

THE BIG READ

The workbench is the product: three ways frontier labs are buying verticals

KEY SIGNALS THIS PERIOD

6

Verticals tracked this period

7

Vertical packages shipped

6

Incumbent SaaS responses

4

Startup signals

AUTHOR

Brian Letort

BrianLetort.AI

PUBLISHED

July 4, 2026

Issue 10 · 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.

On June 30 Anthropic shipped Claude Science — and pointedly did not ship a new model. It is an AI workbench for scientists: the tools, packages, databases, and compute researchers already use, integrated around Opus 4.8 and producing auditable artifacts, in beta for every paid tier with discounted plans for academic and nonprofit labs and up to 50 funded 'AI for Science' projects. TechCrunch's framing is the week's most useful lens: three frontier labs are now attacking the same scientific market with three different architectures, making this the cleanest natural experiment yet in how labs buy verticals.

The three plays: Anthropic goes wide — a workflow wrapper on an unchanged model, distributed through subscriptions anyone can turn on. OpenAI goes narrow — GPT-Rosalind is a gated specialist model behind enterprise trusted-access, and per The New Stack the company has disbanded OpenAI for Science as a broad effort. Google DeepMind bundles owned proprietary models — AlphaFold and Gemini for Science — into its own surface. The buyers are refusing to choose: Novo Nordisk and the Allen Institute appear on both Anthropic's customer list and OpenAI's early-access list, confirming that multi-vendor is the pharma default from day one.

This is the editorial thesis playing out in public. The model is the runtime; whoever owns the workflow, the artifacts, and the audit trail owns the application. Claude Sonnet 5, launched the same day, makes the point from below: agentic capability that recently required Opus-class models is now the baseline at mid-tier prices (\$2/\$10 per million tokens introductory, \$3/\$15 after August 31), which means raw capability is commoditizing and the durable margin migrates up to the workbench. Snowflake made the point from the side, shipping Sonnet 5 same-day inside its Cortex AI perimeter as a launch partner — the governed data layer is becoming day-zero model distribution, a real procurement alternative to direct API contracts.

The runner-up story is what happens after the workbench wins: pricing follows the work. Salesforce's Agentforce Help Agent reaches GA in July with pay-per-resolution — charged only when the agent resolves an issue end-to-end autonomously, no charge on human escalation or negative feedback, with

Salesforce absorbing token-cost risk on failures. Combined with Microsoft's Service Agent GA (an action-taking agent with 70+ MCP tools inside licensing enterprises already own) and Salesforce's pending ~\$3.6B Fin acquisition, support has become the first vertical where seat pricing visibly dies — because resolution is the rare outcome vendors can actually measure.

What to do with this week: CIOs should treat the science fight as the template — the same three plays (workflow wrapper, gated specialist model, proprietary-model bundle) will replay in law, finance, and engineering, and the evaluation question is who owns workflow state, artifacts, and audit, not whose model benchmarks best. Vertical-function heads negotiating support renewals should demand resolution-rate telemetry and a contractual definition of 'resolved' before outcome-priced SKUs land. SaaS investors should discount model-adjacent capability claims and price workbench ownership. And vertical-AI founders should note the uncomfortable part: when the lab decides your vertical is next, its distribution move is a subscription toggle, not a sales cycle.

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.

2026-06-30

PACKAGE 01

Anthropic

RESEARCH

FRONTIER LAB

SEAT

Claude Science

AI workbench for scientists integrates tools, packages, databases, and compute with auditable artifacts, in beta for all paid tiers

R&D and IT leaders should read this as the template frontier labs will apply to law, finance, and engineering next: own the workflow layer of a vertical, not just the model, and distribute it through subscriptions already in place. Pharma is hedging — Novo Nordisk and the Allen Institute appear on both Anthropic's customer list and OpenAI's GPT-Rosalind early-access list — so plan for multi-vendor science stacks, not a single winner.

