

THE BIG READ

A lab's spend commitment became procurement currency, and the credit became the universal meter

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 3, 2026
Issue 24 · 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.

Four vendors changed how enterprise AI is metered in the same week, and none of their credits converts to dollars or to work at the same rate. Microsoft set a 4,000-credit monthly cap per Copilot Business user, about \$40 at its published one-cent rate, and pushed the default-on date to December 1; HubSpot moved new EMEA accounts to one edition plus seats plus credits at one cent per extra credit; ServiceNow's Flow service desk is free now, consumption-priced at GA, and draws on existing customers' pool of 'assists' at a rate it has not disclosed; Anthropic took seat fees out of Claude for Government in favour of prepaid dollar increments under a hard cap. The seat count no longer predicts the bill, and the exchange-rate table (which credit costs what, and what one credit buys) is the operational artifact a finance team now needs; the credit-as-meter entry in the architecture watch is its home. For the engineer, the same table is the spec for an agent gateway: every action an agent takes now has a vendor-specific price, so the gateway has to meter per agent and per workflow, not per user.

The second fact is a payment rail. OpenAI's Marketplace lets an eligible US enterprise apply part of an existing OpenAI spend commitment to software from 32 partners built on OpenAI models, with the partner invoicing directly and OpenAI reconciling the eligible amount against the commitment; the eligible share is a fixed percentage up to a cap, set per contract and not published. The Register (September 30, 'OpenAI apes AWS Marketplace while Anthropic remains stubbornly Microsoftian') made both the comparison and the consequence first: AWS and Azure marketplaces have let customers burn down cloud commitments on third-party software for years, and a model lab doing it makes the lab, not the cloud, the committed-spend relationship an application vendor needs to be inside. The house's addition is the procurement question that follows: the customer's negotiated discount is now shared across two vendors in a ratio nobody has disclosed, and the ratio is the term to negotiate.

The third fact is who got paid: EliseAI (\$350 million at \$4 billion on more than \$200 million of ARR, annual recurring revenue), OpenEvidence (\$250 million at \$15 billion with AdventHealth system-wide), Instinct (\$1 billion at \$10 billion) and ElevenLabs (a \$300 million employee tender at \$22 billion) are in the startup

signals, and Workday's 2.5% cut two weeks before its own conference is in the incumbent responses. The house reads the valuations as paying for proprietary interaction data and workflow embedding rather than for model quality; nothing in the announcements suggests OpenEvidence or EliseAI runs a model a horizontal suite could not license.

For a buyer, the moves are concrete. Procurement should ask each application vendor whether it is in a lab marketplace and, if so, what the commitment burn-down does to the discount already negotiated. Finance should build the credit ledger before December 1, because Microsoft's cap is adjustable down to zero and most organizations have not decided who sets it. Operators should treat ServiceNow's Flow as a 90-day trial of whether a service desk can be bought like a SaaS app on a credit card; if it can, the implementation partner's share of the stack is the thing at risk. Anyone evaluating a vertical agent should price the data moat and the integration depth, not the model. The model economics behind this, the \$2 / \$10 convergence and the Ultrafast tier, are in The Model Pulse; the permission architecture that dots and PACT introduce is in Agent Techniques.

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.

2026-09-29

PACKAGE 01

OpenAI

OPERATIONS

FRONTIER LAB

HYBRID

OpenAI Marketplace, specialist dots, Private Intelligence

ChatGPT Enterprise becomes a procurement platform: 32 partners can be paid from an existing OpenAI commitment

Eligible US enterprises apply part of an OpenAI spend commitment to partner software; the partner invoices, OpenAI reconciles. The eligible percentage (a fixed share up to a cap, per The Register) and any partner fee are not published, which is what a procurement lead needs to ask first. Specialist dots (role-scoped agents with their own credentials, governed through Microsoft Agent 365, Microsoft's control plane for registering and policing agents as it does users) and Private Intelligence (zero data retention now, Private Inference preview this fall) are the enterprise-only features that make the commitment worth concentrating. Custom GPT creation for Enterprise ends October 26, so anyone who built internal tools on Custom GPTs has three weeks to move them to dots or Spaces.

