW 2.0); DEVIN MODEL DOCS

2026-11-02

PACKAGE 07

Palantir

OPERATIONS

INCUMBENT SAAS

HYBRID

AIP (Palantir US commercial)

US commercial revenue of $984M (up 148%) passes US government for the first time and the full-year guide rises to $8.53B, while US commercial contract-value bookings grow 11% against 153% a quarter earlier and RPO is flat sequentially on shorter durations

The largest operations-platform vendor in the field reported the meter and the duration moving in opposite directions: 686 US commercial customers (up 29%), net dollar retention of 165% (up 8 points quarter on quarter), 212 deals of $1M or more, GAAP operating margin of 47% and $10.4B of cash, against closed total contract value up 6% to $2.918B, US commercial bookings of $1.463B, and remaining performance obligations flat at $4.9B. Management attributed the bookings deceleration to shorter contract durations. For a buyer the read is that AIP procurement is cycling faster and committing shorter; that this is what a consumption-weighted product does to a bookings line, and consistent with the HubSpot and ServiceNow prints rather than a Palantir-specific event, is this publication's inference, not Palantir's statement.

SOURCE: PALANTIR INVESTOR RELATIONS; Q3 BUSINESS UPDATE; CALL TRANSCRIPT

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.

2026-11-02

RESPONSE 01

Microsoft

COPILOT CHAT (INCLUDED ENTITLEMENT) AND MICROSOFT 365 COPILOT ADD-ON

Fair-use limits extended to more Copilot Chat features for unlicensed commercial users; numeric limits not published; admins cannot view or override; limit notices carry a licence-request or purchase option

The included tier is converted into the top of the add-on funnel the week after the paid seat count left the earnings deck. A buyer negotiating E3 or E5 renewals should treat Copilot Chat as a metered trial rather than a bundled feature and should ask for the limits in writing, since the product will not show them. Pricing model defended: seat plus add-on with a hidden consumption gate (hybrid).

SOURCE: MICROSOFT SUPPORT; THE REGISTER

2026-11-05

RESPONSE 02

ServiceNow

WORKFORCE (THIRD 2026 REDUCTION)

A new global round of job cuts across the US, UK and India with no number disclosed; severance of $87M in Q3 and $167M year to date (against $24M and $53M a year earlier) under a programme 'to optimize our workforce through AI-driven efficiencies'; shares down about 35% between October 23 and November 5; grade 3

The incumbent that told investors last week that token economics could lift revenue per user fivefold is paying for the transition in severance: the 10-Q severance line is the filed part, the headcount is not disclosed, and the stated goal is to 'end 2026 where we started'. Management's claim of more than $1B of annualised AI efficiency gains is the company's. SUSE's cut of about 10% of roughly 1,900 staff, attributed by its CEO to 'automation and AI', is the smaller instance in the same week. Pricing model defended: seat plus consumption-billed agentic SKUs (hybrid), with the cost base being moved to match.

SOURCE: CALCALIST / CTECH; PEOPLE MATTERS; THE REGISTER (SUSE)

2026-11-04

RESPONSE 03

Uber

THIRD-PARTY AGENT ACCESS (EIGHT ASSISTANTS; 'TERMS OF TRADE' NEGOTIATIONS) AND UBER'S OWN CART ASSISTANT

About 500 Eats orders and 'a grand total of four' rides came from AI agents in a week of about 320 million trips and orders; eight assistants are integrated with more than ten expected; the in-house cart assistant produces baskets about twice the average

The most precise disintermediation datapoint any marketplace has published: agent-originated demand is measurable to the unit and immaterial to the quarter, and the incumbent's own agent is the one moving the basket. Uber is negotiating access terms with the assistants rather than opening the checkout, which places it with Shopify on the pricing-access side of the commerce row. Pricing model defended: take rate on the transaction, with agent access priced separately (hybrid).

SOURCE: UBER Q3 2026 CALL TRANSCRIPT; THE VERGE

2026-11-02

RESPONSE 04

OpenAI

CHATGPT CHAT AND WORK TABS; POOLED CHATGPT AND CODEX ALLOWANCE

Chat and Work tabs being merged and ChatGPT and Codex usage pooled into one allowance for selected users (TechCrunch); dots default-on and billable from November 18

One bucket for chat, delegated work, coding and persistent agents means the seat's allowance is now drawn down by surfaces the user did not open, and the admin's lever is membership rather than quantity. It is the frontier-lab version of the unpublished meter: the quantity is published, the consumers of it multiply. Pricing model defended: seat with a pooled task allowance (hybrid).

