

THE BIG READ

The system of record is where the agent surplus lands: Salesforce tolls the call, Zendesk the resolution, Microsoft the model, and Gainsight gives the call away

KEY SIGNALS THIS PERIOD

8

Verticals tracked this period

7

Vertical packages shipped

3

Incumbent SaaS responses

3

Startup signals

AUTHOR

Brian Letort
BrianLetort.AI

PUBLISHED

October 24, 2026
Issue 27 · 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.

The vendors whose data agents touch are moving their meters to the unit that grows with agent activity, and the week's four meters say it in four different units: the call, the resolution, the model and the account. None of them was a reaction to this week's cache price cut (each was planned before October 20; the connection to cheaper turns is this issue's reading, not the vendors' stated reason). Salesforce will from about November meter every successful MCP (Model Context Protocol, the open standard agents use to call tools) or API call by a registered third-party agent (Claude and ChatGPT are named) as one 'Headless Platform Interaction' drawing Flex Credits, with the multiplier listed as 'TBA'; The Information's October 22 report is what made a help-article notice the week's event, adding the consultants' survey (most of ~150 enterprises already run Anthropic or OpenAI agents on Salesforce data at no charge and may move data out if the toll is high). A standard Agentforce action costs 20 Flex Credits, about \$0.10 at list, so the open question is where between 1 and 20 credits the external call lands; the sales entry below is the fact home for the terms.

Zendesk's Employee Service AI Agents, GA from November 3 to December 1, include 10 (Growth and Professional) or 15 (Enterprise) automated resolutions per agent per month, after which the agent pauses or the buyer pays \$0.99 per additional resolution; about 300 early-access customers ran 6 million workflow executions in five weeks (~4,000 per customer per week; a vendor adoption figure), which Zendesk converts to 500,000 hours saved at five minutes each, a vendor estimate. Microsoft's Message Center notice MC1519122 makes GPT-6 Astra and Claude Fable 5.1 usage-billed 'UBB models' in Copilot Chat, Word, Excel and PowerPoint from early to mid-November for tenants that enable usage-based billing and allocate credits; the same day a Partner Center offer through December 31 gives 'Buy 5 Get 5' Copilot Studio capacity packs to customers taking 300 or more new annual E7 seats (Microsoft 365's top licence tier) with pay-as-you-go enabled. And Gainsight pre-announced that Gainsight CS drops per-seat licensing in 2027 for value-based pricing that scales with the customer's ARR (annual recurring revenue) and the number of customers managed, with MCP access to any external agent free of charge and Staircase AI bundled.

The reading this issue adds, and the reason the four belong in one story, is that the surplus from cheaper agent turns (cache pricing cut the input side of a turn to roughly a tenth of list; Agent Techniques carries the arithmetic) is not landing with the lab that cut the price or with the GPU, but with the vendor that owns the data the agent calls: the scarce complement captures the rent. The toll-road frame itself is prior art, not this issue's: Deloitte described systems of record 'tollgating' agent access and StepTo wrote that the response is to own the read path by syncing the data into a store you control (both per the originality review; outside the graded packet: picked up after source grading closed and not checked by this publication). That is the predictable response here too: agent builders stop crossing the boundary and replicate the CRM records they need into a cache or lake they control, which The Information's consultants already report customers considering. What this issue adds is the price test (one credit per call is noise, five is a budget line, twenty is a migration trigger) and a reading of the quarter's other move: Salesforce agreed the same week to pay ~\$5.66B enterprise value (\$47.50 a share, a 58% premium) for Braze, an event-streaming customer-engagement layer with ~2,770 customers and trillions of messages a year (vendor figure). Nothing in either announcement links the two, and the deal was negotiated before the toll was reported; read together, though, the data layer the toll protects is the data layer the acquisition deepens.

Gainsight's free MCP is the counter-position, and the inference this issue draws (Gainsight has stated no rationale; the change is a pre-announcement with details in 2027) is that a vendor whose new price scales with the customer's ARR wants every agent in the building reading its retention signals, because the more the agent uses the data the more valuable the account the price is indexed to. Zendesk's allowance-and-pause design is the third pattern: an outcome meter with a hard stop, so the buyer sees the ceiling before the bill, and the \$0.99 sits below the \$1.50-\$2.00 on Zendesk's public employee-service price page. Microsoft's is the fourth: the model, not the seat or the action, becomes the metered object, with the admin allocating credits per user and the Copilot Studio offer pushing pay-as-you-go billing as the default commercial posture.

Behind all four sit the week's two enterprise-software quarters: SAP monetises its 300 Joule agents through subscription pull-through rather than a meter (the finance entry carries the quarter), and IBM's Consulting gross margin fell to 27.4% from 28.9% in Q2 (29.3% a year earlier) with AI at 35% of backlog, the pressure that pushes integrators toward asset-based pricing (the IBM entry carries the numbers).

