

THE BIG READ

The enterprise agent bill became a three-layer stack in one week, and the regulator arrived two days later with remedies aimed at the seat

KEY SIGNALS THIS PERIOD

8

Verticals tracked this period

8

Vertical packages shipped

4

Incumbent SaaS responses

3

Startup signals

AUTHOR

Brian Letort
BrianLetort.AI

PUBLISHED

October 17, 2026
Issue 26 · 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.

An enterprise that runs agents on Workday, Microsoft 365 or ChatGPT Enterprise will, from this week's notices, pay on three layers at once: a seat, a meter and an agent identity. Three vendors with three business models published that structure in the same four days. Workday's October 9 Flex Credit rate card prices about 30 actions on credits on top of the per-employee subscription (750 credits per recruiting requisition is the largest line) and tolls third-party agents per tool call; Microsoft's October 14 Partner Center notice puts a default-on consumption meter under the Copilot seat from December 1, with a \$40 per-user monthly ceiling that admins can move, and on October 15 a per-agent licence for agents that hold their own identity in Entra (Microsoft's identity directory); OpenAI's Help Center rate card, dated October 16, replaces per-message Chat pricing with per-token rates on top of contracted seats, including a law-firm-only Astra Law tier. The vertical entries below are the fact homes for each card. Last week's issue documented the toll a system of record charges at the API door. The effect of this week's three cards, read together, is that the structure resists cross-vendor comparison: the seat carries the price history a buyer can benchmark, the meter carries the usage none of the three has given a forward price for, and the agent identity carries a budget line nobody has a prior for. Whether that is design or three vendors independently copying usage-based cloud billing is not shown; Microsoft's own Copilot Credits Guide (September 2026) frames consumption as 'complementing subscription-based services', so the seat-plus-meter half is the vendor's framing, and the agent-identity third layer and the cross-vendor reading are this issue's. Then, on October 16, the UK Competition and Markets Authority proposed Strategic Market Status (SMS) for Microsoft's business software. The proposed perimeter names the surfaces: Microsoft 365 Copilot, Copilot Chat, Cowork, parts of Security Copilot, and (in the roadmap's footnote 67) the Graph API, Work IQ, the Admin Center and Agent 365 where they underpin, manage or deploy functionality, and paragraph 4.18(a) records rivals' concern about the fees they may incur for accessing those APIs. So the meter's surface (Work IQ APIs are what the default-on Copilot Credits bill) and the agent-identity surface (Agent 365) are inside the perimeter. What the roadmap of possible interventions does not yet do is price them: its remedy

vocabulary is parity access for rival assistants, prominence and defaults, minimum API-change notice, licensing transparency and unbundling AI features from suites, which is the language of seats, features and suites that twenty years of suite-bundling precedent supports. No remedy as drafted prices a per-action credit or a per-agent licence, and the parity-access clause is the hook under which one could; that is this issue's reading of the remedies' precedent, not a fact about the document's language. A procurement team should not expect a Copilot unbundling order in February 2027, if one comes, to lower the agent bill by itself, and the consultation closing on November 13 is the window in which to ask that parity be written against the Work IQ meter and the agent account.

The second thing the week settled is which way the front door is moving. Harvey and Legora launched rival 'Intake' products a day apart (October 13 and 14) that let any employee submit legal work from Slack or Teams without holding a licence, with agents completing pre-approved requests such as standard NDAs before a lawyer sees them; Salesforce won an up-to-£2 billion, 15-year HMRC contract that puts Agentforce in front of 34 million taxpayers under a tender condition that core capabilities including AI be first-party Salesforce software; Docusign made its agreement platform a tool inside Agentforce and Slack rather than a standalone agent surface; and HCLTech stood up a dedicated Anthropic business unit with forward-deployed engineers. The common move is that the licence is attaching to the system that holds the workflow (the legal team's policy engine, the CRM, the systems integrator's delivery contract) and detaching from the person who types the request. Suki's \$90M Series E, raised to build post-visit agents for referrals and prior authorizations, says the same thing from healthcare: the note is table stakes and the revenue is in the workflow steps after it. The week's capital colour ran alongside rather than through this story: two sponsor-led take-privates of cash-generative software on the same day (Varonis at about \$6.3B, CCC at about \$6.1B enterprise value, 44-52% premiums to unaffected prices, meaning the share price before deal reports) are private equity buying installed bases; neither target appears in the agent-pricing story, and they are carried here as context, not as a pricing-model datapoint.

Three things to do this week. Put a line in every Workday, Microsoft and OpenAI renewal that fixes the meter rate (credits per action, dollars per credit, tokens per task) for the term, because all three published units and none published a forward price. Count the agent identities you will need before the per-agent licence shows up as a line item, because Copilot Studio's agent accounts and Workday's Agent-Ready Tools both bill them separately from the people who supervise

them. And if you buy Microsoft business software in the UK, file a response to the CMA consultation by November 13 that names the Work IQ meter and the agent-account licence as the surfaces a parity remedy should be priced against, because the roadmap as drafted names them and does not yet price them.

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.