SOURCE: TECHCRUNCH; CHATGPT DOCS (WHAT'S NEW, NOVEMBER 2-6)

2026-11-02

RESPONSE 05

Goldman Sachs (buyer)

INTERNAL GUIDANCE ON CLAUDE CODE AND GITHUB CODING TOOLS

Access limited to engineers and quantitative strategists after a review of 'productivity and risk management'; traders who used the tools to monitor client flows, model positions and draft materials lost access; spokesperson on the record (Bloomberg); grade 4

Recorded as an incumbent response because the enterprise buyer is the incumbent here: the bank ran a non-engineering agent seat for months, reviewed it for 'productivity and risk management' and withdrew it, reversing a July endorsement of non-technical staff building prototypes. It is the first disclosed seat narrowing by a regulated firm in this series and the direct counter-evidence to the vendors' agent-adoption metrics. Pricing model defended: none; the buyer is defending its own seat count.

SOURCE: BLOOMBERG (VIA BLOOMBERG LAW); QUARTZ

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.

2026-11-05

SIGNAL 01

Cognition

ENGINEERING

$0.25-0.75 PER PR (VENDOR-REPORTED)

Devin Review ships a configurable per-PR spend limit alongside an 82% cut in per-PR cost since August (typical $0.25-0.75; the full figures are in the engineering vertical entry)

The code-review category now has a cap as a product feature, which is the engineering vertical's version of the resolution-priced support category. Every number is Cognition's and none is independently measured. Agent Techniques carries the harness-cost study the price cut rests on and the per-bug figure.

SOURCE: DEVIN BLOG

2026-11-04

SIGNAL 02

Conductor (acquired by Meta)

ENGINEERING

UNDISCLOSED

Parallel-agent orchestration tool Conductor joins Meta; no new sign-ups from November 4, Conductor Cloud read-only from November 13 and shut down with data deletion from November 20, local app supported only until early 2027

A sixteen-day window from announcement to data deletion is the shortest export window the series has recorded for an acquired developer tool, and it falls on teams that adopted parallel-agent orchestration early. The pattern (orchestration layers absorbed by platform owners) is the same one Agent Techniques tracks at the harness level; the buyer-side rule is that any orchestration tool without a self-hosted path carries acquisition risk measured in days.

SOURCE: CONDUCTOR ANNOUNCEMENT; THE REGISTER

2026-11-05

SIGNAL 03

Plaid (with Anthropic)

FINANCE

UNDISCLOSED

Plaid supplies the account-linking layer for Personal Finance in Claude, its second frontier-lab integration after OpenAI's in May; read-only, US banks only, 24 months of transactions

The aggregator becomes the data layer for both consumer AI finance surfaces, which is a stronger position than either lab's: the assistant is substitutable, the connector is not. Plaid's own survey figure on consumer willingness (69%) is vendor-reported. The finance row watches whether the same connector reaches an enterprise surface (Cowork, Claude Code), which Anthropic has explicitly excluded for now.

SOURCE: PLAID; ANTHROPIC

PRICING SHIFTS

Seat to outcome, one move at a time.

4 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.

2026-11-04

SHIFT 01

HubSpot

Seat-based → Hybrid

Phase 2 pricing replaces hubs with a go-to-market seat, lowers the Starter-to-Pro step and consolidates all usage-based AI features into one credit; piloted in the Nordics and Benelux (ASPs down, Pro-plus counts up, credit consumption 2-3x other regions, all HubSpot-reported) and extended to all new EMEA customers in Q4. Active core seats are up 34% year on year after seat minimums were removed, which is the seat count rising as the seat price falls and the meter takes the growth.

SOURCE: HUBSPOT Q3 2026 CALL TRANSCRIPT; PREPARED REMARKS

2026-11-02

SHIFT 02

Microsoft

Seat-based → Hybrid

The Copilot Chat entitlement included in Business and Enterprise seats gains unpublished fair-use limits across additional features, with a purchase or licence-request path at the limit. The seat is unchanged; the included feature becomes a consumption-gated on-ramp whose quantity the vendor sets and does not disclose. Dated to the November 2 rollout start; the larger-tenant wave is November 9.

SOURCE: MICROSOFT SUPPORT; COMPUTERWORLD

2026-11-04

SHIFT 03

OpenAI

Seat-based → Hybrid

Persistent dots become generally available to Enterprise, Edu and Healthcare workspaces and billable by default from November 18 unless admins restrict membership during the preview; concurrently ChatGPT and Codex usage is being pooled into one allowance for selected users. The seat remains the contract; a default-on agent now draws on its allowance between sessions.