On the frontier-lab side of the application layer the week was about cost controls as product rather than new workloads. Anthropic made the unified Claude experience (chat plus Cowork, with agentic work running in Anthropic-managed, egress-off cloud sandboxes) GA for Claude Enterprise, off by default for existing organisations and on automatically from November 16, with four admin controls: an Efficient mode that pins the organisation to Sonnet and Haiku (new sessions start on Sonnet 5.5 at low effort), model entitlements, spend limits at organisation, group and member level, and pooled group budgets; Anthropic's claim that Opus 5.5 costs 40% less to run than Opus 5 and Sonnet 5.5 up to 30% less than Sonnet 5 is vendor-reported, and ZDR (zero data retention) for Claude Code, HIPAA-ready configurations and CMEK (customer-managed encryption key) organisations are excluded from the GA. Claude Security moved to public beta for Enterprise with continuous repository scanning, fix tracking (findings auto-close when the fix lands and reopen if it recurs), multi-repo projects, API and MCP access, and \$25,000 of free usage credits per month for two months through December 31; 'over one million vulnerabilities found since April' is a vendor claim. OpenAI added per-project service-tier restrictions and a 12-hour cache option, and Opus 5.5 Fast Mode went GA on Google Cloud at \$8 / \$40, a speed tier at twice list (The Model Pulse carries the model side). Legora closed a \$550M Series E at \$8.8B post-money with ARR past \$200M (the second \$100M in under six months) and in-house legal teams now more than 40% of new customers, so both legal-AI leaders (Harvey at \$15.5B a month earlier) are capitalised to subsidise land-and-expand for years.

What to do differently. If you buy agents: insist on the allowance-and-pause pattern Zendesk shipped (a cap the administrator sets, enforced by the product, with a visible path when it binds) in every agent contract regardless of meter, and price every proposal in three units at once (seat, meter, outcome); for Microsoft, allocate UBB credits per user deliberately rather than letting the mid-November rollout set the default, and model the December 1 standalone SKUs. If you run enterprise architecture: decide before Salesforce publishes its multiplier whether your agents cross MCP boundaries you do not own or cache the records they need, and write the data-egress and replication rights into the next CRM renewal, because the toll says the data layer is where the vendor expects to be paid. If you buy from an integrator: IBM's 27.4% consulting gross margin with AI at 35% of backlog says the vendor's AI work is dilutive, so expect pressure toward asset-based pricing and negotiate the IP and reuse terms now. If you run a SaaS company: Gainsight's ARR-indexed price and free MCP is the model for a vendor

whose data gets more valuable the more agents read it; Salesforce's toll is the model for one whose data is the moat. Pick one on purpose.

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.

2026-10-22

PACKAGE 01

Salesforce

SALES

INCUMBENT SAAS

USAGE

Headless Platform Interactions (external-agent metering on Flex Credits)

Reported October 22: every successful MCP or API call by a registered third-party agent acting on Salesforce data becomes a billable Headless Platform Interaction from about November (notice dated September 17); the multiplier is unpublished

This is the fact home for the Salesforce meter. Salesforce's help article is dated September 17 and sfdcdevelopers.com covered the mechanics on October 9; The Information's October 22 report is what made it the week's event. The article's mechanism: agents (Claude and ChatGPT named) must be registered, each successful call draws Flex Credits as one HPI, the multiplier is 'TBA', 30 days' notice precedes metering, new customers have three months to register and existing customers migrate at renewal, and MCP server use requires registration within three months of notice. Flex Credits list at \$500 per 100,000 and a standard Agentforce action is 20 credits, so the open variable is credits per external call. Consultants say most of ~150 enterprises surveyed already run external agents on Salesforce data for free and may move data out if the toll is high, which is the second-order effect the weekly's synthesis names. For a CRM owner the action is to price the renewal in three scenarios (1, 5 and 20 credits per call) against the agent call volume you can already measure, and to get the egress and replication rights in writing.

SOURCE: SALESFORCE HELP (UNDERSTAND AIFORCE IMPACT); THE INFORMATION; BENZINGA (SYNDICATED)

2026-10-19

PACKAGE 02

Zendesk

SUPPORT

INCUMBENT SAAS

OUTCOME

Employee Service AI Agents and Advanced Action Flows

IT, HR and workplace agents priced by the resolution: 10-15 included per agent per month, then the agent pauses or the buyer pays \$0.99 each; GA November 3 to December 1

The cleanest outcome meter of the week and the first with a hard stop designed in: Growth and Professional plans include 10 automated resolutions per AI agent per month and Enterprise plans 15, after which the agent pauses until the next cycle unless extra resolutions are bought at \$0.99, below the \$1.50-\$2.00 on Zendesk's public employee-service price page. The agents work in the help centre, Slack and Teams across onboarding, access approvals, laptop lifecycle and verification letters with 20+ permission-aware knowledge connectors. The usage datapoint is real and the savings figure is not: ~300 early-access customers ran 6 million workflow executions in five weeks (~4,000 per customer per week; vendor-reported), and the 500,000 hours saved is Zendesk's estimate at five minutes per execution. For an IT or HR leader the allowance-and-pause pattern is the thing to copy into every agent contract, because it makes the ceiling visible before the bill; the Futurum question of whether this pries IT work from ServiceNow is the row to watch.