SOURCE: OPENAI HELP CENTER

2026-09-30

PACKAGE 02

Anthropic

OTHER

FRONTIER LAB

USAGE

Claude for Government

GA for US federal and state agencies with no seat fees: prepaid usage increments under a hard not-to-exceed cap

The procurement design is the news: agencies prepay in fixed increments with a hard cap, department admins allocate to sub-agencies, SCIM group mappings (SCIM is the standard for syncing user groups from an identity provider) set rate limits, dollar caps and allowed models per tier, and burndown alerts fire before balances run out. Delivered in a FedRAMP High environment (the US government's strictest cloud security authorization level) with Claude Code CLI and Claude for Microsoft 365 in early access, and contractable directly without a cloud provider. Any CFO who wants a model for capped AI spend has one to copy, and any engineer wiring an agent platform into an identity provider has a reference for what group-to-budget mapping looks like.

SOURCE: ANTHROPIC

2026-10-01

PACKAGE 03

ServiceNow

SUPPORT

INCUMBENT SAAS

USAGE

Flow by ServiceNow

A standalone, consumption-priced AI service desk that runs in Teams, Slack and email with no implementation project

Flow can run without a ServiceNow instance, is free during controlled availability, and will be consumption-priced at GA with existing AI-entitled customers drawing on their 'assists' pool and new customers buying by credit card. It is ServiceNow attacking the implementation-heavy part of its own category before a startup does. IT leaders should run it against their current service desk for the 90-day window and measure deflection (the share of tickets resolved without a human agent), not demos.

SOURCE: SERVICENOW NEWSROOM

2026-10-01

PACKAGE 04

Decagon

- SUPPORT
- STARTUP
- UNKNOWN

Decagon Voice 3, Agent Modules, Personal Agent Gateway and PACT

Four releases move CX (customer experience) agents into collections and payments and give enterprises a protocol for consumer agents at the front door

Voice 3 runs a duplex architecture, a low-latency conversational model plus a reasoning model for tools and guardrails, speaking through Decagon's own Chord voice model in 70+ languages. Agent Modules extend the same agents into lead qualification, onboarding, payments and collections with outcome analytics. PACT, built on A2A (Google's agent-to-agent messaging protocol) and OAuth 2.0 delegated authorization, is the first vendor protocol for identifying and constraining consumer agents such as dots and Muse when they call a company. CX leaders should ask their vendor how it will authenticate a customer's agent, because that traffic is arriving this quarter; engineers who own a public API or chat front end should plan a separate client class for delegated agents, with its own rate limits and scopes.

SOURCE: DECAGON

2026-10-01

PACKAGE 05

OpenEvidence

- HEALTHCARE
- STARTUP
- UNKNOWN

OpenEvidence

AdventHealth goes system-wide inside EHR (electronic health record) workflows as the company confirms \$250M at \$15B

AdventHealth made the clinical decision-support platform available to physicians, nurses, pharmacists, care managers and therapists across its network, following September deals with Penn Medicine, Memorial Sloan Kettering and UTMB. Becker's reports, citing the company, that more than half of US physicians use it regularly. The moat is the citation-grounded clinical corpus and EHR embedding, not the model; health-system CIOs evaluating alternatives should price the integration, not the LLM, and engineers building clinical tools should expect the EHR vendor's integration surface, not model choice, to set the schedule.

SOURCE: BECKER'S HOSPITAL REVIEW

2026-10-01

PACKAGE 06

ZoomInfo

- SALES
- INCUMBENT SAAS
- SEAT

Agent Teams in GTM Studio

Multi-agent GTM (go-to-market) orchestration ships a day after the DoubleO.ai acquisition, with no new SKU

Agent Teams coordinates agents across a customer's systems and ZoomInfo's GTM Context Graph (100M+ companies, 500M+ contacts) to run plays from signal to action, available immediately in GTM Studio without a new contract. The no-SKU decision is a defensive move: ZoomInfo is bundling orchestration into the data subscription to keep the data the thing customers pay for. Sales-ops leaders get a free multi-agent layer and should test whether it reduces the number of point tools they pay for; sales-engineering teams should check which of their systems the orchestration can write to, not only read from, before letting it run plays.

SOURCE: ZOOMINFO VIA VENTUREBEAT

2026-09-29

PACKAGE 07

EliseAI

OPERATIONS

STARTUP

UNKNOWN

EliseAI housing and healthcare operations agents

\$350M at \$4B on more than \$200M ARR for agents that run leasing, maintenance, billing and clinic intake

EliseAI says it crossed \$200M ARR in June with five consecutive years of doubling and says on its blog that it powers one in five US multifamily apartments (the company's figure, not independently verified); the healthcare unit does scheduling, intake, insurance verification and chart prep for specialty groups. This is what a vertical operations agent at scale looks like: narrow workflows, a proprietary interaction history, and a customer P&L the vendor can point at. Operators in adjacent property and clinic workflows should expect the same playbook to arrive in their category; engineers in those categories should note that the defensible asset is the interaction history, which means logging and retaining the agent's conversations is a product decision, not only a compliance one.