2026-10-13

PACKAGE 01

Workday

HR

INCUMBENT SAAS

HYBRID

Workday Autopilot agents and Flex Credit rate card

Workday Rising prices every new agent, Agent-Ready Tool call and API overage on Flex Credits, with credit counts published and the price of a credit not

This is the fact home for Workday's meters. Workday's Flex Credits page carries the Self-Service Agent rates (1 credit for a retrieval, 5 for an autonomous completion); the fuller October 9 rate card of about 30 meters is reproduced in StackOne's independent guide, which is the source for the rest: Recruiting Agent 750 credits per requisition, Contract Negotiation Agent 500 per document, third-party agents at 10 credits per 100 basic Agent-Ready Tool calls and 50 per 100 advanced, and 60 credits per 10,000 API calls above a 2.5-6.5 million annual allowance once the waiver ends on January 31, 2027. Every subscription includes a complimentary pool of 15,000-200,000 Flex Credits a year by headcount with no rollover, which is the buyer's first arithmetic: the top of that pool covers about 266 recruiting requisitions at 750 credits each before any purchase, and the bottom covers 20. Sana is the new Workday Home, recruiting autopilot is with early adopters now with general availability in early 2027, and payroll, accounting and deployment autopilots follow. Workday's console alerts at 80, 90 and 100% of the pool with no hard stop, and CEO Aneel Bhusri called consumption pricing 'a new muscle for us'. If you integrate with Workday, take three numbers into the renewal: actions per hire and per close on the current process (instrument them now; the console reports credits, not actions), the dollar price per credit (not on any public page), and a hard cap, because the card's unit economics cannot be forecast without the first and cannot be bounded without the third.

SOURCE: WORKDAY INVESTOR PRESS RELEASE; WORKDAY FLEX CREDITS PAGE; STACKONE RATE-CARD GUIDE

2026-10-16

PACKAGE 02

OpenAI

OTHER

FRONTIER LAB

HYBRID

ChatGPT Enterprise token-based rate card

OpenAI replaces per-message Chat pricing with one per-token rate card across Chat, Codex, Office add-ins and Workspace Agents, charged on top of contracted seats

This is the fact home for OpenAI's enterprise prices. Per million tokens (input / cached / output): GPT-6 Sol \$2.00 / \$0.20 / \$10.00; GPT-6 Astra \$10.00 / \$1.00 / \$50.00; GPT-6 Luna \$0.10 / \$0.01 / \$0.50; GPT-6 Astra Law \$12.50 / \$1.25 / \$62.50 (Trusted Access, selected US law firms); GPT-Rosalind-Research \$5 / \$0.50 / \$25 (select Enterprise plans). Add-ons: web search \$10 per 1,000 runs, voice \$0.05 per minute plus tokens, context above 272K at 2x input and 1.5x output, regional data residency 1.1x, Fast mode 2x, and Ultrafast 6x on the two models the card lists it for (GPT-6 Astra and GPT-6.1 Sol). Usage bills on top of seat fees, one agentic pool covers Chat, ChatGPT Work, Codex, Word, Excel, PowerPoint and Workspace Agents, legacy credit contracts stay on the old card until migration, and GPT-5.6 Sol promotional pricing runs at least through November 21. It is the first full public price list for a frontier lab's seat-plus-usage hybrid, and it introduces a vertical price (Astra Law at 1.25x Astra) inside a horizontal product. A team migrating from per-message to per-token should pull last quarter's message counts by model and multiply by measured tokens per message before signing, because the card changes the unit and a speed multiplier applies to the whole token bill for the request, not the increment.

SOURCE: OPENAI HELP CENTER (CHATGPT RATE CARD; TOKEN-BASED BILLING)

2026-10-14

PACKAGE 03

Microsoft

OPERATIONS

INCUMBENT SAAS

HYBRID

Microsoft 365 Copilot standalone SKUs, default-on usage billing and Copilot Studio agent accounts

Copilot decouples from the M365 base subscription on December 1, gains a default-on 4,000-credit per-user monthly meter the same day, and agents get their own Entra identities and licences

This is the fact home for Microsoft's three layers. From December 1 new Copilot and Copilot Business SKUs can be bought without a qualifying M365 subscription at the same price (existing offers become '(add-on)' from October 16 and leave sale through the Cloud Solution Provider (CSP) partner channel on December 1); default-on usage billing for new Copilot Business CSP licences moved from November 2 to December 1 with a default spending limit of 4,000 Copilot Credits per user per month, a \$40 ceiling at the published \$0.01 per credit that admins can adjust and that a 100-user tenant could reach at 400,000 credits, or \$4,000, a month, covering Copilot Cowork, Work IQ APIs and GitHub Copilot Harness; and Copilot Studio's October 15 Frontier preview gives agents their own Entra accounts (Entra is Microsoft's identity directory) in Teams, Outlook and Office comments, each requiring a Microsoft 365 Frontier for AI Teammates licence plus metered credits, with business-process memory shared across users with no privacy boundary and no opt-out during preview. For a tenant administrator the decoupling is the CMA-relevant change (it is the unbundling the roadmap names, done voluntarily two days ahead of the proposed decision), and the default-on meter is the budget-relevant one: the \$40 ceiling is per user per month and applies to every new Business licence unless an admin lowers it before December 1.