SOURCE: ZENDESK NEWSROOM; ZENDESK HELP (GA ANNOUNCEMENT); FUTURUM GROUP

2026-10-21

PACKAGE 03

Microsoft

OPERATIONS

INCUMBENT SAAS

HYBRID

UBB models in Microsoft 365 Copilot (GPT-6 Astra, Claude Fable 5.1) and the E7 + Copilot Studio 'Buy 5 Get 5' offer

Premium models become the metered object in Copilot Chat, Word, Excel and PowerPoint from early to mid-November for tenants with usage-based billing and allocated credits; a CSP (Cloud Solution Provider, Microsoft's partner-resale channel) offer through December 31 pushes pay-as-you-go as the default posture

This is the fact home for Microsoft's November meter. Message Center notice MC1519122 says licensed Copilot users in tenants that enable UBB and allocate credits can select GPT-6 Astra or Claude Fable 5.1 from the model selector, with a user credit-request flow to follow and cost management in the Agent 365 pane; the rollout runs early to mid-November. The Partner Center offer (October 21 to December 31) matches five annual Copilot Studio capacity packs for customers buying 300 or more new annual E7 seats with pay-as-you-go billing on an Azure subscription, Work IQ admin controls for Business Applications reach public preview by November 17, and standalone Copilot SKUs with no base-subscription prerequisite arrive December 1 (existing offers were renamed '(add-on)' from October 16). For an IT finance owner the change is that the model choice is now a cost decision made per user by an administrator; allocate the credits deliberately before the rollout sets the default, and treat December 1 as the date the seat and the meter become separable.

SOURCE: MICROSOFT 365 MESSAGE CENTER ARCHIVE (MC1519122, MC1518036); MICROSOFT LEARN (PARTNER CENTER OCTOBER 2026 ANNOUNCEMENTS)

2026-10-21

PACKAGE 04

Anthropic

SECURITY

FRONTIER LAB

USAGE

Claude Security (continuous scanning, fix tracking, API and MCP)

Rebuilt Claude Security enters Enterprise public beta with baseline and change-driven repository scanning, fixes auto-closed when they land, multi-repo projects, and \$25,000 of monthly credits for two months

The frontier lab's security product became a continuous service rather than a scan: connected GitHub (including Enterprise Server) and GitLab repositories are baseline-scanned, changes re-scanned by risk and complexity, findings marked fixed when the fix lands and reopened if the flaw recurs, with projects spanning up to eight repositories to catch cross-service flaws and findings exposed over API and MCP. It runs on the full model family including Claude Mythos without Cyber Verification Program access, and eligible organisations get \$25,000 of free usage per month through December 31, after which usage bills at list. 'Thousands of organisations' and 'over one million vulnerabilities found since the April beta' are vendor claims. For a CISO the comparison is with Codex Security Cloud (the scorecard leader) on the same two axes, coverage of the change stream and the false-positive rate, and Trail of Bits' regression self-audit (Agent Techniques) is the evidence on what agent patches break.

SOURCE: ANTHROPIC (CLAUDE.COM); HELP NET SECURITY; CYBER SECURITY NEWS

2026-10-23

PACKAGE 05

Anthropic

OTHER

FRONTIER LAB

USAGE

Unified Claude for Enterprise (chat plus Cowork) with Efficient mode and admin cost controls

Chat and Cowork merge for Claude Enterprise with agentic work in egress-off managed sandboxes; four admin cost controls ship with it; off by default until November 16, then on

The frontier lab's knowledge-work product gained its cost-control layer in the same release as its agentic one. Unified Claude carries connectors, skills and plugins across chat and Cowork and runs agentic tasks in Anthropic-managed cloud sandboxes with egress off; pricing stays usage-based, and administrators get Efficient mode (restricts the organisation to Sonnet and Haiku, with new sessions on Sonnet 5.5 at low effort), model entitlements, spend limits at organisation, group and member level, and pooled group budgets. Anthropic's 40% cheaper Opus 5.5 and up-to-30% cheaper Sonnet 5.5 are vendor claims. ZDR for Claude Code, HIPAA-ready configuration and CMEK organisations are excluded from the GA, and the auto-on date of November 16 is the one that matters for existing tenants: an administrator who does nothing gets the unified experience, and one who sets Efficient mode caps the model tier. For an IT leader the move is to decide the entitlement and spend-limit policy before November 16 rather than after the first bill.