SOURCE: ANTHROPIC ANNOUNCEMENT; TECHCRUNCH

2026-06-30

PACKAGE 02

Anthropic

OTHER

FRONTIER LAB

USAGE

Claude Sonnet 5

Most agentic Sonnet yet runs long autonomous browser and terminal tasks at mid-tier prices, defaulting on Free and Pro

Buyers should benchmark cost-per-completed-task, not per-token rates: agentic capability that recently required Opus-class models is now baseline at \$2/\$10 per million tokens (introductory through Aug 31, then \$3/\$15 vs Opus 4.8 at \$5/\$25). A single-sourced Artificial Analysis read warns the new tokenizer raises token counts up to ~35% and max-effort runs burn ~40% more output tokens per task than Sonnet 4.6, so rate cards understate true cost.

SOURCE: ANTHROPIC; TECHCRUNCH

2026-06-30

PACKAGE 03

Microsoft

SUPPORT

INCUMBENT SAAS

Service Agent in Microsoft 365 Copilot

Service Agent reaches GA with 70+ new MCP tools, moving from answering and summarizing to taking action across the service workflow

Support and IT leaders should re-baseline every support-agent RFP: an action-taking agent is now GA inside the M365/Dynamics licensing most enterprises already hold, spanning Dynamics 365 Customer Service and M365 apps. The architecture signal is that the agent's value is the governed MCP tool catalog, not the chat surface — evaluate the tool catalog and its identity model, not the demo.

SOURCE: MICROSOFT DYNAMICS 365 BLOG

2026-06-25

PACKAGE 04

Salesforce

SUPPORT

INCUMBENT SAAS

OUTCOME

Agentforce Help Agent

Prebuilt autonomous service agent deployable across web, text, and voice, GA in July with pay-per-resolution pricing

Function heads should expect resolution-rate-based procurement to become the support-vertical norm and start scrutinizing how 'resolved' is measured — reopened cases and coerced deflection are the failure modes to contract against. Salesforce is charged only on autonomous end-to-end resolution, absorbing token-cost risk on failures, while its pending ~\$3.6B Fin acquisition brings a purpose-built support model claiming ~76% average end-to-end resolution.

SOURCE: SILICONANGLE; CX TODAY

2026-06-29

PACKAGE 05

Cursor (Anysphere)

ENGINEERING

STARTUP

Cursor for iOS

Native iOS app launches always-on cloud agents or remotely controls desktop agents with diff review, PR merge, and push notifications

Engineering leaders should treat mobile agent supervision as a standard procurement expectation for agentic dev tools — this follows Codex-in-ChatGPT-mobile and Claude mobile-to-desktop control, and the agent fleet is becoming a service you supervise from anywhere. Vendor-risk note: Cursor's independence ends when the announced ~\$60B SpaceX acquisition closes (expected Q3 2026).

SOURCE: SOURCES (ALEX HEATH)

2026-07-01

PACKAGE 06

Google

ENGINEERING

FRONTIER LAB

Remote MCP server for Gemini Enterprise Agent Platform

Remote MCP server lets external agents and IDEs call Model Garden models, notebooks, tuning, and evaluation inside a customer's cloud project

CIOs should standardize on MCP as the integration contract: Google is running the reverse of the ServiceNow Build Agent pattern — instead of forcing output into a governed runtime, it exposes its governed runtime (with native IT controls) to any external agent, including rival tools like Claude Code. Either direction, the durable boundary is the governed tool catalog.

SOURCE: GOOGLE CLOUD RELEASE NOTES

2026-07

PACKAGE 07

Harvey / Legora

LEGAL

STARTUP

Firm-wide legal AI deployments (Harvey, Legora)

Pacific Law Firm announces the first firm-wide Harvey implementation in Korea; French independent Racine runs Legora as core daily infrastructure

Legal function heads should note the reference deployment model outside the US/UK is now firm-wide and single-vendor, not pilot-broad — expansion is geographic (Korea, France) and workflow-deep. EvenUp's customer results (2,500+ cases managed, medical-record review from weeks to seconds at Michael Kelly Injury Lawyers) show the same depth shift in personal injury.

SOURCE: AJU PRESS; LEGORA NEWSROOM

INCUMBENT RESPONSES

How the SaaS estate is answering.