SOURCE: ELISEAI VIA GLOBENEWSWIRE

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.

2026-10-01

RESPONSE 01

ServiceNow

FLOW BY SERVICENOW

Launches a 'startup within ServiceNow': a standalone AI service desk, free in controlled availability, consumption-priced at GA, purchasable by credit card

ServiceNow is disrupting its own implementation economics before a startup does, and the credit-card path is aimed at buyers who never had a Now instance. Existing customers draw on their 'assists' pool with no new license, so the first-year revenue effect is small and the strategic effect is to make the service desk a consumption product.

SOURCE: SERVICENOW NEWSROOM

2026-09-29

RESPONSE 02

Workday

WORKFORCE

REORGANIZATION (FORM 8-K)

Cuts about 2.5% of staff (roughly 500 roles on the last reported headcount; the 8-K gives only the percentage), concentrated in Product and Technology, with \$65-80M of charges, two weeks before Workday Rising

Workday did not attribute the cuts to AI and reiterated guidance, but this is the second round of 2026 and it lands two weeks before the company pitches Sana-branded agents at its own conference. The house does not know what the cut funds; buyers should watch whether the October 12-15 announcements come with pricing.

SOURCE: WORKDAY FORM 8-K VIA SEC

2026-10-01

RESPONSE 03

Microsoft

MICROSOFT 365

COPILOT BUSINESS

USAGE-BASED BILLING

Delays default-on usage billing for new CSP licenses (sold through Cloud Solution Provider partners rather than directly) from November 2 to December 1; sets an adjustable 4,000-credit per-user monthly cap, about \$40 at Microsoft's published rate of one cent per credit (the figure Microsoft and CRN both print); excludes 11 countries from the first rollout

The delay and the cap are concessions to partners who could not explain the bill. A \$40 default on top of \$18-32 seats can more than double the per-user cost if agents are used as designed, and the cap can be set to zero, which means the real decision is who in the customer sets it. Finance teams should decide that before December 1. The 11 excluded countries (Australia, Belgium, Brazil, France, Germany, India, Italy, Korea, Netherlands, Poland, Spain, per CRN) include seven EU members, so EU tenants get the delay twice over while HubSpot runs its credit test in EMEA only; buyers in both regions should ask whether they are the pricing lab or the pricing exception.

SOURCE: MICROSOFT LEARN (PARTNER CENTER ANNOUNCEMENTS)

2026-10-01

RESPONSE 04

HubSpot

FLEXIBLE SEATS-AND-
CREDITS PRICING
(EMEA)

New EMEA accounts move from hub subscriptions to one edition plus GTM and Ops seats plus HubSpot Credits, with SKUs consolidated from 73 to 23 (the SKU count is from partner write-ups, not HubSpot's supplement)

Credits meter emails, workflow actions and AI agents at one cent per extra credit with no rollover; existing customers stay on legacy pricing. HubSpot is testing in one region what Microsoft is rolling out globally, and the effective date and SKU count come from partner reports rather than a dated HubSpot announcement, so buyers should confirm terms in writing.

SOURCE: HUBSPOT PRICING SUPPLEMENT

2026-09-30

RESPONSE 05

ZoomInfo

DOUBLEO.AI
ACQUISITION AND
AGENT TEAMS

Acquires a ~10-person orchestration startup for undisclosed terms and ships its architecture as Agent Teams the next day inside the existing subscription

A tuck-in that becomes a shipped feature in 24 hours says the product was ready and the deal was about the team. Bundling orchestration at no extra charge is the data incumbent's answer to agent startups: make the agents free and keep charging for the graph.

SOURCE: ZOOMINFO VIA MARKETSCREENER

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.

2026-09-29

SIGNAL 01

EliseAI

OPERATIONS

\$350M AT \$4B

\$350M round co-led by a16z and Bessemer, doubling the August 2025 valuation per the company's blog; \$200M+ ARR with five years of 100% growth

A vertical operations agent reaching \$200M ARR on housing and clinic workflows is the clearest proof this week that narrow, data-rich operations beat horizontal suites on growth. The unit economics claim (up to 20% higher net operating income for some operators) and the prior valuation come from the company's blog, not the funding release.

