THE APPLICATION LAYER
ISSUE 21 · WEEK 37 OF 2026

THE BIG READ

Seven agents shipped with no price, and the market lost one independent per-resolution comparator in the same week

KEY SIGNALS THIS PERIOD

7	6	4	4
Verticals tracked this period	Vertical packages shipped	Incumbent SaaS responses	Startup signals

AUTHOR
Brian Letort
BrianLetort.AI

PUBLISHED
September 12, 2026
Issue 21 · Weekly read

WEB
brianletort.ai/industry/applica
The Application Layer archive

THE BIG READ

The thesis this issue defends.

Independent industry analysis. Compiled from public vendor announcements, SEC filings, earnings releases, and reputable trade press. The body sets the read for the rest of the brief.

Salesforce released seven named pre-built agents on September 11 and published no pricing, packaging, or edition placement for any of them. Zscaler made an agentic security operations centre generally available worldwide on September 9 and published no rate. Automation Anywhere shipped a generally available procure-to-pay product and published no per-invoice figure. Adobe put structured contract extraction into an enterprise document tier and published no enterprise price. In a market that spent 2026 arguing about seats against consumption against outcomes, four of the week's most significant releases moved the price off the pricing page and into the contract.

The day before, Salesforce closed its acquisition of Fin. That puts the two largest per-resolution priced customer agents under one vendor, which is one fewer independent comparator rather than the last one; per-resolution pricing still exists elsewhere, it is simply no longer quoted by an independent leader. It matters more than the acquisition multiple, which was not disclosed either. If you are renewing support automation in 2027, the contractual definition of a resolution is now the single most negotiable term you have, and there is one fewer independent quote to anchor it against.

The outcome numbers that were published are almost uniformly unbaselined. Six customer percentages in the Salesforce release, a 76% resolution rate at the Fin close, and a 90% straight-through-processing ceiling from Automation Anywhere all arrive without a denominator, a measurement window, or a definition of the operative verb. That is not a reason to dismiss the products, which are real and in several cases genuinely available. It is a reason to insist that the metric you will be billed on is defined in the agreement rather than in the vendor's telemetry. The buying unit this week moved from a platform you configure to a role you hire. The measurement of whether that role did its job stayed with the vendor.

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.

6 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.

2026-09-11
PACKAGE 01
Salesforce
SUPPORT
INCUMBENT SAAS
UNKNOWN

Agentforce pre-built job agents and long-horizon runtime

Seven named role-shaped agents, six generally available immediately, with no published price for any of them

The agents cover service, internal service, commerce, outbound sales, supply chain, pipeline generation and customer experience, and the accompanying runtime adds cross-session memory, durable execution and dynamic steering so an agent can hold a goal across weeks. Only one of the seven runs on that runtime today. A buyer now acquires a role rather than a builder, which shortens deployment but shifts diligence onto verifying what the vendor means by resolved before signing anything priced on outcomes.

SOURCE: SALESFORCE NEWSROOM

2026-09-08
PACKAGE 02
Forus
HEALTHCARE
STARTUP UNKNOWN

Prescription-to-treatment navigation agents

Agents that own insurance navigation, financial assistance and pharmacy supply between decision and first dose

Each agent works as a case manager reasoning over medical history, coverage and financial situation, embedded directly in provider workflow. The commercial structure is the decision point: the service is free to providers and patients and monetised on the biopharmaceutical side, which means it can enter a health system without conventional budget approval and may bypass the purchase trigger that normally starts review. It does not remove security, privacy, integration or legal review. The concrete question for a chief information officer is narrower than who pays: which fields reach the biopharmaceutical sponsor, specifically whether protected health information, coverage determinations, treatment-start dates and financial-assistance decisions are visible to it, under what patient consent, and under what contractual restriction on secondary use. Ask that before clinicians adopt it, not after.

SOURCE: REUTERS

2026-09-09
PACKAGE 03
Automation Anywhere
FINANCE
INCUMBENT SAAS
UNKNOWN

Agentic Procure-to-Pay

Generally available procurement-lifecycle orchestration that explicitly does not replace the system of record

It coordinates vendor onboarding, requisition, purchase orders, goods receipt, invoice processing and payment tracking as a layer across existing enterprise resource planning rather than inside it, and routes exceptions to a human with context rather than acting unilaterally. That posture is the pragmatic path for a finance function with entrenched systems. The 90% straight-through-processing figure attached to the announcement belongs to the adjacent accounts-payable product, is an upper bound rather than an average, and has no published baseline.