SOURCE: ANTHROPIC (CLAUDE.COM); CLAUDE HELP CENTER (ENTERPRISE ADMIN SETUP GUIDE); NEW RUNTIME

2026-10-21

PACKAGE 06

SAP

FINANCE

INCUMBENT SAAS

HYBRID

Joule Work (40 assistants, 300 agents) inside RISE and cloud subscriptions

SAP's first EUR 10B quarter with cloud backlog of EUR 24.1B (+26% cc, constant currency); nine of the ten largest deals include AI; Joule is monetised through subscription pull-through rather than a separate agent meter

The one large incumbent not building a toll booth this week. SAP reported current cloud backlog of EUR 24.14B, cloud revenue of EUR 6.67B (+26%), non-IFRS operating profit of EUR 3.03B at a 29.5% margin and free cash flow up 36%, lifted its 2026 cloud outlook to the upper end of the range, and said AI was the main driver of cloud migrations with nine of the top ten of its 50 largest deals including AI and Business Data Cloud. Joule Work now spans 40 assistants and 300 agents with 70+ models available in Joule Studio, and SAP monetises them through RISE and cloud subscription pull-through, with the 0.02 AI Units per action rate from W41 as the only metered component. For a CFO or CIO negotiating with SAP the window is now: entitlements for Joule agents belong inside the cloud contract while SAP is using AI to pull migrations rather than to meter them, and the Business Data Cloud terms decide which external agents can read the finance data without a Salesforce-style toll later.

SOURCE: SAP INVESTOR RELATIONS (Q3 2026 QUARTERLY STATEMENT); REUTERS VIA STREETINSIDER; INVESTING.COM (EARNINGS CALL TRANSCRIPT)

2026-10-20

PACKAGE 07

Gainsight

SALES

INCUMBENT SAAS

HYBRID

Gainsight CS value-based pricing (ARR and customers managed) with free MCP access; Atlas AI-native renewals

Gainsight drops per-seat licensing for a price indexed to the customer's ARR and customers managed (available 2027), makes MCP access for any external agent free, and reports Atlas managing >\$75M of customer ARR on outcome-tied contracts

The counter-position to Salesforce's toll, from a vendor whose data gets more valuable the more agents read it. Gainsight CS moves from seats to a price that scales with the buyer's own ARR and the number of customers managed in Gainsight, with Staircase AI bundled and MCP access to any external agent at no charge; Agent Studio (natural-language reusable skills with human approval on every action) is in beta. Atlas, the services business launched in May that pairs agents with human Renewal Agent Managers on outcome-tied contracts, now manages more than \$75M of customer ARR in long-tail renewals with 7,000+ agentic interactions, which is the one services datapoint of the week where the vendor's revenue is tied to the outcome rather than the hour. The event is a pre-announcement at Pulse Europe (grade 3; pricing details arrive in 2027). For a post-sale leader the question to ask Gainsight is how the ARR index is measured and capped; for a SaaS operator it is the model for pricing when agents, not seats, consume the product.

SOURCE: GLOBENEWSWIRE (GAINSIGHT PRESS RELEASE); GAINSIGHT BLOG

INCUMBENT RESPONSES

How the SaaS estate is answering.

3 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.

2026-10-20

RESPONSE 01

Salesforce

BRAZE ACQUISITION
(~\$5.66B ENTERPRISE
VALUE; \$47.50 A SHARE
CASH)

Definitive agreement to acquire Braze at a 58% premium, closing in Salesforce FY2028 with no change to FY27 guidance; \$300M reverse termination fee, \$191.4M Braze break fee

The data layer the Headless Platform Interaction toll protects is the data layer this acquisition deepens: Braze brings ~2,770 customers (Canva, Instacart, Venmo), an event-streaming architecture handling trillions of messages a year and BrazeAI decisioning that Salesforce says will plug into Data 360, Customer 360, Agentforce and SIForce. This issue's read, not anything either announcement says, is that the two belong to one strategy: own the real-time B2C engagement stream, and charge external agents to read the records it produces. It extends the run of AI-software take-outs after Proofpoint-Varonis (\$6.3B) in W42.

SOURCE: SALESFORCE NEWSROOM; SEC EDGAR (SALESFORCE 8-K; BRAZE DEFA14A); CNBC

2026-10-21

RESPONSE 02

IBM

Q3 2026: CONSULTING
AI MIX AND THE
DROPPED GENAI BOOK-
OF-BUSINESS METRIC

AI is ~30% of Consulting revenue and 35% of Consulting backlog; Consulting revenue flat at \$5.34B, signings -2%, gross margin down to 27.4% from 28.9% in Q2 and 29.3% a year earlier; the GenAI book-of-business figure omitted for a third straight quarter

The integrator's AI work is now the majority of new backlog and dilutive to the vendor, which is why the CFO described a shift to 'services as software' and cited an OpenAI strategic partnership and an expanded Google Cloud relationship. For buyers of SI-led AI programmes the margin compression is the leverage: expect pressure toward asset-based and outcome-based pricing from the vendor, and negotiate IP, reuse and model-portability terms before they arrive. Software (\$7.80B, +8%, ARR ~\$25B) beat consensus, so the AI pressure sits in services, not products.

