

THE BIG READ

Three systems of record took three positions on agent traffic, and one of them pays the lab

KEY SIGNALS THIS PERIOD

6

Verticals tracked this period

7

Vertical packages shipped

5

Incumbent SaaS responses

5

Startup signals

AUTHOR

Brian Letort
BrianLetort.AI

PUBLISHED

October 10, 2026
Issue 25 · Weekly read

WEB

[brianletort.ai/industry/applications](#)
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.

Three systems of record took three positions on agent traffic this week: SAP charges its own users per agent action, Workday will toll third-party agents at the API door, and Atlassian tolls nothing and instead pays the lab for the intelligence behind its own agents. The first two legs are the pattern Ben Bajarin (Creative Strategies, The Diligence Stack) named 'The Agentic Data Toll' on October 6, four days before this issue and a day before The Information's Workday report; his read-toll-versus-write-toll map across Salesforce, ServiceNow, Microsoft, SAP and Workday is the fuller version and is credited here. The third leg is this issue's addition, and together they make per-action and per-call rates the newest margin levers a system of record controls, alongside the seat and the storage tier it already had. SAP retired its per-user Joule Premium packages at SAP Connect and now prices every SAP-delivered autonomous agent at 0.02 AI Units per action, idle agents free. Workday, per The Information, will charge when outside agents reach its data through the API; the overage meter that starts on February 1, 2027 is already in its Flex Credit policy at 60 credits per 10,000 calls above a 2.5-6.5 million annual entitlement, with no price per credit published, and on Workday's own rate card the toll is small enough (one first-party agent run is priced like 125,000 third-party calls) to read as a bargaining and visibility meter rather than a revenue line, which is this issue's inference from the policy, not Workday's statement. Atlassian chose not to toll: its rebuilt MCP server (Model Context Protocol, the open standard agents use to call tools) handles 15 million calls a day from Claude, ChatGPT and Cursor, per Atlassian, and Atlassian is paying OpenAI under what an OpenAI representative told VentureBeat is 'effectively a spend commitment', a reversal of last week's pattern in which a lab's spend commitment was currency to buy software. If you run procurement for a company that uses SAP, Workday or Atlassian, the number to find this week is actions per task, because that is the unit your next renewal will be priced in and none of the three vendors has published it; inventory agent-driven API traffic before February and put actions-per-task and per-credit rates into renewal terms now, while the vendors are still deciding them.

The lab side moved the other way on price and the same way on data. Google's Gemini agent, in private preview, is one persistent coworker with its own @agents

email, calendar and Drive, with no per-seat fee for coworker agents and metering details promised closer to a late-October or early-November GA (general availability); Anthropic's Claude Dashboards puts live, refreshable charts with exposed SQL over six warehouses and CRMs inside existing plan allowances, and hands Claude Motion output to Adobe as a finish step. Both are a frontier lab shipping the interface layer over the enterprise's data without a new SKU, which is the squeeze the systems of record are tolling against. The one customer-facing standard of the week came from the application side rather than the labs: Sierra and Meta's Personal Agent Protocol published draft 0.1 on October 9 with a further 35 design partners on top of the launch set, including OpenAI, Visa, Mastercard, Bank of America and Walmart, Anthropic still absent, and Decagon open-sourced its rival PACT under Apache 2.0 and joined the working group. Healthcare's signal was procurement form: Abridge's enterprise-wide agreement across the University of Texas System and Suki's 'Ask' on the UpToDate Expert AI API across four EHRs (electronic health record systems), both without disclosed pricing. The restructuring 8-Ks (the SEC filing a public company makes for a material event) put the pricing conversion in the same quarter as the cuts without saying one caused the other: HubSpot's Exhibit 99.1 says the shift to delivering outcomes with AI 'is transforming product, pricing and how we serve our customers', that the ~7% cut is 'not driven by AI-related efficiencies', and that product teams are reorganizing around customer outcomes; FICO's cites AI-driven product development. The inference here, from the filing read against Breeze's outcomes pricing since April, is that the conversion from seats to outcomes is dilutive first; the exhibit names pricing as part of the transformation, which supports the pairing, and denies the AI-efficiency cause, which limits it.

Three decisions are already in this issue for a buyer. Pull Workday API call counts by integration, agent-initiated separated from human-initiated, before Rising on October 12-15, and use the entitlement baseline as leverage over first-party agent pricing rather than as a cost to avoid. Review which warehouse credentials Claude Dashboards runs under before Docs, Slides and Design turn on by default for Enterprise on October 15. Ask SAP for the action count on your top ten workflows during the free custom-runtime window, because the AI Unit model cannot be forecast without it.