SOURCE: ELISEAI VIA GLOBENEWSWIRE; PRIOR VALUATION AND NOI CLAIM FROM THE ELISEAI BLOG POST 'INTELLIGENCE IS GETTING CHEAPER, LIFE SHOULD TOO'

2026-09-28

SIGNAL 02

OpenEvidence

HEALTHCARE

\$250M AT \$15B

CEO confirms \$250M led by a16z and Byers Capital at \$15B, up from \$12B in January; total funding near \$1B

Same round size, 25% higher valuation in eight months, and a system-wide AdventHealth rollout three days later. The repricing is for distribution inside the EHR, which is the asset a competitor cannot buy with a better model.

SOURCE: BECKER'S HOSPITAL REVIEW

2026-09-28

SIGNAL 03

Instinct

OTHER

\$1B AT \$10B

\$1B Series C from Sequoia, Benchmark and Coatue at \$10B, about a month after a \$2.5B mark, for an invite-only personal agent launched in August

A 4x step-up in a month for a product most people cannot use is a bet that the personal-agent category OpenAI's dots just entered will have more than one winner. Enterprises will meet Instinct, like dots and Muse, as inbound traffic from customers' agents; the question is whether a protocol like Decagon's PACT is in place when it arrives.

SOURCE: TECHCRUNCH

2026-09-30

SIGNAL 04

ElevenLabs

SUPPORT

\$300M TENDER AT \$22B

\$300M employee tender co-led by Wellington and T. Rowe Price at \$22B, double February's \$11B primary mark (the valuation set when the company last sold new shares)

A doubled mark on a secondary (a sale of existing shares rather than new ones), with crossover investors (funds that buy both private and public stock), says the voice layer is being priced as infrastructure for every agent that talks. Decagon shipping its own Chord voice model the same week is the counter-signal: the largest CX vendors are internalizing voice rather than renting it.

SOURCE: TECHCRUNCH CITING COMPANY ANNOUNCEMENT

2026-09-30

SIGNAL 05

Cognition

ENGINEERING

Hires ex-Snowflake CRO Chris Degnan; Factory's CEO publicly accuses him of sharing confidential information while a Factory adviser, which all parties deny

Cognition at a \$48B valuation and Factory at \$5B are now fighting over enterprise sales leadership in public, which tells buyers the coding-agent market has moved from model quality to go-to-market. Expect aggressive enterprise terms from both this quarter and read any exclusivity clause carefully.

SOURCE : TECHCRUNCH

PRICING SHIFTS

Seat to outcome, one move at a time.

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

2026-10-01

SHIFT 01

Microsoft

Seat-based → Hybrid

Usage-based billing becomes default-on for new Microsoft 365 Copilot Business CSP licenses on December 1 (delayed from November 2) with an adjustable default cap of 4,000 Copilot Credits, about \$40 at one cent per credit, per user per month on top of \$18-32 seats; initially unavailable in 11 countries listed in the incumbent response above.

SOURCE: MICROSOFT LEARN (PARTNER CENTER ANNOUNCEMENTS)

2026-10-01

SHIFT 02

HubSpot

Seat-based → Hybrid

New EMEA accounts move to one edition plus GTM seats (EUR 20 / 100 / 175) and Ops seats (EUR 50) plus 5,000-15,000 included monthly credits, with extra credits at EUR 0.010 each and no rollover; marketing emails cost from 15 credits each for the first 10,000 sends, stepping down at higher volumes. Existing customers stay on legacy pricing.

SOURCE: HUBSPOT PRICING SUPPLEMENT

2026-09-30

SHIFT 03

Anthropic

Seat-based → Usage-based

Claude for Government GA drops per-seat fees in favour of prepaid usage in fixed increments under a hard not-to-exceed cap, with admin allocation to sub-agencies, per-tier rate limits and dollar caps via SCIM groups, and burndown alerts.

SOURCE: ANTHROPIC

2026-10-01

SHIFT 04

ServiceNow

Seat-based → Usage-based

Flow is free during controlled availability and consumption-priced at GA; existing AI-entitled customers draw on their 'assists' pool with no new license, new customers can buy by credit card. Specific consumption rates were not disclosed.

SOURCE: SERVICENOW NEWSROOM

2026-09-29

SHIFT 05

OpenAI

Seat-based → Hybrid

Pro 500 at \$500 a month with 25x the Plus Work/Codex allowance and exclusive Ultrafast access; Pro 200's included Work/Codex allowance halves from 20x to 10x Plus on October 30 with a one-time \$2,500 usage-credit grant. Marketplace purchases burn down an enterprise's spend commitment at an undisclosed eligible percentage.