SOURCE: OPENAI HELP CENTER; CHATGPT DOCS

2026-11-02

SHIFT 04

Anthropic (S-1/A customer metrics; first-entry classification, not a shift)

Hybrid → Hybrid

The amended S-1 restates the mix through September 30: consumption revenue $26.3B (79%) against subscription $6.9B (21%) for nine months, more than 11,900 customers at $100K-plus trailing revenue (more than 1,900 at $1M-plus, about 250 at $10M-plus) generating $24.5B or about 74% of nine-month revenue, net revenue retention of about 900% stated to exclude Pro, Max and Ultra subscribers, and Opus 5.5 reducing customer cost per task about 40% against Opus 5 (company-reported). This is the first Anthropic entry in the pricing table and it records a classification, not a shift: the model is hybrid (consumption plus subscription) and the consumption share moved from 80% of H1 revenue (W44) to 79% of nine-month revenue, a point toward subscription, so no from-to move is claimed. The Weekly owns the revenue, loss and commitment figures.

SOURCE: SEC EDGAR (ANTHROPIC S-1/A NO. 1)

VERTICAL SCORECARD

Who leads each vertical.

Leader-vs-challenger by vertical, useful for procurement shortlists when matching workload to vendor cohort. As of Nov 7, 2026.

VERTICAL	LEADER	CHALLENGER	READ
Commerce	Amazon Seller Assistant	Sierra-led Poppy standard partners (Meta, Shopify, Stripe, Walmart)	Held, with the row's thesis now quantified by two earnings calls: Shopify's AI-channel orders up 3.3x on a base its president calls 'still small', and Uber's four rides and about 500 Eats orders in a 320-million-trip week. Both incumbents price access (Catalog, Shop Pay, 'terms of trade') rather than ceding checkout; Shopify is a Poppy member and its Catalog is the kind of structured access the standard would formalise. Amazon did not update Seller Assistant in the window.
Engineering	Codex (on GPT-6.1 Terra and Sol)	Claude Code (on Opus 5.5)	Held, with the first buyer-side narrowing on record: Goldman Sachs limited both Claude Code and GitHub's tools to engineers and strats after months of wider use, which cuts against the seat-expansion story for both leader and challenger. Devin Review's $0.25-0.75 per PR (vendor-reported) sets a published floor for the review sub-category; Conductor's absorption into Meta with a sixteen-day data window is the orchestration layer's risk made concrete; dots on Terra extend Codex's model into a persistent Enterprise surface. The Model Pulse carries the SWE-2 Flash model and its withheld weights.
Healthcare	Abridge	Suki	Held; no healthcare package shipped in the window. OpenAI's dots are HIPAA-eligible in covered Healthcare workspaces, which puts a persistent agent surface inside the row's compliance perimeter without a clinical product attached to it.
Support	Sierra	Salesforce Agentforce	Held. HubSpot's Customer Agent passed 2 million resolutions in the quarter (vendor-reported) under the $0.50-per-resolution list that anchors the low end of the row's price field; the company's 101% net revenue retention is the reminder that resolutions and retention are different lines. Sierra and Salesforce had no disclosures in the window; the Flex Credit multiplier for headless interactions remains unpublished (W44 watchlist, carried).
Sales	ZoomInfo	Salesforce Agentforce with Docusign IAM	Held on the standing rule (a change needs a second datapoint), with HubSpot entering the field: 20% of Pro-plus customers using agents, prospecting-agent users reporting 76% more qualified leads and 80% more

VERTICAL	LEADER	CHALLENGER	READ
			meetings, and the Phase 2 seat-plus-credit structure rolling to EMEA, all HubSpot's figures. HubSpot becomes the challenger if a second quarter shows the agent cohort moving net revenue retention, which this one did not.
Security	**Codex Security Cloud**	AWS Continuum	Held. The week's security events were model-side and threat-side rather than package-side: GPT-6 Cyber to approved Daybreak Red organisations (The Model Pulse) and Google Threat Intelligence and Mandiant's LLMjacking and Microsoft's ART EX disclosures (Agent Techniques). Neither changes the row; both raise the price of not having a package in it.
Legal	**Harvey**	Legora	Held; no legal package or customer disclosure in the window.
Research	**Claude Dashboards**	Gemini agent	Held, with OpenAI's dots entering the field as a persistent research-and-monitoring surface for Enterprise workspaces from November 18 by default. The row changes when dots carry a named deployment or a usage figure; at GA they carry an admin preview and a billing date.
Finance	**SAP Joule**	Workday agentic finance	Held. Anthropic's Personal Finance hub is a consumer product excluded from Team and Enterprise and from every agentic surface, so it does not enter the finance-function row; it is recorded for the boundary it draws. Palantir's US commercial print is an operations event with a finance-procurement reading (shorter durations).