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.

2026-10-06

PACKAGE 01

SAP

OPERATIONS

INCUMBENT SAAS

USAGE

Joule Work, Joule Assistants, AI Agent Hub

Joule Work rolls out to every SAP customer this month; per-user Joule Premium packages are retired for 0.02 AI Units per agent action

This is the most explicit seat-to-action conversion by a major incumbent so far, and the budget consequence is immediate: the cost of an SAP agent is now actions per task times the AI Unit price, and SAP has published neither the actions-per-task figure for any workflow nor a public AI Unit price. Base AI stays included with eligible cloud subscriptions, Premium AI is bought as annually expiring AI Units, idle agents are free, custom-agent runtime is free through end-2026, and Joule for Consultants (35 AI Units per user per month) is the only per-seat SKU left. SAP's own evidence is internal: 110,000 employees on Joule Work with a claimed 20% productivity gain in finance, HR and procurement, a vendor figure. Finance and procurement leads should ask for the action count on their top ten workflows before the Q4 supply-chain agents reach GA, and treat the free custom-runtime window as the period to measure it.

SOURCE: SAP NEWS CENTER; SAP BUSINESS AI PRICING PAGE

2026-10-08

PACKAGE 02

Google Cloud

OPERATIONS

FRONTIER LAB

UNKNOWN

Gemini agent (private preview)

One universal work agent with its own Workspace identity, no per-seat fee for coworker agents, metering to be announced

CIOs get a single SKU in place of Gemini Enterprise, Workspace AI and the agent-builder tier, and a cost model that moves from seats to activity before the activity meter has been published. The agent takes objectives, runs for hours or days in the cloud, routes across Gemini and Claude models, and can be provisioned as a persistent coworker with its own @agents address, calendar and Drive. Google told Futurum it does not intend to charge a per-seat licence for coworker agents, with metering details closer to a late-October or early-November GA (the pricing table carries the editions). Until the meter is published, treat the \$0 seat as a deferred price, not a free one, and note that this agent reaches ServiceNow and Salesforce data through a tools registry and MCP, which is the traffic the systems of record are now tolling.

SOURCE: GOOGLE (THE KEYWORD); FUTURUM GROUP

2026-10-08

PACKAGE 03

Anthropic

RESEARCH

FRONTIER LAB

HYBRID

Claude Dashboards and Claude Motion (beta); Docs, Slides and Design GA

Live BI over six warehouses and CRMs with the SQL exposed, inside existing plan allowances, plus a 'Send to Adobe' handoff

Analytics leads now have a self-serve BI layer from a frontier lab that connects to BigQuery, Databricks, Snowflake, Redshift, ClickHouse or Salesforce, answers plain-language questions with refreshable charts, shows the underlying SQL, and pushes to Grafana, Hex, Mixpanel or Sigma, with Looker and Tableau listed as coming. It is bundled into Pro, Max, Team and Enterprise allowances (Enterprise is seat plus usage at API rates), not a new SKU, which is the pressure point for seat-priced BI. For Enterprise admins Dashboards and Motion are off by default; Docs, Slides and Design turn on by default on October 15, so admins have a week to decide. Motion generates code-based, editable animations with no video model and no synthetic people and hands off to Adobe Firefly Video Editor, Descript, HeyGen, Runway and Luma; Adobe confirmed the flow and said its Claude plugin exposes more than 80 tools. The data-governance question to answer first is which warehouse credentials the connector runs under, because the exposed-SQL design means the agent is querying production data on the user's permissions.

SOURCE: ANTHROPIC (CLAUDE BLOG); IT BRIEF (ADOBE HANDOFF)

2026-10-06

PACKAGE 04

Sierra (with Meta)

SUPPORT

STARTUP

OUTCOME

Personal Agent Protocol (draft 0.1), Horizon revenue agents, Curie and Fleming models

The CX (customer-experience) vendors standardize the agent-to-agent front door before the labs do: a further 35 design partners on draft 0.1, OpenAI in, Anthropic absent

Heads of digital in banking, retail, insurance and travel are being pulled into a design partnership now, and the choice of which protocol a brand's front door speaks will decide whether the brand or the personal-agent platform owns the customer session. Draft 0.1, nicknamed Poppy, is built on OAuth, JWTs, OpenAPI and MCP; the October 9 list adds 35 design partners (Visa, Mastercard, PayPal, Bank of America, Wells Fargo, Target, United Airlines, Zendesk, Notion and Okta among them) to the launch set of Walmart, Stripe, Shopify, Genesys, Rocket and Instinct (Meta's list also names NiCE and Decagon). Sierra's Fleming model detects when a caller is another agent, which is the capability every contact centre will need once consumer agents are common; Horizon revenue agents sell on Sierra's outcomes model. The spec's owner is a vendor, so enterprises should ask for the working-group charter before standardizing.