SOURCE: MICROSOFT LEARN (PARTNER CENTER ANNOUNCEMENTS, OCTOBER 2026); MICROSOFT COPILOT STUDIO BLOG

2026-10-14

PACKAGE 04

Harvey; Legora

LEGAL

STARTUP

UNKNOWN

Harvey Intake and Legora Intake

The two best-funded legal AI vendors launch seat-less front doors for in-house teams a day apart, with agents completing standard NDAs before a lawyer sees the request

Legora's Intake went live in Slack on October 13 (Teams, email and MCP, the Model Context Protocol tool interface, to follow); Harvey's followed on October 14 in early access, designed with Carvana, GE Aerospace and Dentsu. Both capture requests in the channel the employee already uses, route them to an approved agent, workflow or lawyer under rules the legal team sets, return the answer in the same channel, and let business users submit and track requests without a Harvey or Legora licence. Harvey says it serves 1,600 in-house teams including five of the Fortune 10; Legora says the median in-house lawyer now supports 367 employees, up 26% in a year (both vendor-reported). The pricing consequence is that the unit of value moves from the lawyer's seat to the request the policy engine handles, and the competitive one is that the \$15.5B and \$5.55B vendors are now in the corporate legal department, where contract lifecycle management (CLM) incumbents and Docusign's Agentforce integration (below) meet them from the CRM side. If you run the Slack or Teams workspace these products plug into, the list of requests pre-approved for agent completion is now product configuration and the thing a renewal will be priced on; get it into a versioned, reviewable artifact while it is still being written in early access.

SOURCE: ARTIFICIAL LAWYER; HARVEY PRODUCT PAGE; LEGORA NEWSROOM; LAW360 PULSE

2026-10-15

PACKAGE 05

Salesforce

SUPPORT

INCUMBENT SAAS

UNKNOWN

Agentforce for HMRC (Enterprise CRM, up to £2 billion over 15 years)

Salesforce wins a 15-year HMRC contract that puts Agentforce in front of 34 million taxpayers under a tender that requires first-party Salesforce AI

The contract (period July 31, 2026 to April 30, 2036 per the UK Find a Tender record, extendable to 2041) bundles Agentforce, Public Sector Solutions, Marketing Cloud and Data 360 to give caseworkers a single taxpayer view, with agents triaging inbound queries, automating routine administration and routing complex cases while staff keep oversight of every decision; Quantexa is a delivery partner and the target, per Salesforce's release, is 90% digital service delivery by 2030 from about 76%. Two things matter for other buyers. The tender's requirement that core capabilities including AI be first-party Salesforce software is a procurement condition that locks the agent layer to the system of record for fifteen years, which is the strongest form of the pattern this issue tracks. And the scale (34 million individuals, 5 million businesses) makes it the largest disclosed Agentforce deployment, which puts the support row's challenger slot on a public-sector workload rather than a commercial one. Pricing is a ceiling (up to £2 billion), not a rate; the per-conversation economics will not be visible until HMRC's transparency data publishes.

SOURCE: SALESFORCE NEWSROOM; TENDERLINE (UK CONTRACT AWARD RECORD)

2026-10-13

PACKAGE 06

Suki

HEALTHCARE

STARTUP

UNKNOWN

Suki post-visit agents (referrals, prior authorizations, orders)

Suki raises \$90M to charge for the workflow steps after the note, with encounters up more than 190% year over year

The Series E (Healthier Capital, Crosslink, Hercules; all existing institutional investors participating) takes total funding past \$250M and is earmarked for agents that complete work after the visit (referrals, prior authorizations for new medications, order filling, follow-up scheduling) and for deeper embedding with 30-plus partners including Optum, Zoom and athenahealth rather than new ones. Suki reports encounters up more than 190% in 2026, revenue doubling in 2025 and on track to double again, partner revenue up more than tenfold in two years and 450-plus customer organizations (Fierce reported 250-plus); all vendor-reported. The message is the same as last week's Abridge and Suki notes with the pricing consequence made explicit: ambient capture is table stakes, and the vendors (Suki, Abridge, Ambience, OpenEvidence) will price the next contract on completed downstream steps, which have a measurable dollar value in prior-authorization labour that the note never had. If you integrate a scribe vendor, the billable event is about to move from the note to the completed downstream step (referral created, authorization submitted, order filled), so instrument those steps with an attribution both you and the vendor can audit before the contract changes unit.

SOURCE: SUKI PRESS RELEASE; HEALTHCARE DIVE; FIERCE HEALTHCARE

2026-10-13

PACKAGE 07

DocuSign

SALES

INCUMBENT SAAS

SEAT

DocuSign Intelligent Agreement Management inside Salesforce, Agentforce and Slack

DocuSign chooses to be a tool inside Agentforce and Slack rather than a standalone agent surface for contracts

Sellers can ask Agentforce about agreements and generate or send contracts in the CRM; legal and contract teams review, track renewals and see post-execution analytics there; and in Slack a team can request a contract, pull live pricing, apply approved terms and route to legal in one thread. Salesforce's Bill Patterson framed it as 'people and AI agents' executing agreements. The decision an e-signature and contract lifecycle management (CLM) incumbent took here, to surface through the CRM vendor's agent rather than compete with it, is the opposite of Harvey's and Legora's (above), who built their own front door into the same in-house contract workflow; the two approaches will meet on the same NDA, with the legal-side product completing it under the legal team's rules and the sales-side product completing it under the CRM's. Whoever owns the Salesforce and Slack integrations should decide which system's approval rules govern a standard contract before both agents are live, because the integration gives each a path to send one. No pricing was disclosed; the integration runs on existing DocuSign IAM (Intelligent Agreement Management, its agreement platform) and Salesforce licences.