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 unpublished meter: the quantity on the invoice is set after the purchase and is not visible to the administrator

- Microsoft Copilot Chat fair-use limits (numeric limits not published; admins cannot view)
- OpenAI dots (default-on and billable from November 18 unless membership is restricted)
- HubSpot single AI credit (consumption rate learned from the pilot: 2-3x)
- Zendesk resolution allowance with a pause (W44)

Last week's pattern was the vendor ceasing to report the seat; this week's is the vendor declining to publish the meter. Microsoft's support page says in terms that numeric limits are not shared and that administrators have no view of them; OpenAI's preview period is the only interval in which an admin sets the dots population before billing starts; HubSpot has disclosed the ratio of credit consumption between cohorts and not the amount. The common design is a meter whose quantity the vendor controls and the buyer discovers by hitting it (The Register's phrase for the Microsoft case), which is rational for a vendor whose revenue is now consumption and whose interest is in the buyer consuming before the buyer budgets. The buyer-side control is procedural rather than technical: ask for the limit in writing at renewal, set agent membership during every preview window, and treat any credit whose consumption rate the vendor reports as a multiple rather than a number as unpriced.

SOURCE: MICROSOFT SUPPORT; OPENAI HELP CENTER; HUBSPOT Q3 CALL TRANSCRIPT

PATTERN 02

Agent-originated demand is measurable to the unit and immaterial to the quarter, and the incumbents price access rather than cede checkout

- Uber (4 rides, ~500 Eats orders of ~320M in a week; eight assistants negotiating 'terms of trade')
- Shopify (AI-channel orders +3.3x, 'still small'; Catalog and Shop Pay as the priced surfaces)
- Sierra-led Poppy standard (W44; Shopify a member)

Two marketplaces answered the disintermediation question with numbers in the same week and the numbers agree: the agent channel is real, growing fast from a base near zero, converting well where the incumbent structures the data for it, and not yet large enough to move revenue. Both companies responded the same way, by selling the assistants access to the catalogue, the wallet and the transaction terms, and by building an in-house agent that outperforms the third-party ones on basket size. For a retailer or marketplace the operating conclusion is that the agent channel is a distribution partnership to be priced, not a checkout to be defended, and that the structured-catalogue investment (Shopify's billion product pages) is what converts agent traffic. The commerce row's challenger, the Poppy standard, is the attempt to make that access a protocol rather than a bilateral deal.

SOURCE: UBER Q3 2026 CALL TRANSCRIPT; SHOPIFY Q3 2026 CALL TRANSCRIPT

PATTERN 03

The buyer narrows the seat after a review, and the central bank's number is the context

- Goldman Sachs (Claude Code and GitHub tools to engineers and strats only)
- Bank of England (firms report ~0.5% productivity gain over three years against ~0.9pp a year expected)
- HubSpot (net revenue retention 101%; net new ARR to decline on budget sensitivity and longer sales cycles)
- Palantir (US commercial bookings +11% against +153%, shorter durations; net dollar retention 165%)

The enterprise-return gap moved this week from surveys to policy. Goldman is the first regulated firm on record to withdraw an agent seat from a non-engineering population after a review it describes as covering productivity and risk management; the Bank of England's panel puts realised AI productivity at an order of magnitude below expected (0.1-0.2 points a year realised against 0.9 expected) and attributes the measurable growth to the supply side; HubSpot's retention fell, on causes it names as budget sensitivity and longer sales cycles, while its agent metrics rose; Palantir's buyers are signing shorter even as net dollar retention rose to 165%. None of these contradicts the adoption metrics the vendors publish, because the vendors measure usage and the buyers measure return, and the week's lesson is that the two can diverge for a long time before the buyer acts. The board-level reading is to budget agent seats by population against measured return rather than by availability, which is the shape of what Goldman's review produced, and to expect vendors to answer with meters that make the population harder to see.