SOURCE: IBM INVESTOR RELATIONS (3Q26 PRESENTATION AND TRANSCRIPT); CNBC

2026-10-21

RESPONSE 03

SAP

Q3 2026 CLOUD BACKLOG
AND JOULE WORK
ROLLOUT

Current cloud backlog EUR 24.1B (+26% cc), first EUR 10B revenue quarter, 2026 cloud revenue guided to the upper end of EUR 25.8-26.2B; Joule Work to all customers

SAP's response to the agent wave is to use it as the migration driver rather than the meter: nine of the ten largest deals included AI and Business Data Cloud, and Joule's 300 agents ride the RISE subscription. The finance entry carries the negotiation point; for the scorecard it keeps SAP Joule as the finance leader on the strength of a backlog figure, not a product release.

SOURCE: SAP INVESTOR RELATIONS (Q3 2026 QUARTERLY STATEMENT); REUTERS VIA STREETINSIDER

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.

2026-10-21

SIGNAL 01

Legora

LEGAL

**\$550M AT \$8.8B
POST-MONEY**

\$550M Series E led by Atreides at \$8.8B post-money (from \$5.6B in April); ARR past \$200M (company-reported) after going \$100M to \$200M in under six months; in-house legal teams >40% of Q3 new customers

Both legal-AI leaders are now capitalised to subsidise land-and-expand for years (Harvey raised \$550M at \$15.5B a month earlier), and the customer mix is shifting from law firms (more than half the AmLaw 50) to corporate legal departments, which is where last week's seat-less Intake products pointed. For a legal-ops team the move is to flag the round to whoever owns the renewal, because multi-year price locks and agentic-intake commitments are available while the two compete for the in-house logo; 140,000+ professionals at 2,400+ firms and teams across 80+ markets is the installed base the price will be defended on.

SOURCE: LEGORA NEWSROOM; TECHCRUNCH; BLOOMBERG LAW

2026-10-21

SIGNAL 02

Recursive

ENGINEERING

First model, Flint, released under Apache 2.0 as a 39.8B / ~2.5B-active coding MoE sized to run on a laptop, in Ollama the same day

A new lab entered the coding-agent market with a model whose distribution is the product: no hosted price, no licence gate, and a size chosen to fit a developer's own machine, which routes around every meter this issue documents. The Model Pulse carries the row and the vendor-reported scores (89.7% SWE-bench Pro v2, unreproduced); for an engineering leader the signal is that the on-device coding tier now has an unencumbered entrant to test against Codex and Claude Code on cost per merged change.

SOURCE: HUGGING FACE MODEL CARD; RECURSIVE; TECHCRUNCH

2026-10-22

SIGNAL 03

Subquadratic

OTHER

SubQ Snap 1.1 ties for the lead on Artificial Analysis's Decision Index at 65 with a 384K context, When2Call up from 75.2 to 96.2, and a 20% price cut to \$0.04 per million tokens

The decision-model category (small models that pick an action or a tool rather than write prose) now has three vendors tied at the top and a price point two orders of magnitude below the chat tier, which is the layer that will sit in front of every metered system of record deciding whether to make the call at all. Knowledge accuracy ranges from 11.6 to 31.6 across the three leaders, so the choice is a workload decision, not a leaderboard one; The Model Pulse carries the table and keeps the category deferred on the tree.

SOURCE: SUBQUADRATIC; ARTIFICIAL ANALYSIS DECISION INDEX

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.

2026-10-22

SHIFT 01

Salesforce

Seat-based → Usage-based

External agents acting on Salesforce data move from unmetered access under the customer's seats to per-call metering on Flex Credits from about November, multiplier unpublished; the sales entry carries the terms.

SOURCE: SALESFORCE HELP (UNDERSTAND AIFORCE IMPACT); THE INFORMATION

2026-10-19

SHIFT 02

Zendesk

Seat-based → Outcome-based

Employee-service agents are packaged by automated resolution rather than by agent seat: 10 (Growth, Professional) or 15 (Enterprise, Enterprise Plus) resolutions per AI agent per month included, the agent pauses when the allowance is exhausted, and additional resolutions cost \$0.99 each, below the \$1.50-\$2.00 on Zendesk's public employee-service price page.