SOURCE: DOCUSIGN NEWS CENTER

2026-10-12

PACKAGE 08

HCLTech

ENGINEERING

INCUMBENT SAAS

UNKNOWN

HCLTech Anthropic Business Unit; Claude and Claude Code in BigFix, AppScan and AI Force

HCLTech stands up a dedicated Anthropic business unit with forward-deployed engineers and embeds Claude Code in its own software line

The Global Premier Partnership (press release and Indian exchange filing, the morning of HCLTech's Q2 results) creates an Anthropic Business Unit and Center of Excellence with sandboxed Claude environments delivered by HCLTech forward-deployed engineers, makes HCLSoftware an Anthropic ISV partner embedding Claude and Claude Code into BigFix and AppScan, and puts Claude Code into the AI Force platform for IT support, code remediation and digital-employee services with human oversight for high-impact changes. It extends the pattern of frontier labs distributing through global systems integrators rather than direct sales alone (Anthropic's Daniela Amodei pointed to HCLTech's bank, hospital, telco and manufacturing estates); the twist is that the SI is also an ISV, so the same deal is a distribution channel and a product embed. For a CIO already on BigFix or AppScan, Claude Code arrives inside a product that is already licensed, which is the quietest way a frontier model enters a regulated estate and the one a model-vendor allowlist is least likely to catch. No commercial terms were disclosed.

SOURCE: HCLTECH PRESS RELEASE; PR NEWSWIRE; INDIAN EXPRESS / PTI (EXCHANGE FILING)

INCUMBENT RESPONSES

How the SaaS estate is answering.

4 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.

2026-10-16

RESPONSE 01

Microsoft

MICROSOFT 365 COPILOT (STANDALONE SKUS, USAGE BILLING, AGENT ACCOUNTS) AND THE CMA PROPOSED SMS DECISION

Microsoft publishes Copilot decoupling and default-on metering on October 14, agent accounts on October 15, and the CMA proposes Strategic Market Status (SMS) on October 16 with Copilot, Copilot Chat, Cowork, Agent 365 and the Work IQ APIs inside the proposed perimeter

The sequence is the point: the vendor unbundled the seat voluntarily two days before the regulator proposed to require it, and in the same notice moved the agent margin onto a consumption meter and a per-agent licence. The CMA's perimeter reaches both surfaces (the roadmap's footnote 67 names Work IQ and Agent 365; paragraph 4.18(a) records rivals' concern about API fees), but its roadmap of possible interventions (parity access, prominence and defaults, API-change notice, licensing transparency, unbundling AI features from suites) is drafted in the vocabulary of seats, features and suites, and no remedy yet prices a per-action credit or a per-agent licence; parity access is the clause under which one could be. Microsoft argued Copilot is a separate 'AI layer'; the CMA provisionally rejected that. Consultation closes November 13, the statutory decision deadline is February 13, 2027, and the CMA kept AWS and Microsoft cloud under informal monitoring rather than opening a cloud probe. UK tenants that want the meter priced in a remedy have one filing window; everyone else should note that the first binding remedy on an enterprise AI assistant, if it comes, is being drafted with seat-era language and will need the meter written in.

SOURCE: CMA PRESS RELEASE; CMA ROADMAP OF POSSIBLE INTERVENTIONS; MICROSOFT LEARN PARTNER CENTER ANNOUNCEMENTS

2026-10-13

RESPONSE 02

Workday

SANA AGENTS AND AGENT-READY TOOLS INSIDE CLAUDE (AND MICROSOFT 365 COPILOT, TEAMS, SLACK, GEMINI ENTERPRISE)

Workday and Anthropic put Sana agents and Workday Agent-Ready Tools inside Claude over MCP (the Model Context Protocol) with Workday permissions enforced on every action, early access now and general availability late 2026; Workday names OpenAI, ElevenLabs and TypeSafe as additional partners and declines to commit to one model vendor

Employees can approve invoice exceptions and analyze HR and finance data in Claude, and Sana calls Claude for steps in compensation cycles and the financial close without data leaving Workday; admins 'connect once' with no new permission model. Read with the rate card, every Agent-Ready Tool call from Claude meters on Flex Credits, so the partnership is the toll's distribution deal: the system of record keeps the permission boundary and the meter, and lets the frontier lab own the chat surface. A buyer gets a model-neutral stance from Workday and a clear line of accountability (Workday's audit trail covers the action, whichever model requested it); the cost is that the same action can now be triggered from five surfaces, each of which bills the same credit pool, which is why the console's 80/90/100% alerts need a hard cap behind them.