SOURCE: BLOOMBERG; BANK OF ENGLAND MONETARY POLICY REPORT; HUBSPOT AND PALANTIR Q3 CALLS

PATTERN 04

Headcount for AI efficiency, second week: the severance line is where the consumption transition is being paid for

- ServiceNow ($167M severance year to date against $53M; third 2026 round, number undisclosed)
- SUSE (~10% of ~1,900 staff, 'automation and AI')
- HubSpot ($65-75M plan charge, mostly Q4; October reorganisation with no AI attribution, listed for the cost shape only)
- Microsoft (~60,000 eligible for voluntary exits, W44)

Three application vendors in two weeks (ServiceNow, SUSE, Microsoft) have taken headcount actions attributed in their own words to AI efficiency, and the filed part of each is the severance line rather than the efficiency claim: ServiceNow's 10-Q shows $167M against $53M a year earlier while management claims more than $1B of annualised gains it does not itemise, and SUSE cuts a tenth of its staff a month after an under-subscribed voluntary offer. HubSpot's $65-75M charge for an October reorganisation, mostly in Q4, is listed alongside because it lands in the quarter its retention fell, but HubSpot attributes it to no AI cause and this issue does not either. The pattern is consistent with a seat-to-meter transition in which the vendor's revenue per employee has to rise to carry the compute bill, and it is the cost-side companion to the unpublished meter: the buyer sees less of the quantity and the vendor employs fewer people to sell it. The claims of efficiency are the vendors'; the severance is in the filings.

SOURCE: SERVICENOW 10-Q (VIA CTECH); HUBSPOT Q3 CALL; THE REGISTER (SUSE)

PATTERN 05

The data boundary is the product boundary: a frontier lab ships a connector and names the agent surfaces it will not serve

- Anthropic Personal Finance (Plaid; excluded from Team, Enterprise, scheduled tasks, Cowork, Claude Code)
- OpenAI dots (UAE residency workspaces excluded; HIPAA-eligible in covered workspaces)
- Plaid as the shared connector for both labs' consumer finance surfaces

Anthropic's finance hub is notable less for what it does than for the list of what it cannot touch: the linked accounts are unreachable from the surfaces that act autonomously, and the feature is withheld from every business plan. OpenAI's dots carry the mirror image, a residency exclusion and a HIPAA eligibility that define where the persistent agent may run. Both are the labs drawing the perimeter between data a chat surface may read and data an agent may act on, in product terms rather than policy terms, and both put the connector vendor (Plaid, in the finance case) in the position of serving multiple assistants from one integration. For an enterprise architect the exclusion lists are a template: the surfaces that act should have a smaller data perimeter than the surfaces that answer, and the vendor's own product boundaries now say so.

SOURCE: CLAUDE HELP CENTER; OPENAI HELP CENTER; PLAID

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 9

WATCH 01

Microsoft Copilot Chat fair-use limits reach tenants above 2,000 seats

The larger-tenant wave puts the unpublished caps in front of every enterprise administrator; the first disclosed limit (from a user notice, a support ticket or a Message Center post) is the datapoint that turns the meter from hidden to known, and the add-on conversion rate Microsoft reports next quarter is the test of whether the on-ramp works.

NOVEMBER 18

WATCH 02

OpenAI dots billing begins by default in Enterprise, Edu and Healthcare workspaces

The end of the admin preview is the moment a persistent agent starts drawing on Enterprise allowances for every eligible member whose admin did not restrict access; the first month's usage reports will show whether default-on agents consume at the rate the pooled allowance assumed.

NOVEMBER 13 AND NOVEMBER 20

WATCH 03

Conductor Cloud read-only, then shutdown and data deletion

A sixteen-day export window for an orchestration tool's cloud data; teams that do not export by November 20 lose it. The acquisition's terms and Meta's intended use of the product are undisclosed.

DECEMBER 4

WATCH 04

SWE-2 Flash free period ends in Devin; enterprise 25% discount ends November 18

The $0.10 / $0.50 list goes live inside Devin only; whether Cognition opens an API or a partner weights tier decides whether the price enters the model market or stays a Devin feature (The Model Pulse carries the model).

Q4 (REPORTED FEBRUARY)

WATCH 05

HubSpot EMEA Phase 2 pricing cohort

The first full quarter of seat-plus-credit pricing for all new EMEA customers; the figures to look for are credit consumption per customer in currency rather than as a multiple, and whether net revenue retention moves off 101%.

30 DAYS BEFORE METERING BEGINS

WATCH 06

Salesforce Headless Platform Interaction Flex Credit multiplier

Carried from W43 and W44: still unpublished after three weeks of the pattern shipping; it remains the price of cross-platform agent traffic and the test of whether a platform incumbent can toll agents it does not originate.

WEEK OF NOVEMBER 9-20

WATCH 07

Anthropic roadshow disclosures on customer metrics

The S-1/A's 11,900-plus $100K customers, 79% consumption mix and ~900% net revenue retention (excluding consumer subscribers) are as of September 30; any roadshow or further amendment that updates them to October is the first read on whether the Goldman-style narrowing shows up in a frontier lab's cohort data.

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 SCORECARD	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 29

Week 45 of 2026 · November 7, 2026

WEB

brianletort.ai/industry/applica

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