SOURCE: ZENDESK HELP (GA ANNOUNCEMENT)

2026-10-21

SHIFT 03

Microsoft

Seat-based → Hybrid

Premium models (GPT-6 Astra, Claude Fable 5.1) in Copilot Chat, Word, Excel and PowerPoint become usage-billed on admin-allocated credits on top of the Copilot seat from early to mid-November; the Copilot Studio 'Buy 5 Get 5' offer requires pay-as-you-go billing to qualify; standalone Copilot SKUs decouple the seat from a base subscription on December 1; the operations entry above carries the sources.

SOURCE: MICROSOFT LEARN (PARTNER CENTER OCTOBER 2026 ANNOUNCEMENTS); MICROSOFT 365 MESSAGE CENTER ARCHIVE (MC1519122)

2026-10-20

SHIFT 04

Gainsight

Seat-based → Hybrid

Gainsight CS drops per-seat licensing for value-based pricing that scales with the customer's ARR and the number of customers managed in Gainsight (available 2027), with Staircase AI bundled and MCP access for any external agent free of charge; pre-announced at Pulse Europe with pricing details to follow.

SOURCE: GAINSIGHT BLOG; GLOBENEWSWIRE

VERTICAL SCORECARD

Who leads each vertical.

Leader-vs-challenger by vertical, useful for procurement shortlists when matching workload to vendor cohort. As of Oct 24, 2026.

VERTICAL	LEADER	CHALLENGER	READ
Commerce	Amazon Seller Assistant	Personal Agent Protocol retail partners (Sierra, Meta, Walmart, Visa, Mastercard)	Unchanged in order. Salesforce's agreement to buy Braze (~\$5.66B, ~2,770 B2C customers, closing in FY2028) is the week's commerce-adjacent move and enters the note rather than the row until it closes: it puts a real-time engagement stream under Agentforce and behind the Headless Platform Interaction toll, which is a different route to the consumer than the PAP partners' agent-to-merchant protocol.
Engineering	Codex (on GPT-6.1 Sol)	Claude Code (on Sonnet 5.5)	Unchanged. The Model Pulse held its coding leader and entered Kimi K3.1 Pro as a provisional challenger on the independent Terminal-Bench 4.0 result (67%); this row tracks the shipped product, and neither Codex nor Claude Code offers K3.1 Pro as a routable model. Both products shipped runtime-control releases this week (Agent Techniques carries them). Flint is the unencumbered on-device entrant to test; GitHub's usage API now attributes requests to claude_code, codex and opencode, which is the first vendor-side measurement of which harness enterprises run.
Healthcare	Abridge	Suki	Unchanged. The week's healthcare datapoint is a benchmark, not a product: Artificial Analysis's HealthBench Professional leaderboard (GPT-6 Astra 66.0% length-adjusted; Claude Opus 5.5 first raw and sixth adjusted) uses OpenAI's benchmark, length penalty and a GPT-5.4 judge, so a procurement citing it should state the methodology; The Model Pulse carries the table. GPT-6 Luna at 61.9% for \$0.004 per task is the cost floor for a documentation assistant's model layer.
Support	Sierra	Salesforce Agentforce	Unchanged in order; Zendesk enters the note on the strength of the week's cleanest outcome meter. Its Employee Service AI Agents (10-15 resolutions included, \$0.99 above, agent pauses at the ceiling; GA November 3 to December 1) extend the support row into IT and HR requests, with 6 million early-access executions in five weeks as the vendor-reported usage figure. Agentforce holds the challenger slot on the W42 HMRC award (the UK tax authority's £2 billion contract); Salesforce's HPI toll applies to external agents acting on its data, not to Agentforce itself.

VERTICAL	LEADER	CHALLENGER	READ
Sales	ZoomInfo	Salesforce Agentforce with DocuSign IAM (its agreement platform)	Unchanged in order; the row's economics changed. From about November every external agent reading or writing Salesforce records pays a Headless Platform Interaction toll with an unpublished multiplier, which affects every sales-agent vendor built on Salesforce data including the leader; Gainsight went the other way with free MCP and an ARR-indexed price for its customer-success platform. The sales row is now the clearest place to watch whether agent builders cross the MCP boundary or replicate the data.
Security	Codex Security Cloud	Claude Security	Challenger changes from Elastic AlertZero to Claude Security on the Enterprise public beta: continuous repository scanning, change-driven re-scans, fix tracking with auto-close and reopen, multi-repo projects, API and MCP access, and \$25,000 of monthly credits for two months. 'Over one million vulnerabilities found' is a vendor claim. The leader holds on coverage and its Patch the Planet record (OpenAI's open-source patching programme); Trail of Bits' self-audit finding that about one in six of its merged agent-assisted patches carried a regression (Agent Techniques is the fact home) is the caveat for both. Elastic AlertZero moves to the note on the strength of its W41 launch.
Legal	Harvey	Legora	Unchanged in order; Legora's \$550M Series E at \$8.8B (ARR past \$200M, a press-release figure; in-house teams >40% of new customers) closes the funding gap to Harvey's \$15.5B and shifts the contest to corporate legal departments, where both launched seat-less Intake products in W42. For buyers the two are now capitalised to subsidise land-and-expand for years; flag the round to whoever owns the renewal.
Research	Claude Dashboards	Gemini agent	Unchanged. The unified Claude Enterprise experience (chat plus Cowork, GA October 23, auto-on November 16) is the leader's platform change and carries the first admin cost controls (Efficient mode, entitlements, spend limits, pooled budgets) in this row; the Gemini agent's general availability and metering remain the pending event for the challenger.