SOURCE: SIERRA; CNBC

2026-10-08

PACKAGE 05

Abridge

HEALTHCARE

STARTUP

SEAT

Abridge clinician intelligence platform (UT System enterprise agreement)

A state-system master agreement becomes the procurement vehicle for ambient AI: every UT health institution, local opt-in, no price disclosed

Health-system CIOs should read this as the shape of the next buying cycle: a legislatively funded system initiative (UT REAL Health AI) signs one framework, each institution adopts under it, and the point-solution window closes. The agreement centres on ambient documentation and the clinical intelligence Assistant with a stated path into care-gap surfacing, payer administrative workflows and clinical-trial matching. Abridge says the platform will support more than 100 million patient-clinician conversations across 300+ US health systems in 2026 and is validated in 28+ languages, both vendor figures. Pricing is per-clinician enterprise subscription and no contract value was disclosed, so the economic signal is the procurement form, not the rate.

SOURCE: ABRIDGE

2026-10-08

PACKAGE 06

Suki

HEALTHCARE

STARTUP

SEAT

Suki Ask on the UpToDate Expert AI API

The ambient scribe adds clinical reasoning: chart queries plus UpToDate-grounded answers inside Epic, Oracle Health, athenahealth and MEDITECH

Clinical informatics leads get the first external product on Wolters Kluwer's UpToDate Expert AI API, generally available this week to Suki's 400+ health-system and health-tech customers, some enabled automatically. Clinicians query medications, allergies, vitals, labs and chronic conditions by voice or text and ask clinical questions answered only from UpToDate content with sourcing, dosing considerations and monograph links; UpToDate is used by more than three million clinicians. No incremental price was disclosed and Ask requires an existing UpToDate licence, so the bundling is the point: documentation, reference content, ordering (Suki reports 72% month-over-month growth in ambient order staging, a vendor figure) and coding converge in one subscription, which is the squeeze on standalone clinical-reasoning tools and a new distribution channel for the content incumbent.

SOURCE: SUKI; NEWSWEEK

2026-10-06

PACKAGE 07

Atlassian

ENGINEERING

INCUMBENT SAAS

HYBRID

Rovo and platform agents on OpenAI models; Atlassian MCP (rebuilt)

A SaaS incumbent commits spend to a lab: GPT-6 Astra and GPT-5.6 behind Rovo, 220+ MCP tools at 15 million calls a day

Engineering leaders get Jira and Confluence context inside ChatGPT and Codex through the Atlassian MCP plugin, and Atlassian gets frontier models behind Rovo and its platform agents under an agreement an OpenAI representative described to VentureBeat on background as 'effectively a spend commitment' (VentureBeat also carries the MCP permission and risk-tier details). Atlassian says Astra is not Rovo's default and that its gateway routes across OpenAI and other providers on capability and cost, so the commitment is a floor on OpenAI spend, not an exclusivity. Two numbers matter for the architecture question: 15 million MCP calls a day is the first disclosed measure of MCP as a distribution channel from a major system of record, and 3,000+ Atlassian developers on Codex is a reference deployment. No Rovo price change was announced; the Teamwork Graph context layer is the asset Atlassian is keeping while renting the intelligence.

SOURCE: OPENAI; ATLASSIAN; VENTUREBEAT

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.

2026-10-07

RESPONSE 01

Workday

AGENT DATA-ACCESS
TOLL AND FLEX CREDIT
API OVERAGE POLICY

The Information reports Workday will charge when third-party AI agents access customer data; API-overage billing stays waived through January 31, 2027, then bills from February 1 at 60 Flex Credits per 10,000 calls above a 2.5-6.5 million annual entitlement; no price per credit published; Workday Rising is October 12-15

On Workday's own Flex Credit card the toll is small: 60 credits per 10,000 calls against 1 credit for a self-service question and 750 for a talent-rediscovery run, so one first-party agent run is priced like 125,000 third-party calls, and the 15,000-200,000 complimentary credits a year cover roughly 2.5-33 million overage calls. The toll is a visibility and bargaining meter rather than a revenue line, and the waiver period is the window in which customers can measure their agent traffic before it costs anything. CIOs should pull API call counts by integration now, separate human-initiated from agent-initiated traffic, and treat the entitlement baseline as leverage over first-party agent pricing at Rising rather than as a cost to avoid. For investors, Bhusri's earlier 'considerable financial upside' remark now has a meter behind it, if not yet a price.