SOURCE: WORKDAY NEWSROOM; ITWIRE (BHUSRI, KAZMAIER AND HELLERMARK PRESS Q&A)

2026-10-15

RESPONSE 03

Salesforce

AGENTFORCE (HMRC
ENTERPRISE CRM;
DOCUSIGN IAM
INTEGRATION)**Salesforce books an up-to-£2 billion, 15-year public-sector Agentforce deployment with a first-party-AI tender condition, two days after making Docusign's agreement platform a native tool for Agentforce and Slack**

The CRM incumbent's response to agent competition is to become the system the agents are contractually required to run on. The HMRC condition that core capabilities including AI be first-party Salesforce software means no third-party assistant sits between the caseworker and the taxpayer record for fifteen years; the Docusign deal brings a category incumbent inside the same boundary as a tool. For a team that builds on Salesforce the HMRC terms are a template for what the vendor will ask for in large renewals (first-party AI as a contract term, not a product choice), which means a third-party model or agent runtime in the architecture becomes a contract question; the counter is a per-conversation rate for agent traffic that survives the term, because an up-to figure over fifteen years is a ceiling that leaves the unit rate open.

SOURCE: SALESFORCE NEWSROOM; TENDERLINE; DOCUSIGN NEWS CENTER

2026-10-15

RESPONSE 04

Thoma Bravo (Proofpoint);
GTCR and ElliottVARONIS (ABOUT \$6.3B
CASH) AND CCC
INTELLIGENT SOLUTIONS
(ABOUT \$6.1B
ENTERPRISE VALUE)
TAKE-PRIVATES**Two sponsor-led buyouts of public AI-adjacent software on one day: Proofpoint pays \$49.75 per share for Varonis (52% premium to the June 22 unaffected price, the share price before deal reports; about 8.4x trailing revenue and 53x trailing free cash flow); GTCR and Elliott pay \$7.75 per share for CCC (44% premium to July 9) with no financing condition. This entry is the fact home for both deals**

This is capital colour rather than a pricing-model datapoint: Varonis sells data security and CCC sells insurance-claims software, neither appears in the week's agent-pricing story, and 53x free cash flow is not a depressed multiple. What the deals show is private equity paying 44-52% premiums for installed bases of seat-based, cash-generative software in the same week the vendors above moved their agent margin onto meters and licences; whether either sponsor runs a seat-to-meter conversion under private ownership is the thing to watch (watchlist, 1H 2027), not a conclusion to draw now. For an engineer on either product the practical effect is an ownership change in the middle of the industry's pricing transition, which is the time to get a multi-year rate into the contract.

SOURCE: PROOFPOINT / VARONIS PRESS RELEASE; CCC 8-K EXHIBIT 99.1 (THE PRESS RELEASE ATTACHED TO CCC'S SEC FILING)

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.

2026-10-13

SIGNAL 01

Suki

HEALTHCARE

\$90M

\$90M Series E led by Healthier Capital with Crosslink and Hercules, total funding above \$250M, proceeds for post-visit agents and deeper embedding with 30-plus existing partners

The round is sized to build the billable steps after the note rather than to acquire new distribution, which is the clearest statement yet from an ambient-documentation vendor that the note itself no longer carries a price. Vendor-reported growth (encounters up more than 190%, revenue doubling two years running) is consistent with a product that is being given away inside partner EHRs (electronic health record systems) to earn the downstream workflow; expect the next contract to be priced per completed referral or authorization, and the integration work that pays off is the attribution for those steps (the healthcare entry above has the instruction).

SOURCE: SUKI PRESS RELEASE; HEALTHCARE DIVE

2026-10-14

SIGNAL 02

Harvey and Legora

LEGAL

Rival Intake launches within 24 hours (Legora October 13 in Slack; Harvey October 14 in early access), both letting employees submit and track legal requests without a licence and both completing pre-approved NDAs end to end

Two vendors with \$15.5B and \$5.55B valuations converging on the same seat-less front door in the same week is a market telling itself that the per-lawyer seat has stopped growing fast enough; the next contest is who holds the in-house team's routing rules. The matter for buyers is that the vendor who holds the rules holds the renewal, and the rules are being configured now in early access, before pricing for the request-based model has been published by either.

SOURCE: LEGORA NEWSROOM; HARVEY PRODUCT PAGE; LAW360 PULSE

2026-10-14

SIGNAL 03

Modal

ENGINEERING

\$750M

\$750M Series D at a \$15.75B post-money valuation led by Accel, five months after a \$355M Series C at \$4.65B; Modal says run-rate crossed \$1B in September and its Sandboxes product (isolated code execution for agents, about 3 billion sandboxes run) is at a \$300M run-rate on its own

The application layer's compute is being bought by the second rather than by the reserved cluster, and a 3.4x valuation step-up inside five months says investors believe the burst pattern (an agent fleet that spins up a sandbox per task) is the default shape of application demand; the run-rate figures are vendor-reported. For an engineering leader the relevance is budgeting: usage-based compute under a usage-based agent bill under a usage-based model rate card is three meters that move together, and none of the three vendors publishes the others' rates. The round's terms are carried here; the AI Stack Weekly's capital lens reads the round against the week's other financings.

SOURCE: MODAL BLOG; ACCEL; SILICONANGLE

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.

2026-10-16

SHIFT 01

OpenAI

Seat-based → Hybrid

ChatGPT Enterprise moves from per-message Chat pricing on top of seats to a per-token rate card on top of seats, with one agentic usage pool across Chat, Codex, Office add-ins and Workspace Agents, a vertical SKU (Astra Law) and speed multipliers (Fast 2x across the card; Ultrafast 6x on GPT-6 Astra and GPT-6.1 Sol only). Legacy credit contracts remain on the old card until migration. The 'other' vertical entry above has the card.