VERTICAL	LEADER	CHALLENGER	READ
Finance	SAP Joule (per-action AI Units)	Workday agentic finance (Flex Credits)	Unchanged in order; SAP's quarter (EUR 24.1B cloud backlog, nine of ten largest deals with AI, Joule Work's 300 agents to all customers) is the leader's week, and its decision to monetise agents through subscription pull-through rather than a separate meter is the contrast with Workday's Flex Credit rate card from W42. IBM's AI-heavy consulting at a 27.4% gross margin is the integrator-side datapoint for finance transformation programmes.

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

MCP as a billing boundary: the protocol call becomes the metered unit, and the system of record decides whether to toll it or give it away

Salesforce: every successful MCP or API call by a registered external agent is one Headless Platform Interaction on Flex Credits, multiplier unpublished

Gainsight: MCP access for any external agent free, with the platform price indexed to the customer's ARR instead

Anthropic: Claude Security findings exposed over API and MCP inside the Enterprise plan's usage billing

Claude Code MCP Tasks: server-side long-running jobs that survive the client (Agent Techniques)

The same protocol produced opposite pricing decisions in one week. This issue's reading of the difference, which neither vendor has stated, is what the data is worth to the vendor when agents read it: Salesforce's records are the moat, so the call is tolled; Gainsight's retention signals get more valuable the more agents act on them, so the call is free and the price moves to the account's ARR. Both are inferences from the pricing design, not disclosed motives. For an architect the consequence is that the MCP server you connect to now carries a commercial term as well as a schema, and the agent's decision to call it is a cost decision; the response where the toll is high is the one Deloitte ('tollgating') and StepTo ('own the read path') described before this notice (both per the originality review, outside the graded packet): replicate the records into a store you control, which The Information's consultants report customers already considering and which the weekly's synthesis names as the second-order effect. Expect the next generation of enterprise MCP servers to ship with a rate card, and expect data-out and replication rights to become the contested clause in CRM renewals.

SOURCE: SALESFORCE HELP (UNDERSTAND AIFORCE IMPACT); GAINSIGHT BLOG; ANTHROPIC (CLAUDE.COM)

PATTERN 02

Allowance-and-pause: outcome meters with a hard ceiling the buyer sees before the bill

Zendesk: 10-15 resolutions per agent per month included, the agent pauses at the ceiling, \$0.99 per additional resolution

Microsoft: admin-allocated UBB credits per user with a credit-request flow for the user who runs out

Anthropic: spend limits at organisation, group and member level with pooled group budgets; Efficient mode pins the tier

OpenAI: per-project service-tier restrictions at the organisation level

Four vendors shipped the same control in four products: a cap set by the administrator, enforced by the product, with an explicit path (buy more, request credits, raise the limit) when it binds. This is the application layer's answer to the cost unpredictability that has kept agent pilots from scaling, and it changes who decides: the administrator sets the ceiling, the agent or the user hits it, and the vendor's revenue above the allowance is a visible decision rather than a surprise. For a buyer the pattern is the thing to demand in every agent contract regardless of the meter (resolution, credit, token or call), and Zendesk's version, where the agent stops rather than the bill growing, is the strictest of the four. For a vendor it is the mechanism that makes usage pricing sellable to a CFO.

SOURCE: ZENDESK HELP (GA ANNOUNCEMENT); MICROSOFT 365 MESSAGE CENTER ARCHIVE; CLAUDE HELP CENTER; OPENAI DEVELOPER CHANGELOG

PATTERN 03

The model as the metered object: premium model choice becomes a per-user cost decision made by an administrator

Microsoft: GPT-6 Astra and Claude Fable 5.1 as UBB models selectable per user with allocated credits

Anthropic: model entitlements and Efficient mode (Sonnet and Haiku only; new sessions on Sonnet 5.5 at low effort)

Anthropic on Google Cloud: Opus 5.5 Fast Mode at \$8 / \$40, a speed tier at twice list

Kimi K3.1 Pro Highspeed at \$4 / \$20 (The Model Pulse)

The meters of W42 priced the action and the token; this week's price the model and the speed. Microsoft's selector puts the premium model behind credits the administrator allocates, Anthropic's entitlements let an administrator remove Opus from an organisation's menu, and two labs shipped speed tiers at exactly twice list, so a 'faster' toggle is now a 2x multiplier on the same card. For an IT finance owner this is the first time the model choice shows up as a line the administrator controls rather than a default the vendor sets; the decision to make before the mid-November rollouts is which users get which models, because the default is the vendor's. For a vendor it means the model name is becoming a SKU, and the Model Pulse's cost-per-task table (band members at the same list differing ~3x per completed task) is the reason to check cost per task rather than the card.