SOURCE: AUTOMATION ANYWHERE

2026-09-09
PACKAGE 04
Zscaler
SECURITY
INCUMBENT SAAS
UNKNOWN

Agentic SOC

Globally available the day it was announced, with agents that can isolate users and sever network paths

Agents perform triage, root-cause investigation, verdict assignment and response triggering, with closed-loop remediation that can isolate compromised users, block command-and-control traffic and cut lateral movement, investigating on-platform rather than forwarding raw logs. This is the highest-blast-radius agent category an enterprise will deploy, and the launch published no detection rate, containment time, or false-positive figure, and did not state which containment actions require human approval. Ask about the approval gate and the rollback path, not the detection quality.

SOURCE: ZSCALER

2026-09-10
PACKAGE 05
OpenAI
OPERATIONS
FRONTIER LAB
UNKNOWN

Data agent in ChatGPT Work

Conversational warehouse analytics whose queries inherit the connected account's row and column permissions

It connects to a named list of warehouses and observability tools, investigates what changed, and builds shareable dashboards from conversation rather than a query language. The governance model is the whole story: administrators choose which connections exist and which roles may use them, and queries are stated to enforce table, row and column level restrictions. Permission inheritance at that granularity is the single control that decides deployability in a regulated environment, and it is currently asserted in documentation rather than demonstrated, so test it against a real restricted table.

SOURCE: OPENAI

2026-09-09
PACKAGE 06
Adobe
LEGAL
INCUMBENT SAAS
UNKNOWN

Acrobat Knowledge Base and Analyzer

Structured extraction of renewal dates and termination clauses from large contract sets, inside an already-licensed tool

Knowledge Base makes document collections queryable across formats with cloud-storage sync, and Analyzer pulls structured fields out of contract populations with source citations. This is a procurement and legal-operations capability arriving in software most organisations already pay for, restricted to the enterprise document tier whose price was not disclosed. It is a shadow-adoption question before it is a budget line, so the governance policy should be written before a business unit discovers the feature.

SOURCE: HEISE ONLINE

INCUMBENT RESPONSES

How the SaaS estate is answering.

4 moves from established SaaS vendors reacting to the agentic shift — product launches, repositioning, earnings color, partnerships. The system-of-record incumbents defending turf against systems of action.

2026-09-10
RESPONSE 01
Salesforce
FIN

Acquisition of the leading independent per-resolution customer agent closes

More than 30,000 companies and a separate proprietary model suite move in-house, and Salesforce positions the two products as an explicit two-tier portfolio rather than a migration: Fin as the fast-deploy option over third-party help desks, Agentforce as the customisable enterprise option. The market effect is that the two largest per-resolution priced agents now share a vendor, which removes an independent price anchor from every 2027 renewal negotiation.

SOURCE: SALESFORCE

2026-09-09
RESPONSE 02
Adobe
ACROBAT

Document capabilities extended into third-party assistants and browsers

Acrobat functionality is being pushed into external chat assistants, a messaging platform and two browsers, which moves the document layer to wherever the user already is rather than requiring them to open the application. For a governance lead that is a data-egress surface to enumerate, because the same capability reached through an external assistant carries a different data path than the same capability reached inside the desktop tool.

SOURCE: HEISE ONLINE

2026-09-08
RESPONSE 03
Workday
DECISION INTELLIGENCE AND FINANCE AGENTS

Strategy refocused from many small agents onto fewer agents handling large processes

An executive defined the new posture as AI systems that automate large parts of the human-resources and finance value chains, with decision intelligence generally available and audit and compliance agents still in progress for later in 2026. The same remarks disclosed that employees had built 22,000 agents internally in a matter of weeks, which is a build-volume figure with no production filter attached and the clearest available picture of what agent sprawl looks like before a registry exists.

SOURCE: CLOUD WARS

2026-09-10
RESPONSE 04
ServiceNow
AI CONTROL TOWER

Commitment to govern third-party agents alongside first-party ones

Executive remarks at an investor conference described an agentic front door supporting both the vendor's own agents and external ones, with model-to-task matching framed as token-cost control rather than capability. The strategically useful commitment is third-party agent governance, which is the position every enterprise should press each incumbent platform on before accepting a single-vendor agent estate. The accompanying claim that agents handle 90% of certain internal tasks is an unscripted remark reported secondhand with no baseline or method and should be treated as directional intent.

SOURCE: MARKETS DAILY

STARTUP SIGNALS

The insurgent vertical cohort.

4 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.

2026-09-09
SIGNAL 01
Harvey
LEGAL $550M

Raises $550 million at a $15.5 billion valuation

Legal is now the most heavily capitalised vertical AI category by a wide margin, and the company states that 80% of the largest hundred US law firms are customers, though use is undefined and no distinction is drawn between pilot and firmwide deployment. No revenue, recurring revenue, retention, or pricing figure was disclosed. An in-house legal function weighing build against buy in 2027 is competing with a nine-figure annual development budget.