SOURCE: OPENAI HELP CENTER

VERTICAL SCORECARD

Who leads each vertical.

Leader-vs-challenger by vertical, useful for procurement shortlists when matching workload to vendor cohort. As of Oct 3, 2026.

VERTICAL	LEADER	CHALLENGER	READ
Commerce	Amazon Seller Assistant via the Selling Partner plugin for Claude	Decagon Agent Modules (payments and collections)	No change at the top; Amazon's plugin documentation did not move this week. Decagon extending CX agents into payments and collections is the first credible challenger from the service side of commerce.
Engineering	Codex on GPT-6.1 Sol (default) with Codex Cloud for Enterprise	Claude Code on Sonnet 5.5 with managed Mods	Codex gets Sol as default, reusable cloud environments and the Ultrafast tier; Claude Code gets Sonnet 5.5 and a Mods extension system. Cognition and Factory are fighting over sales leadership, not features. Cursor, last week's leader on Grok 4.7, is displaced under the same rule applied in the healthcare and support rows: a challenger that shipped in-window product evidence takes the slot from a leader with none. It returns on its own shipped evidence, not on the absence of a competitor's.
Healthcare	OpenEvidence	EliseAI healthcare operations	OpenEvidence at \$15B with AdventHealth system-wide and more than half of US physicians (per Becker's) takes the slot from Basecamp Research, which had no in-window news. EliseAI's clinic-operations unit is the challenger on the administrative side.
Support	Decagon (Voice 3, Agent Modules, PACT)	Flow by ServiceNow	Decagon shipped a voice model, outcome-priced modules and the first consumer-agent protocol in one event; ServiceNow's Flow is the incumbent's consumption-priced answer for the employee service desk. Gemini TTS, last week's leader, is a component, not a package.
Sales	ZoomInfo Agent Teams	Salesforce (OpenAI Marketplace launch partner)	A shipped multi-agent orchestration layer at no extra charge beats a marketplace listing this week. Salesforce's presence among the 32 Marketplace partners is a distribution signal, not a product.
Security	Codex Security Cloud	CrowdStrike and Palo Alto Networks via the OpenAI Marketplace	Codex Security runs on Sol by default and OpenAI claims a 1% rollback rate on its fixes (vendor claim, see Agent Techniques). Elastic's Agentic SOC is the negative example this week with an unpatched prompt-injection finding from PromptArmor.
Legal	Harvey	Legora	Both are OpenAI Marketplace launch partners, which means a law firm's OpenAI commitment can now pay for either. No feature news; the slot is held on distribution.

VERTICAL	LEADER	CHALLENGER	READ
Research	Claude Sonnet 5.5	GPT-6.1 Sol	On GDPval-AA (Artificial Analysis's Elo-style rating of models on paid knowledge-work tasks), Sonnet 5.5's 1844 (Anthropic-reported) sits two points behind Opus 5.5; two Elo points is noise, so the default rests on parity at half the price, not on the score. Sol at one-fifth of Astra's price is the cross-vendor alternative. Opus 5.5 gives up the slot on price, not quality, and Grok 4.7 drops from the challenger slot on no new evidence.

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 Commitment as currency

OpenAI Marketplace (32 partners, commitment burn-down)

AWS and Azure marketplaces (the precedent)

Claude for Government prepaid increments

A model lab's spend commitment can now be spent on third-party software, the way a cloud commitment has been for a decade. The effects are structural: application vendors have a new reason to build on the lab whose commitments their customers hold, the lab gains a view of which partner software enterprises value, and the customer's negotiated discount is now shared across two vendors in a ratio OpenAI has not disclosed. Procurement teams should treat the eligible percentage as a term to negotiate, and application vendors outside the marketplace should expect to be asked why. Anthropic's prepaid-increment model for government is the same idea from the other side: concentrate spend with the lab, let the lab meter it out.