SOURCE: THE INFORMATION; WORKDAY FLEX CREDIT AND PLATFORM ENTITLEMENT POLICY (VIA REDRESS COMPLIANCE)

2026-10-06

RESPONSE 02

SAP

JOULE PREMIUM
PACKAGES RETIRED; AI
UNITS PER AGENT
ACTION

Per-user-per-month Joule Premium packages for Finance, HCM, Supply Chain, Spend and CX are being retired; all SAP-delivered autonomous agents in 2026 cost 0.02 AI Units per action, nothing when idle; AI Units are bought annually and expire after 12 months

SAP has moved the whole agent catalogue to one meter, which is the cleanest comparable for the ERP tier and the hardest to forecast, because the expiring annual purchase forces customers to pre-buy actions they have not yet measured. Procurement should buy the smallest AI Unit block that covers measured Q4 usage and negotiate rollover, and finance should model the expiry as breakage in the vendor's favour.

SOURCE: SAP BUSINESS AI PRICING PAGE

2026-10-06

RESPONSE 03

HubSpot and Fair Isaac

WORKFORCE REDUCTIONS
(FORMS 8-K)**HubSpot cuts about 7% (~660 roles) with \$65-75M of charges, reorganizing product teams around customer outcomes and stating the cuts are 'not driven by AI-related efficiencies'; FICO eliminates about 15% of positions 'integrating AI-driven product development' with ~\$27M of charges**

HubSpot's Exhibit 99.1 says the company has 'shifted our strategy from building software that helps customers grow to delivering outcomes for them with AI', that this shift 'is transforming product, pricing and how we serve our customers', and that the cuts are 'not driven by AI-related efficiencies'. So the filing does name pricing as part of the transformation; what it does not do is attribute the cuts to the pricing change. The inference here, reading the exhibit against Breeze's outcomes pricing since April, is that this is the first restructuring filed in the same quarter as an outcome-pricing conversion, which, if the reading holds, tells investors the conversion costs margin before it earns it and tells customers that renewal terms will be restructured around outcomes rather than Hubs. FICO's filing is the scoring incumbent reorganizing product development around AI. With Workday's W40 filing (a large-cap, and one that did not attribute its cut to AI), that is three restructuring 8-Ks in two weeks, one of which ties the cuts to an outcomes pivot and one to AI product development.

SOURCE: SEC EDGAR (HUBSPOT 8-K); FAIR ISAAC 8-K

2026-10-06

RESPONSE 04

ServiceNow

FLOW BY SERVICENOW

Flow reaches existing North American customers on the Now Assist assists pool with no new licence; GA for all NA and EMEA customers including credit-card purchase is targeted later in Q4; a competitor-published rate card of \$2 per conversation and \$1 per automation run with a \$10K annual minimum is unconfirmed by ServiceNow

The availability date held, which is the only confirmed fact; the rate card is a competitor's number and should be treated as a hypothesis until ServiceNow publishes. If it holds, \$2 per conversation is the first public per-conversation benchmark from a tier-one ITSM vendor and becomes the comparison point for every usage-priced service-desk negotiation; buyers should hold the question open until the Q4 GA.

SOURCE: SERVICENOW NEWSROOM; REZOLVE.AI (COMPETITOR ANALYSIS, UNVERIFIED PRICING)

2026-10-08

RESPONSE 05

Adobe

SEND TO ADOBE HANDOFF
FROM CLAUDE MOTION;
CLAUDE PLUGIN WITH
80+ TOOLS**Adobe confirms a 'Send to Adobe' flow from Claude Motion into Firefly Video Editor and says its Claude plugin now exposes more than 80 tools**

Adobe is accepting the position of finish step rather than starting point for a class of creative work, which protects the editing seat while ceding the first draft to the lab. Creative-ops leads should expect the same pattern from other incumbents whose value is in the last 20% of a workflow: integrate as the destination, keep the professional tool, let the agent generate.

SOURCE: IT BRIEF

STARTUP SIGNALS

The insurgent vertical cohort.

5 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.