SOURCE: OPENAI HELP CENTER (CHATGPT RATE CARD)

2026-10-14

SHIFT 02

Microsoft

Seat-based → Hybrid

Microsoft 365 Copilot Business gains default-on usage billing for new licences sold through the Cloud Solution Provider (CSP) channel from December 1, with an adjustable per-user monthly default cap; Copilot itself becomes purchasable without an M365 base subscription at the same price; and Copilot Studio agent accounts each carry their own licence plus metered credits, adding a per-agent layer to the per-user one. The operations entry above has the figures.

SOURCE: MICROSOFT LEARN (PARTNER CENTER ANNOUNCEMENTS, OCTOBER 2026); COPILOT STUDIO PRICING

2026-10-13

SHIFT 03

Workday

Seat-based → Hybrid

Everything shipped at Rising bills through Flex Credits on top of the per-employee subscription: about 30 per-action meters, a per-call toll on third-party agents and an API overage meter from February 1, 2027, at an unpublished dollar price per credit, so the effective rate is set per contract. The HR entry above has the card.

SOURCE: WORKDAY FLEX CREDITS PAGE; STACKONE RATE-CARD GUIDE (RATES AS OF OCTOBER 9, 2026)

2026-10-14

SHIFT 04

Harvey; Legora

Seat-based → Unknown

Both Intake products let business users submit, track and receive completed legal work without holding a Harvey or Legora licence, which moves the billable unit away from the per-lawyer seat toward the request the legal team's rules route; neither vendor has published a price for the request-based model, and the seat remains the unit for the lawyers who configure the rules.

SOURCE: ARTIFICIAL LAWYER; LEGORA NEWSROOM

VERTICAL SCORECARD

Who leads each vertical.

Leader-vs-challenger by vertical, useful for procurement shortlists when matching workload to vendor cohort. As of Oct 17, 2026.

VERTICAL	LEADER	CHALLENGER	READ
Commerce	Amazon Seller Assistant	Personal Agent Protocol retail partners (Sierra, Meta, Walmart, Visa, Mastercard)	Unchanged. No commerce-specific release in the graded packet this week; PAP draft 0.1's design-partner set from W41 stands, Anthropic still absent from the working group.
Engineering	Codex (on GPT-6.1 Sol)	Claude Code (on Sonnet 5.5)	Unchanged. OpenAI's rate card folds Codex into the same per-token pool as Chat; HCLTech's embed puts Claude Code inside BigFix, AppScan and AI Force; Agent Techniques carries the week's harness releases (Claude Code 2.1.299-302, Codex 0.163, Copilot CLI, Cursor). Neither changes the slot order.
Healthcare	Abridge	Suki	Unchanged, with the challenger's thesis now funded: Suki's \$90M is earmarked for post-visit agents (referrals, prior authorizations, orders), which is the first explicit move by either vendor to price the workflow after the note. Abridge's UT System agreement from W41 remains the largest disclosed enterprise footprint.
Support	Sierra	Salesforce Agentforce	Challenger changes from Decagon to Agentforce on the HMRC award: up to £2 billion over 15 years, 34 million taxpayers and 5 million businesses, with a first-party-AI tender condition. Sierra holds the leader slot on commercial customer-service deployments and the Personal Agent Protocol (PAP); Decagon moves to the note on the strength of its W41 open-sourcing of PACT (Personal Agent Consent and Trust, its consent standard for consumer agents) under Apache 2.0.
Sales	ZoomInfo	Salesforce Agentforce with Docusign IAM (its agreement platform)	Challenger changes from HubSpot Breeze. Docusign's native integration lets Agentforce generate, negotiate and route contracts from the CRM and Slack, which closes the quote-to-signature loop HubSpot's outcomes-priced Breeze does not yet cover; HubSpot's 8-K restructuring from W41 stays in the note.

VERTICAL	LEADER	CHALLENGER	READ
Security	Codex Security Cloud	Elastic AlertZero	Unchanged. The security news was ownership rather than product: Varonis to Proofpoint (Thoma Bravo) at about \$6.3B, and HCLSoftware embedding Claude in BigFix and AppScan; the CMA proposed decision brings parts of Security Copilot inside the proposed Strategic Market Status perimeter. The slots wait for a product release.
Legal	Harvey	Legora	Unchanged in order; both moved. Rival Intake launches a day apart put each in the corporate legal department with a seat-less front door; OpenAI's Astra Law tier (\$12.50 / \$62.50, selected US law firms) is the first vertical SKU on a frontier lab's enterprise rate card and is noted as a third path into the same workload.
Research	Claude Dashboards	Gemini agent	Unchanged. OpenAI's GPT-Rosalind-Research (\$5 / \$25, select Enterprise plans) appears on the rate card without a product release; the Gemini agent's late-October or early-November general availability and metering details remain the pending event for the row.
Finance	SAP Joule (per-action AI Units)	Workday agentic finance (Flex Credits)	Unchanged in order; Workday's rate card and Claude integration (Sana approving invoice exceptions, calling Claude for close steps) make the challenger's pricing the better documented of the two this week, with credit counts public and the price per credit not. SAP's 0.02 AI Units per action from W41 stands.