SOURCE: MICROSOFT 365 MESSAGE CENTER ARCHIVE (MC1519122); CLAUDE HELP CENTER (ENTERPRISE ADMIN SETUP GUIDE); ANTHROPIC (CLAUDE.COM)

PATTERN 04

AI-heavy services are margin-dilutive to the vendor, which pushes integrators toward asset-based and outcome-tied pricing

IBM: AI ~30% of Consulting revenue and 35% of backlog with gross margin down to 27.4% and the GenAI book-of-business figure dropped for a third quarter

Gainsight Atlas: agents paired with human Renewal Agent Managers on outcome-tied contracts, >\$75M of customer ARR under management

SAP: nine of ten largest deals include AI, monetised through subscription pull-through rather than a services line

The week's two enterprise-software quarters and one services pre-announcement describe the same economics from three positions. The integrator whose AI work is the majority of new backlog is seeing margin compress and has stopped publishing the metric it used to lead with; the software vendor whose AI drives migrations is booking it as cloud backlog at a 29.5% operating margin; and the customer-success vendor that pairs agents with humans is pricing the outcome rather than the hour. For a buyer of SI-led programmes the compression is leverage now and a pricing change later: expect proposals that bundle the vendor's own agents and assets at a fixed fee tied to an outcome, and negotiate the IP, reuse and model-portability terms before the model changes. For the application layer it is the first evidence that the agent wave moves services revenue into software and software revenue into usage, which is where every meter in this issue sits.

SOURCE: IBM INVESTOR RELATIONS (3Q26 EARNINGS TRANSCRIPT); GLOBENEWSWIRE (GAINSIGHT); SAP INVESTOR RELATIONS

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 (30 DAYS' NOTICE PROMISED)

WATCH 01

Salesforce publishes the Headless Platform Interaction multiplier and the terms of the 'unlimited' headless licence

The number that decides whether enterprise agents keep crossing the MCP boundary or replicate the data: at 1 credit per call the toll is noise, at 20 (a standard Agentforce action) it is material for any agent with thousands of daily calls. The weekly's synthesis treats a multiplier other than one credit per call as confirming its non-list-meter pattern.

NOV 3 TO DEC 1

WATCH 02

Zendesk Employee Service AI Agents GA rollout and the first customer disclosures of resolution volumes against the 10-15 allowance

The first test of allowance-and-pause in production: if most customers exhaust the included resolutions and buy at \$0.99, the allowance is a trial; if most do not, the product is bundled. Either answer is the datapoint for every outcome meter that follows.

**EARLY TO MID-NOV;
NOV 17; DEC 1**

WATCH 03

Microsoft's UBB model rollout in Copilot, Work IQ admin controls public preview, and standalone Copilot SKUs

Three dated steps toward usage billing as Copilot's default posture; the share of tenants that enable UBB and allocate credits before mid-November decides how many users get GPT-6 Astra and Claude Fable 5.1 by default. December 1 is also the weekly's meter-cluster checkpoint (Argon reversion, Copilot credits).

NOV 16

WATCH 04

Unified Claude Enterprise turns on by default for existing organisations

The date an administrator who has set no policy gets the merged chat-plus-Cowork experience with agentic sandboxes; the share of organisations that set Efficient mode or model entitlements before the switch is the first measurement of how enterprises use the new cost controls. ZDR for Claude Code, HIPAA-ready and CMEK configurations remain excluded.

**THROUGH DEC 31,
THEN LIST**

WATCH 05

Claude Security's \$25,000 monthly credit window and the first independent comparison with Codex Security Cloud on change-stream coverage and false-positive rate

The security scorecard's new challenger is on subsidised usage until year-end; the first customer disclosure of cost at list, or an independent evaluation on the two axes that matter, decides whether the row holds. Trail of Bits' regression-hunt baseline (Agent Techniques) is the methodology to apply.

JAN 2027 (IBM
Q4); 2027
(GAINSIGHT
PRICING)
WATCH 06

IBM's Q4 disclosure on the GenAI book of business and Consulting margin; Gainsight's published ARR-indexed price

A fourth quarter without the book-of-business figure makes the omission policy; a Consulting gross margin below 27% makes the dilution structural. Gainsight's actual index and cap on the ARR-based price is the first published example of a SaaS price tied to the customer's own revenue rather than to seats or usage.

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 27

Week 43 of 2026 · October 24, 2026

WEB

brianletort.ai/industry/applications

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