2026-10-08

SIGNAL 01

Manus (Butterfly Effect)

OTHER

\$500M+ AT ~\$4B
(REPORTED)

Closes a \$500M+ round co-led by Boyu Capital and IDG Capital with Tencent, HSG and ZhenFund participating, the first raise since Chinese authorities ordered Meta's \$2B+ acquisition unwound in April; valuation undisclosed, ~\$4B per Bloomberg's September report; The Information put annualized revenue run-rate near \$500M in June

A general-purpose agent company doubling its valuation six months after a blocked acquisition, with a Hong Kong listing under consideration, is the clearest signal that the horizontal agent category has a public-market path outside the US. The revenue figure is a single-outlet report; the round and its leads are company-confirmed.

SOURCE: TECHCRUNCH; REUTERS

2026-10-09

SIGNAL 02

TypeSafe (maker of Jev)

OTHER

\$870M AT \$7.5B

Raises \$870M at a \$7.5B valuation led by a16z (TechCrunch), weeks after launching Jev, the hosted decision model priced at \$0.042 per million input tokens (price per Simon Willison's write-up), the same figure Microsoft-Decision-1 listed on October 9

The decision-model category now has a venture-scale pure play priced at \$7.5B in the same week that Microsoft, Cloudflare and OpenAI entered or cut price, which tells application vendors that the approve, route and classify call is being treated as its own market rather than a feature. The valuation and lead are as reported by TechCrunch; the round documents are not public. The Model Pulse carries the category's price table.

SOURCE: TECHCRUNCH; SIMON WILLISON (JEV PRICE)

2026-10-09

SIGNAL 03

Decagon

SUPPORT

Open-sources PACT (Personal Agent Consent and Trust) under Apache 2.0 with Instinct and joins the Sierra-Meta Personal Agent Protocol working group one week after shipping PACT as its own standard

A one-week reversal from rival protocol to open contribution is the fastest consolidation of an agent standard so far and reduces the risk that brands must implement two front doors. For buyers it means the Sierra-led spec is the one to track; for Decagon it trades spec ownership for a seat at the table and keeps its Agent Modules as the differentiator.

SOURCE: DECAGON

2026-10-08

SIGNAL 04

Suki

HEALTHCARE

Reports 72% month-over-month growth in clinicians adopting ambient order staging across 300+ organizations alongside the Ask launch, and names coding, RCM and operational task completion as the next steps

Order staging is the step where an ambient scribe becomes an agent that changes the record, and a growth rate like that (vendor-reported, no base disclosed) is the evidence that clinicians accept it. The next step, coding and revenue-cycle actions, is where the pricing will move from per-clinician to per-outcome, and where health-system CFOs should expect the first outcome-based offers.

SOURCE: SUKI BLOG

2026-10-06

SIGNAL 05

Sierra

SUPPORT

Announces two in-house models, Curie for orchestration and reply generation and Fleming for detecting when a caller is another agent; repositions Ghostwriter as an always-on Slack and Teams teammate that monitors agents and runs A/B tests; expands Horizon revenue agents on outcomes pricing

A CX startup training its own orchestration and bot-detection models, rather than renting everything from a lab, is the application layer building the pieces the labs have not: agent-caller detection is a security control every contact centre will need once consumer agents call in, and no frontier lab sells it. Model quality claims are vendor-only; the models are not publicly available and are not on the LLM tree.

SOURCE: SIERRA

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.

2026-10-06

SHIFT 01

SAP

Seat-based → Usage-based

Per-user-per-month Joule Premium packages (Finance, HCM, Supply Chain, Spend, CX) retired; every SAP-delivered autonomous agent in 2026 priced at 0.02 AI Units per agent action with no charge when idle; AI Units purchased annually and expiring after 12 months; Joule for Consultants (35 AI Units per user per month) remains the only per-seat Premium SKU; custom-agent runtime free through end-2026. No public AI Unit price and no actions-per-task figures.

SOURCE: SAP BUSINESS AI PRICING PAGE

2026-10-07

SHIFT 02

Workday

Seat-based → Hybrid

Reported plan to charge when third-party AI agents access customer data through the platform; production API calls above a 2.5-6.5 million annual entitlement (by company size) draw Flex Credits at 60 credits per 10,000 calls; overage billing suspended from May 30, 2026 through January 31, 2027 and billed from February 1, 2027; price per credit not published. Formal agentic HR and finance pricing expected at Workday Rising, October 12-15.