ARCHITECTURE WATCH

Patterns to track.

4 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 three-layer agent bill: seat, meter, agent identity

Workday: per-employee subscription + Flex Credits per action + a per-call toll on third-party agents' Agent-Ready Tool use (the HR entry has the card)

Microsoft: Copilot seat (standalone from Dec 1) + default-on Copilot Credits per user per month + a licence per agent account (the operations entry has the figures)

OpenAI: ChatGPT Enterprise seat + per-token rate card with a vertical SKU and speed multipliers + one shared agentic pool across Chat, Codex, Office and Workspace Agents (the 'other' entry has the card)

SAP (W41): per-user packages retired for 0.02 AI Units per action, idle agents free

Three vendors with three business models published the same structure in four days, and the structure is more informative than any one rate. The seat layer carries the price history buyers can benchmark; the meter layer carries the usage that none of the three has given a forward price for (Workday publishes credit counts without a dollar per credit, Microsoft publishes a per-credit price without a credits-per-task rate, OpenAI publishes per-token rates for a product whose tokens per message buyers have never measured); and the agent-identity layer carries a budget line with no prior at all, because enterprises have never bought licences for non-human workers. The observable is that the three cards are mutually non-comparable; whether that is design or three vendors independently copying usage-based cloud billing is not shown, and Microsoft's own Copilot Credits Guide frames consumption as complementing the subscription. The practical defence does not depend on which: negotiate a blended unit (dollars per completed workflow, measured on your own top ten processes) that spans all three layers. The CMA's proposed perimeter names the meter's and the agent layer's surfaces (Work IQ, Agent 365) but its remedy roadmap prices neither yet; the regulator entry below has the detail. The fact homes for each vendor's figures are the vertical-movement entries above; this entry is the home for the structure.

SOURCE: WORKDAY FLEX CREDITS PAGE; MICROSOFT LEARN PARTNER CENTER ANNOUNCEMENTS; OPENAI HELP CENTER RATE CARD

PATTERN 02

Seat-less front doors: the licence attaches to the routing rules, not the requester

Harvey Intake and Legora Intake: any employee submits from Slack or Teams without a licence; the legal team's rules decide agent, workflow or lawyer; pre-approved NDAs completed end to end

Workday Sana in Claude, Microsoft 365 Copilot, Teams, Slack and Gemini Enterprise: the same approval from five surfaces, metered to one credit pool, permissioned by Workday

Docusign IAM in Agentforce and Slack: request, price, apply approved terms and route to legal in one thread

Copilot Studio agent accounts: the agent has the Entra (Microsoft identity directory) account and the licence; the human is a supervisor in a Teams thread

Last week's issue tracked where the toll sits; this week's launches show where the front door is moving. In four separate products the person who starts a task no longer needs the licence; the licence attaches to the system that holds the rules (the legal team's policy engine, the HR system's permission model, the CRM's approved terms, the agent's own identity). That inverts the per-seat assumption that revenue scales with headcount, and it explains why the same week's rate cards price actions and agents instead: the vendor needs a unit that grows when the requester count is unlimited. For the buyer the consequence is governance before economics. Whoever configures the rules (which requests an agent may complete unattended; which approvals a Workday action needs when it is triggered from Claude; which contract terms Agentforce may apply) is configuring the product, and in every case above that configuration is happening in early access, preview or a tender's first year, before the request-based price is published. Write the rules down as an owned artifact with a change log now; the renewal will be priced against them.

SOURCE: HARVEY PRODUCT PAGE; LEGORA NEWSROOM; WORKDAY NEWSROOM; DOCUSIGN NEWS CENTER; MICROSOFT LEARN (AGENT ACCOUNTS)

PATTERN 03

The regulator meets the meter: the SMS perimeter names the surfaces, the remedies are drafted for the seat

CMA proposed decision (Oct 16): Strategic Market Status (SMS) across productivity suite, PC OS, security software, server OS, relational database; Microsoft 365 Copilot, Copilot Chat, Cowork/Agents and parts of Security Copilot inside the designated activities, with the Graph API, Work IQ, the Admin Center and Agent 365 included 'where they underpin, manage or deploy functionality' (roadmap footnote 67); paragraph 4.18(a) records rivals' concern about API access fees

Roadmap of possible interventions: parity access for rival AI assistants, prominence and default rules, minimum API-change notice, licensing transparency, unbundling AI features from suites, scrutiny of security bundling and cloud-licensing restrictions; none priced against a per-action credit or a per-agent licence

Microsoft's own Dec 1 Copilot decoupling, published two days earlier, delivers the unbundling remedy voluntarily while adding a default-on meter and a per-agent licence

Consultation closes 5pm Nov 13, 2026; statutory decision by Feb 13, 2027; AWS and Microsoft cloud stay under informal monitoring