SOURCE: OPENAI HELP CENTER

PATTERN 02 The credit as the universal meter

Microsoft Copilot Credits (4,000 per user per month default cap)

HubSpot Credits (EUR 0.010 each)

ServiceNow assists

Anthropic prepaid increments

This is the fact home for the week's credit exchange rates. Each of the four incumbent pricing moves replaced or supplemented a seat with a credit, and no two credits convert to dollars or to work at the same rate. Microsoft's Copilot Credit is one cent (4,000 credits, about \$40, per user per month as the default cap; Microsoft and CRN both print the figure), HubSpot's is EUR 0.010 with 15 credits buying one marketing email at the first tier, ServiceNow's 'assist' has no disclosed rate, and Anthropic skips the credit and caps dollars directly. The pattern lets vendors raise effective price without raising list price and lets buyers cap spend without buying fewer seats, which is why both sides accept it. The operational consequence is that finance needs a ledger that tracks credit consumption per agent and per workflow across vendors, because the seat count no longer predicts the bill; the engineer who owns the agent gateway is the one who can emit that feed, because the gateway is the only place every metered action passes through. Build it before December 1, when Microsoft's default flips.

SOURCE: MICROSOFT LEARN; HUBSPOT PRICING SUPPLEMENT; SERVICENOW NEWSROOM; ANTHROPIC

PATTERN 03

Consumer agents arrive at the enterprise front door

Decagon Personal Agent Gateway and PACT (A2A plus OAuth 2.0)

OpenAI dots and specialist dots

Instinct (\$10B)

Meta Muse

Three personal-agent products reached scale or funding this week, and one CX vendor shipped a protocol for handling them. PACT lets an enterprise identify a consumer's agent, verify delegated authorization and constrain what it can do, which turns 'a customer's bot called us' from a fraud signal into an authenticated channel. The architectural shift is that the customer-facing surface now has to serve two clients, humans and their agents, with different authentication and different rate limits. Any company with a support or commerce front end should ask its vendors how they will tell a delegated agent from a scraper, because dots is on Pro and Business Premium now and Instinct just raised a billion dollars to put more agents in the queue.

SOURCE: DECAGON

PATTERN 04

Vertical data compounds; horizontal headcount reorganizes

EliseAI (\$200M+ ARR, \$4B)

OpenEvidence (\$15B, AdventHealth)

Workday (2.5% cut, \$65-80M charges)

The two application companies that raised at the highest multiples this week sell narrow operations agents sitting on proprietary interaction data: leasing and clinic workflows at EliseAI, citation-grounded clinical answers inside the EHR at OpenEvidence. The horizontal suite that reorganized sells a system of record and is cutting product headcount two weeks before it pitches agents at its own conference. The lesson for buyers is that the agent's model is the least defensible part of the stack; the data it was trained and tuned on, and the workflow it is embedded in, are what justify the price. Evaluate vertical agents on the data moat and the integration depth, and evaluate suite agents on whether they come with pricing.

SOURCE: ELISEAI; BECKER'S HOSPITAL REVIEW; WORKDAY FORM 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 6

WATCH 01

Flow by ServiceNow release to current North American customers

The first real deployment data for a consumption-priced service desk that needs no implementation project. Watch for disclosed consumption rates; the GA price is the number that decides whether it is a trial or a product.

OCT 12-15

WATCH 02

Workday Rising: agentic HR and finance announcements and whether they ship with pricing

Two weeks after a 2.5% cut concentrated in product and technology, Workday's conference tells buyers whether the suite's agents are a roadmap or a priced SKU, and whether Sana is metered in seats or credits.

OCT 26 AND OCT 30

WATCH 03

Custom GPT creation ends for Enterprise; Pro 200's Codex allowance halves

Two forced migrations in one week: Enterprise builders move to dots and Spaces, Pro 200 users choose between Pro 500 at \$500 a month and moving work to Sonnet 5.5. Both are tests of how much switching cost OpenAI's platform has accumulated.

DEC 1

WATCH 04

Microsoft Copilot Business usage-based billing becomes default-on

The 4,000-credit cap is adjustable to zero and most tenants have not decided who sets it. The partner channel's readiness, and the first invoices, decide whether the hybrid model sticks or slips again.

Q4 2026

WATCH 05

OpenAI Marketplace eligible percentage and partner-fee disclosure

Without the burn-down ratio, no customer can price a Marketplace purchase against its negotiated discount and no partner can price its margin. The first disclosure, or the first leaked term sheet, resets every partner's economics.

Q4 2026

WATCH 06

Decagon PACT adoption by a second CX vendor or a personal-agent platform

A protocol with one implementer is a feature. If Sierra, Salesforce or OpenAI's dots adopt PACT or publish a competing delegated-authorization scheme, the enterprise front door gets a standard; if not, every company builds its own agent gate.

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 24

Week 40 of 2026 · October 3, 2026

WEB

brianletort.ai/industry/applications

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