SOURCE: THE INFORMATION; WORKDAY FLEX CREDIT POLICY (VIA REDRESS COMPLIANCE RESEARCH NOTE)

2026-10-08

SHIFT 03

Google Cloud

Seat-based → Unknown

Gemini agent coworker agents carry no per-seat licence in Workspace and are included for Gemini Enterprise customers (Business edition lists \$21 per seat per month); a \$0-seat-fee pay-as-you-go edition is rolling out; activity-metering details are promised closer to a late-October or early-November GA. Until the meter is published the effective price is unknown.

SOURCE: FUTURUM GROUP; GOOGLE CLOUD PRICING PAGE

2026-10-07

SHIFT 04

Anthropic

Hybrid → Hybrid

Monthly Claude Platform API credits added to subscriptions: \$100 for Max 5x, \$200 for Max 20x, and per-seat Team credits (\$20 Standard, \$100 Premium) pooled up to \$500 per month, usable on any model, expiring each cycle; Free, Pro and Enterprise excluded. Dashboards and Motion are bundled into existing plan allowances rather than sold as a SKU. The subscription now funds metered API use, which blurs the seat-versus-token boundary from the lab side.

SOURCE: ANTHROPIC

VERTICAL SCORECARD

Who leads each vertical.

Leader-vs-challenger by vertical, useful for procurement shortlists when matching workload to vendor cohort. As of Oct 10, 2026.

VERTICAL	LEADER	CHALLENGER	READ
Commerce	Amazon Seller Assistant via the Selling Partner plugin for Claude	Personal Agent Protocol retail design partners (Walmart, Shopify, Stripe, Target) via Sierra	Leader unchanged with no in-window news. The challenger moves from Decagon Agent Modules to the retail bloc of the Personal Agent Protocol, because the week's commerce movement was the standard for how consumer agents transact with merchants, and Walmart, Shopify and Stripe are its launch partners.
Engineering	Codex on GPT-6.1 Sol (default) with Codex Cloud for Enterprise	Claude Code on Sonnet 5.5 with managed Mods	Unchanged. Both shipped governance releases this week (fail-closed hooks in Claude Code 2.1.295, turn-bound permissions and enterprise MCP auth in Codex 0.162), which Agent Techniques carries; Atlassian's 3,000 developers on Codex and the 15 million daily MCP calls are the week's distribution datapoints for the leader.
Healthcare	Abridge	Suki (Ask on UpToDate Expert AI)	Both slots change. Abridge takes the leader slot from OpenEvidence on the UT System enterprise-wide agreement and the 300+ health-system footprint; Suki takes the challenger slot from EliseAI on the first external UpToDate API product across four EHRs. OpenEvidence had no in-window product evidence and is the standalone tool the Suki bundle targets; EliseAI's last signal was a funding round.
Support	Sierra (Personal Agent Protocol, Horizon, Curie and Fleming)	Decagon (Agent Modules; PACT open-sourced)	Leader and challenger swap. Sierra published the protocol 35 design partners signed, shipped two in-house models and expanded outcome-priced revenue agents; Decagon open-sourced its own protocol and joined Sierra's working group. Flow by ServiceNow drops to the incumbent-response table pending a confirmed rate card.
Sales	ZoomInfo Agent Teams	HubSpot Breeze (outcomes pricing, reorganized around outcomes)	Leader unchanged with no in-window news. The challenger moves from Salesforce to HubSpot on pricing posture, not product evidence: HubSpot's 8-K says the outcomes shift is transforming pricing and reorganizes product and go-to-market around outcome delivery, with Breeze already on outcomes pricing; Salesforce's only in-window appearance is as a Claude Dashboards data source.

VERTICAL	LEADER	CHALLENGER	READ
Security	Codex Security Cloud	Elastic AlertZero (Technical Preview)	Challenger changes from CrowdStrike and Palo Alto Networks to Elastic, whose AlertZero agent team shipped with per-agent autonomy levels and mandatory approval for consequential actions, the first SOC product to expose that control; the PromptArmor disclosure against Elastic's agentic SOC remains unaddressed, which Agent Techniques covers.
Legal	Harvey	Legora	Unchanged; no in-window legal movement. Carried pending evidence.
Research	Claude (Dashboards over Snowflake, Databricks, BigQuery, Redshift, ClickHouse, Salesforce)	Gemini agent (hours-long objective runs, private preview)	Both slots change from model names to products, because the week's research movement was interface, not model: Claude Dashboards ships live BI with exposed SQL inside existing plans, and the Gemini agent runs multi-day research objectives but is private-preview only. Sonnet 5.5 and GPT-6.1 Sol remain the underlying models.
Finance	SAP Joule (finance agents at 0.02 AI Units per action)	Workday agentic finance (Rising, October 12-15)	New row. SAP's Joule Work rollout and per-action pricing are the first fully metered finance-agent catalogue from an ERP incumbent; Workday's toll and Rising agenda make it the challenger pending its pricing announcement next week.