SOURCE: REUTERS

2026-09-08
SIGNAL 02
Forus
HEALTHCARE $150M

Raises $150 million at a $3 billion valuation, roughly tripling in four months

Prior authorisation and benefits navigation is the highest-velocity funded healthcare AI category of the year, and this round marks it at roughly three times its valuation four months earlier. Coverage claims of all fifty states and 85% of residential postal codes are company statements with no denominator or method. The free-to-provider model is the structural point: adoption does not require conventional budget approval and may not trigger a purchase review, so the capability needs an explicit security, privacy, integration and legal review.

SOURCE: REUTERS

2026-09-08
SIGNAL 03
GenHealth.ai
HEALTHCARE $16.5M

Raises $16.5 million to run revenue-cycle agents provisioned like employees

The agents work inside electronic health records, payer portals, fax lines and phone systems for intake, eligibility, prior authorisation, billing and denials. The pattern worth separating from the funding is the deployment model: staff train an agent by demonstrating the workflow, then provision it a user account the way they would a new hire. That should trigger the same access review, least-privilege and offboarding controls as a human account, and most identity programmes do not currently name agents at all. Source quality caveat: this item rests on a single secondary outlet with no company release opened, so treat the deployment pattern as the finding and the round itself as unconfirmed.

SOURCE: MEDSPARK

2026-09-08
SIGNAL 04
Desert Ant Labs
OTHER

Launches eighteen on-device models free to 100,000 monthly active devices per software development kit

Twelve stable and six beta task-specific models run entirely on device across mobile and web runtimes, delivered through three language software development kits, with no tokens and no logins below the stated threshold. The commercially interesting member is an approximately 11.6 MB reversible redaction model covering twenty-seven languages. For any product handling regulated or personal data, a bundled on-device redactor with no per-inference bill changes the build-versus-application-programming-interface calculus for preprocessing before anything reaches a frontier model. Pricing above the threshold is unpublished.

SOURCE: DESERT ANT LABS

PRICING SHIFTS

Seat to outcome, one move at a time.

2 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.

2026-09-10
SHIFT 01
OpenAI

Usage-based → Hybrid

GPT-Live-1 prices the full-duplex voice layer at $0.05 per minute and delegates reasoning and tool calls to a separately selected, separately billed text model. A voice agent that used to sit on one token meter now sits on a per-minute meter plus a per-token meter, which allows cheap models on routine turns and expensive ones on escalations but exposes the deployment to two independently changeable rate cards. Four controls have to be written down before this ships: which turns route to which backend model, the end-to-end latency budget the voice layer will tolerate while waiting on it, the recording and retention posture for the audio stream separately from the text, and the behaviour when the backend stalls, because a full-duplex model that keeps listening while its reasoning layer hangs will hold the conversation open rather than fail.

SOURCE: OPENAI

2026-09-08
SHIFT 02
Desert Ant Labs

Usage-based → Hybrid

On-device inference billed per monthly active device rather than per inference, free to 100,000 monthly active devices per software development kit with no stated limit on runs per user and no expiry date published. Below that threshold there is no per-inference bill and no data leaving the device. Above it the rate is unpublished, so the free tier is currently the only known commercial term and should not be modelled as a permanent one.

SOURCE: DESERT ANT LABS

VERTICAL SCORECARD

Who leads each vertical.

Leader-vs-challenger by vertical, useful for procurement shortlists when matching workload to vendor cohort. As of Sep 12, 2026.

VERTICAL	LEADER	CHALLENGER	READ
Healthcare	Forus	GenHealth.ai	Both attack the administrative gap rather than the clinical decision, and both are priced away from the provider, so neither requires conventional budget approval to enter a health system. That can bypass a purchase-triggered review, not the need for security, privacy, integration and legal review. The challenger placement is provisional: GenHealth.ai is sourced to a single secondary outlet.
Sales	Salesforce Agentforce	Independent outbound agent vendors	The outbound agent is the one member of the new seven-agent portfolio still in pilot, with general availability stated for November, so the incumbent's sales claim is the least shipped part of its own release.
Support	Salesforce	Zendesk and per-resolution independents	With the Fin close, the two leading per-resolution agents share one vendor, and the contractual definition of a resolution becomes the most negotiable term in a 2027 renewal.
Operations	OpenAI ChatGPT Work data agent	Established business-intelligence vendors	A frontier lab is now shipping governed warehouse analytics directly rather than through a data platform, competing on permission inheritance rather than on visualisation.
Engineering	GitHub Copilot	Cursor	Both shipped unattended agent work this week, and the differentiator moved to controls: override-resistant managed permissions against a persistent coordinator that delegates to subagents.
Finance	Automation Anywhere	Workday	One shipped a generally available procure-to-pay layer over existing systems of record; the other refocused on fewer larger agents with its audit and compliance agents still unshipped until late 2026.
Security	Zscaler	Incumbent security-operations platforms	An agentic security operations centre went globally available with autonomous containment and no published efficacy metric, which makes the approval gate rather than the detection rate the procurement question.