The first competition proceeding to put an enterprise AI assistant inside a designated perimeter arrived in the same week the assistant's vendor moved its margin onto a meter and an agent licence. The perimeter reaches both: footnote 67 of the roadmap puts Work IQ (whose APIs are what the default-on Copilot Credits bill) and Agent 365 (the agent-identity layer) inside the proposed activity, and paragraph 4.18(a) records rivals' concern about API fees. The remedies do not yet price either. The roadmap's interventions are the ones twenty years of suite-bundling law supports: give rivals equal access, stop defaulting, give notice of API changes, publish licence terms, sell the AI feature separately. Microsoft's December 1 standalone SKUs satisfy the last of those in advance, at the same price, and the Partner Center notice that announced them also announced the consumption meter and the agent-account licence. The reading here, which is this issue's reading of the remedies' precedent and not a statement about the document's language, is that an unbundling order in February 2027, if it comes, lowers the seat and leaves the meter and the agent licence where they are unless the parity-access remedy is written against the Work IQ meter rate and the agent account, which is the clause under which it could be. That is not a criticism of the CMA, which is working from a May 2026 investigation scope and has the surfaces in view; it is a timing fact a customer can act on. The consultation is open until November 13 and is the only public process in which a customer can put the meter on the record; the EU's GPAI powers, live since August 2, are the other binding lever and are aimed at model providers rather than at suite pricing. The AI Stack Weekly's policy coverage carries the US enforcement actions of the same week (eight state AG suits, FTC civil investigative demands) and its second-order section reads the same remedy question; this issue carries the proceeding's application-layer pricing detail.

SOURCE: CMA PRESS RELEASE; CMA ROADMAP OF POSSIBLE INTERVENTIONS (PDF); CMA PROPOSED DECISION (PDF)

PATTERN 04

Distribution through the system that already has the licence: SIs as ISVs, tenders as lock-ins

HCLTech: Anthropic Business Unit with forward-deployed engineers, plus Claude and Claude Code embedded in HCL BigFix, AppScan and AI Force

HMRC: core capabilities including AI must be first-party Salesforce software, 15 years with extensions to 2041

Docusign: agreement platform as a tool inside Agentforce and Slack rather than a standalone agent

Workday: Sana inside Claude, Copilot, Teams, Slack and Gemini Enterprise, with Workday's permissions and meter on every action

The week's distribution deals share a shape: the frontier model or the agent enters the enterprise through something that is already licensed, approved and audited, and the commercial terms attach there. HCLTech's deal makes a systems integrator into an ISV, so Claude Code arrives inside BigFix and AppScan under existing software licences rather than as a new vendor on an allowlist. The HMRC tender's first-party-AI condition makes the system of record the only permitted agent runtime for the contract's life. Docusign and Workday both chose to be tools inside someone else's agent surface while keeping their own permission boundary and (for Workday) their own meter. For an architect this is the practical answer to the question of where the control plane for agents will live: not in a new agent-management product, but in whichever existing system holds the permissions the agent needs, because that is where every vendor above has put the toll gate. For a procurement lead it means the agent bill will arrive on invoices from vendors you already have, under contract terms written before agents existed; the renewal, not the RFP, is where the three-layer structure gets negotiated.

SOURCE: HCLTECH PRESS RELEASE; TENDERLINE (HMRC AWARD RECORD); DOCUSIGN NEWS CENTER; WORKDAY NEWSROOM

WATCHLIST

On the radar next.

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

NOV 13

WATCH 01

CMA consultation on Microsoft SMS closes; count of customer submissions naming Copilot Credits or agent-account licences

The only public process in which the Work IQ meter and the agent-account licence, already inside the proposed perimeter, can be put inside a priced remedy before the February 13, 2027 decision; if no customer response names them, the first binding remedy on an enterprise assistant will be drafted for the seat alone.

NOV 21

WATCH 02

End of OpenAI's GPT-5.6 Sol promotional pricing on the Enterprise rate card and first customer migrations from the legacy credit card

The first observable repricing on the new card; the migration path for credit-based agreements decides whether per-token billing lands as a cost increase or a neutral conversion for existing enterprise accounts.

DEC 1

WATCH 03

Microsoft 365 Copilot standalone SKUs and default-on usage billing go live; share of tenants that lowered the 4,000-credit default before the date

The first default-on consumption meter under a mainstream productivity seat; whether admins adjusted the \$40 per-user ceiling before it took effect is the first measure of whether the meter layer is governed or simply paid. The AI Stack Weekly's watchlist groups the December 1 dates (this meter and Gemini 4 Argon's rate reversion) in one entry.

LATE OCT TO EARLY NOV

WATCH 04

Gemini agent general availability and metering details; Google Workspace skills rollout date

Google is the third suite vendor converging on seat-plus-meter-plus-agent, and its coworker-agent pricing (no per-seat fee promised in W41) is the missing column in the three-layer comparison; the delayed Workspace skills rollout decides when enterprise admins get controls.

EARLY 2027

WATCH 05

Workday recruiting autopilot general availability and the first published dollar price per Flex Credit; API overage waiver ends Jan 31, 2027

The rate card has credit counts without a price; the first customer disclosure of a per-credit rate, or Workday publishing one, converts the HR row's challenger pricing from a structure into a number, and the overage meter starts billing the next day.

1H 2027

WATCH 06

Varonis (Proofpoint) and CCC (GTCR, Elliott) closings and any pricing-model changes announced under private ownership

Two sponsor-owned seat-based franchises entering the seat-to-meter transition without public reporting; a meter or agent licence introduced within two quarters of close would confirm that the installed base, not the pricing model, was the asset being bought.

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 26

Week 42 of 2026 · October 17, 2026

WEB

brianletort.ai/industry/applications

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