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

Toll, meter or pay the lab: three positions on agent traffic at the data door

Workday agent data-access charge; API overage at 60 Flex Credits per 10,000 calls from February 1, 2027

SAP 0.02 AI Units per agent action, idle free

Atlassian MCP at 15 million calls a day, untolled; Atlassian pays OpenAI under a spend commitment

Google Gemini agent reaching ServiceNow and Salesforce via tools registry and MCP (the traffic being tolled)

The data-toll frame is Ben Bajarin's ('The Agentic Data Toll', Creative Strategies / The Diligence Stack, October 6), including the read-toll-versus-write-toll map across Salesforce, ServiceNow, Microsoft, SAP and Workday; this entry adds the third position and the Workday arithmetic. Two systems of record now meter agent activity at the point where agents touch their data, and a third chose not to: SAP charges per action by its own agents, Workday per call by anyone else's, and Atlassian, whose rebuilt MCP server handles 15 million calls a day (per Atlassian; the first disclosed measurement of that traffic from a major incumbent), tolls nothing and instead pays the lab under a spend commitment, monetizing the traffic indirectly through Rovo. ServiceNow's Flow, priced by usage for its own service desk, is a usage-priced product rather than a toll on external agents and belongs in the incumbent table. The architectural fact underneath is that MCP and tool registries have made the API the agent's user interface, so the API baseline, which was a fair-use allowance when humans drove integrations, becomes the pricing surface when agents do. For architects the design implication is to put a metering proxy in front of every system-of-record connector now, attribute calls to agents and tasks, and cache aggressively, because the per-call meter rewards fewer, larger reads. For procurement the implication is that entitlement baselines and per-credit prices are the negotiable terms, and the February waiver window is the time to measure.

SOURCE: WORKDAY FLEX CREDIT POLICY; SAP BUSINESS AI PRICING; ATLASSIAN

PATTERN 02

The agent-to-agent front door standardizes from the application side

Personal Agent Protocol draft 0.1 'Poppy' (OAuth, JWTs, OpenAPI, MCP; 35 design partners incl. OpenAI, Visa, Mastercard, PayPal, Bank of America, Walmart)

Decagon PACT open-sourced under Apache 2.0, joins the working group

Sierra Fleming (detects agent callers)

Rocket mortgage pre-approval via Meta Muse

The protocol for how a consumer's agent authenticates and transacts with a business is being written by a CX vendor and a social platform, with one lab inside the tent and one outside, and the rival standard folded in within a week. The composition is familiar (OAuth for delegation, JWTs for identity, OpenAPI for capability description, MCP for tool invocation), which is why adoption can be fast and why the governance question matters more than the technical one: the working group's charter, the licensing of the spec, and whether the brand or the personal-agent platform holds the session token are the terms that decide who owns the customer. Fleming is the operational counterpart: once consumer agents call in, every contact centre needs to know it is talking to one, and that detection becomes a security control. Enterprises in banking, retail, travel and insurance should join as design partners now if they want a say, and should build their front door to speak the draft rather than wait for 1.0.

SOURCE: SIERRA; DECAGON; CNBC

PATTERN 03

The lab ships the interface layer over the warehouse and the suite

Claude Dashboards (BigQuery, Databricks, Snowflake, Redshift, ClickHouse, Salesforce; exposed SQL; push to Grafana, Hex, Mixpanel, Sigma)

Gemini agent (@agents identity; Gmail, Microsoft 365, Slack, ServiceNow, headless API)

Claude Motion to Adobe Firefly Video Editor; Adobe plugin with 80+ tools

Two frontier labs shipped products this week whose value is the connector, not the model: a BI layer that queries six data platforms on the user's credentials, and a coworker agent with a directory entry that reaches the incumbents' data through tools and MCP. Neither carries a new SKU, which is the commercial design: the lab bundles the interface into a subscription it already sells and lets the metered usage grow. The incumbents' responses are the toll (above) and the finish step (Adobe), and the question for a CIO is which layer holds the user's permissions. Dashboards runs SQL on production data under the user's warehouse credentials with the SQL exposed, which is good for auditability and bad if the credentials are over-scoped; Enterprise admins have Dashboards and Motion off by default and Docs, Slides and Design on by default from October 15. The data-access review should happen before the 15th, not after.