VERTICAL	LEADER	CHALLENGER	READ
Legal	Harvey	Adobe Acrobat Analyzer	A $15.5 billion specialist and a contract-extraction feature inside already-licensed document software now address overlapping work, and only one of them requires a purchase decision.

ARCHITECTURE WATCH

Patterns to track.

5 cross-vendor patterns this period. Each pattern names the trend, the exemplar moves, and what it changes for procurement, cost, or vendor-risk posture.

PATTERN 01

The buying unit is a role, not a builder

`Salesforce seven named job agents` `Salesforce Fin two-tier portfolio`

The dominant packaging move of the week was to name the job rather than the toolkit: seven agents named for service, commerce, sales, supply chain and pipeline roles, plus an explicit two-tier split between a fast-deploy option and a customisable one. This shortens time to value and it relocates the diligence. When you bought a platform you evaluated configurability; when you buy a role you must evaluate the vendor's definition of the role's output, because that definition is what gets measured and, under outcome pricing, what gets billed.

SOURCE: SALESFORCE NEWSROOM

PATTERN 02

Permission inheritance is the deployability test, not the feature list

`OpenAI Data agent row and column enforcement` `Adobe enterprise document tier`

The only governance property that decides whether a conversational data agent can run in a regulated environment is whether its queries execute under the connected account's existing restrictions down to row and column level. One vendor asserted exactly that this week, in documentation, without demonstration. The practical test is cheap and specific: point the agent at a table where a named account is restricted at row level and confirm the agent cannot see what that account cannot see. Run it before rollout, not after.

SOURCE: OPENAI HELP CENTER

PATTERN 03

Orchestrate across the system of record rather than replacing it

`Automation Anywhere agentic procure-to-pay` `ServiceNow agentic front door`

Two vendors shipped or described the same architecture: an agent layer coordinating across existing enterprise resource planning, procurement and service systems without becoming the system of record, with exceptions surfaced to a human rather than resolved unilaterally. For an organisation with entrenched core systems this is the only deployable shape, because it does not require a migration to realise value. The risk it introduces is that the orchestration layer accumulates the effective business logic while the auditable record stays in systems that never see the reasoning.

SOURCE: AUTOMATION ANYWHERE

PATTERN 04

On-device preprocessing before anything reaches a frontier model

`Desert Ant Labs redaction and language-identification models`

A set of small task-specific models running entirely on device, with no per-inference bill below a stated device threshold, changes where the trust boundary sits for regulated data. Redaction, language identification and topic tagging done locally means the payload that crosses into a hosted frontier model is already minimised, which is a cheaper control than negotiating retention terms. The gap to check is that the lab publishes no minimum device specification, so the latency claims have no stated hardware floor and must be measured on the devices your users actually carry.

SOURCE: GROUND TRUTH

PATTERN 05

Adoption routed around procurement by moving who pays

`Forus free-to-provider model` `Adobe enterprise features inside licensed software`

Two different mechanisms produced the same outcome this week: capability entering an enterprise without conventional budget approval. One vendor gives the product away to providers and patients and bills an upstream industry; another ships enterprise contract extraction inside software the organisation already licenses. Either may bypass the purchase trigger that normally starts review, which means the governance question can arrive after adoption rather than before it. The countermeasure is to write the policy against the capability rather than against the budget line, because there may be no new budget line to attach it to.

SOURCE: REUTERS

WATCHLIST

On the radar next.

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

SEP 15-17
WATCH 01

Dreamforce pricing and packaging for the seven job agents

Six agents went generally available with no published price, edition placement, or statement on whether the long-horizon runtime is separately licensed; the conference is where those terms should appear.

OCT 2026
WATCH 02

Agent skills and agent optimiser general availability

Both were announced as pilot-now with October general availability, so the platform half of the agent portfolio is not buyable until then and should not be assumed in a 2026 plan.

NOV 2026
WATCH 03

Outbound sales agent general availability

It is the only one of the seven running on the long-horizon runtime and the only one still in pilot, which makes its general availability the first real test of week-scale goal persistence in production.

LATER IN 2026
WATCH 04

Reconciliation of per-resolution and credit-based pricing after the Fin close

Salesforce stated no change to the acquired commercial model, leaving two differently metered customer agents under one vendor with no published mapping between them.

2028-2029
WATCH 05

California AI auditor registry commencement

Reported statutory dates would make unregistered covered AI audit work unlawful, which would render today's informal assurance arrangements unusable as later evidence; verify the dates against the enrolled statute.

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 21

Week 37 of 2026 · September 12, 2026

WEB

brianletort.ai/industry/applica

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