SOURCE: ANTHROPIC (CLAUDE BLOG); GOOGLE (THE KEYWORD)

PATTERN 04

Master agreements and content APIs as the healthcare distribution layer

Abridge enterprise-wide agreement with UT REAL Health AI (all UT health institutions)

Suki Ask as the first external product on Wolters Kluwer's UpToDate Expert AI API (Epic, Oracle Health, athenahealth, MEDITECH)

Abridge 300+ health systems, 100M+ conversations in 2026 (vendor figures)

Healthcare's application layer is consolidating through two mechanisms that have nothing to do with model quality: system-level master agreements that let each institution opt in under one framework, and reference-content APIs that make a scribe into a reasoning tool without the vendor building a knowledge base. Both favour the vendors already in the chart (Abridge, Suki, Ambience, Microsoft) and compress the window for standalone tools, including the clinical-reasoning category OpenEvidence leads. The pricing consequence is that per-clinician subscriptions will bundle more functions before they move to outcomes; the point at which coding and revenue-cycle actions enter the bundle is where outcome-based pricing will appear, and Suki named that as the next step.

SOURCE: ABRIDGE; SUKI; NEWSWEEK

PATTERN 05

Restructuring and pricing conversion in the same quarter

HubSpot 8-K, Exhibit 99.1: ~7%, \$65-75M, reorganized around customer outcomes; the outcomes shift 'is transforming product, pricing and how we serve our customers'; cuts 'not driven by AI-related efficiencies' (Breeze on outcomes pricing since April)

Fair Isaac 8-K: ~15%, ~\$27M, AI-driven product development

Workday 8-K (W40; no AI attribution)

SAP per-action conversion (no 8-K; the pricing is the restructuring)

Three restructuring 8-Ks in two weeks, one of which (HubSpot) ties the cuts to an outcomes pivot while denying that AI efficiencies drove them, one (FICO) to AI-driven product development, and one (Workday, W40) to neither. HubSpot's Exhibit 99.1 does name pricing: the shift to delivering outcomes 'is transforming product, pricing and how we serve our customers'. What the exhibit does not do is say the cuts are caused by the pricing change, so the link between the restructuring and the outcome-pricing conversion is this issue's inference from their co-occurrence, not the filing's claim. The pattern investors should price, on that inference, is that the conversion from seats to actions or outcomes is dilutive first: revenue recognition moves from contracted seats to measured activity, go-to-market is rebuilt around delivery of results, and the vendor absorbs the measurement cost before the customer pays for the outcome. SAP is running the same conversion without a restructuring filing because it is converting pricing rather than people, which is the alternative path. Customers should expect renewal terms that reference outcomes and should ask for the measurement method in writing; engineers who own the integrations should expect the outcome definitions to arrive as metering requirements on their side of the API.

SOURCE: SEC EDGAR (HUBSPOT 8-K); FAIR ISAAC 8-K

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.

OCT 12–15

WATCH 01

Workday Rising: formal agentic HR and finance pricing, the agent data-access toll, and a per-credit price

The first public rate for tolling third-party agents at a system of record; a per-credit price and the entitlement baselines decide what the February 1 meter costs.

OCT 15

WATCH 02

Claude Docs, Slides and Design turn on by default for Enterprise admins

The data-access review for Dashboards and Motion (off by default) should be done by the same date; the exposed-SQL design runs on the user's warehouse credentials.

LATE OCT TO EARLY NOV

WATCH 03

Gemini agent GA and its activity meter

The \$0 seat is a deferred price until the meter is published; the metering unit (actions, tokens, runtime hours) decides whether it undercuts or matches SAP's per-action model.

Q4 2026

WATCH 04

SAP supply-chain Joule Agents GA and the first customer-disclosed actions-per-task figures

The AI Unit model is unforecastable without actions per task; the first customer or partner that publishes a count sets the benchmark for every SAP renewal.

Q4 2026

WATCH 05

Flow by ServiceNow GA with a confirmed rate card

Settles whether the reported \$2 per conversation holds; it would be the first tier-one per-conversation benchmark for service-desk pricing.

OPEN

WATCH 06

Personal Agent Protocol working-group charter and Anthropic's position

A vendor-owned spec with one lab in and one out; the charter and licence decide whether enterprises can standardize on it, and Anthropic joining or publishing an alternative changes the map.

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 25

Week 41 of 2026 · October 10, 2026

WEB

brianletort.ai/industry/applications

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