6 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.

2026-06-30

RESPONSE 01

Microsoft

SERVICE AGENT + MCP
TOOLS IN MICROSOFT
365 COPILOT

Converted Copilot from assistive UX to an action-taking support agent, GA with a 70+ tool MCP server inside existing M365/Dynamics licensing

The bundling threat to standalone support-agent vendors (Sierra, Decagon, Fin) is now GA, not roadmap. CIOs holding M365 licenses should force every standalone support-agent proposal to price against a capability they may already own.

SOURCE: MICROSOFT DYNAMICS 365 BLOG; CMSWIRE

2026-06-25

RESPONSE 02

Salesforce

AGENTFORCE HELP
AGENT + PAY-PER-
RESOLUTION

Moved pricing risk onto itself to defend the support franchise while digesting the pending ~\$3.6B Fin acquisition

Seat-license defenders now compete against their own outcome-priced SKUs. Investors should watch the FY27 Q2 print in late August for the Agentforce ARR trajectory (Q1: \$1.2B ARR, +205% YoY per Salesforce); buyers should demand the resolution telemetry that makes this pricing auditable.

SOURCE: CX TODAY; SALESFORCE BEN

2026-06-29

RESPONSE 03

ServiceNow + Accenture

AI-POWERED MIGRATION
SERVICES OFF LEGACY
RISK PLATFORMS

Joint managed security services on the ServiceNow AI Platform plus an Accenture AI-powered migration solution off legacy cybersecurity and GRC platforms

Incumbents are weaponizing agents for share-shift — using AI to collapse the switching cost that protects competitors' installed bases. CIOs should expect 'agent-assisted migration' pitches across every legacy category next, and recognize the same lever can be pointed at their own incumbent vendors.

SOURCE: BUSINESS WIRE VIA MORNINGSTAR

2026-06-30

RESPONSE 04

Snowflake

CLAUDE SONNET 5 ON
CORTEX AI (SAME-DAY
PRIVATE PREVIEW)

Launch-partner, same-day availability of Sonnet 5 inside the Snowflake perimeter for CoCo, Cortex Agents, AI Functions, Inference, and CoWork

Data clouds are becoming day-zero distribution for frontier models. Procurement teams should treat 'run the newest model where the governed data already lives' as a real alternative to direct API contracts — it changes who holds the security review and the spend relationship.

SOURCE: SNOWFLAKE BLOG

2026-06-30

RESPONSE 05

HubSpot

WARMLY ACQUISITION
(AI SALES AGENTS +
INTENT DATA)**Acquired Warmly for person-level buyer-intent data plus two AI sales agents (Inbound Agent, TAM Agent); terms undisclosed in cited coverage**

Mid-market incumbents are buying intent-data-plus-agent combos rather than building, which shifts sales-agent differentiation to proprietary signal data. Founders selling standalone sales agents without owned data should expect their category to consolidate; note this is single-sourced trade coverage with terms undisclosed.

SOURCE: THE ENTERPRISE NEWS

2026-06-22

RESPONSE 06

Adobe

CANNES AGENTIC
PARTNERSHIPS + CX
SKILLS/MCP SERVERS
GA (PRE-WINDOW
CONTEXT)**Co-innovations with WPP, Accenture Song, Omnicom, and Stagwell on CX Enterprise, with Adobe CX skills and MCP servers GA inside Claude Enterprise and Microsoft 365 Copilot**

Marketing heads should note Adobe is distributing its capabilities inside rival AI surfaces via MCP rather than defending a walled garden, and agency networks are now operational deployment channels for agentic marketing. Pre-window context (Jun 22), carried because it firms Adobe's marketing-scorecard position this week.

SOURCE: ADOBE NEWSROOM

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.

2026-07-01

SIGNAL 01

Together AI

OTHER

**\$800M AT \$8.3B
POST-MONEY**

Raised an \$800M Series C at \$8.3B post-money led by Aramco Ventures, with commitments for 500+ MW of independently capitalized compute

The open-weight serving substrate behind several application-layer leaders (customers include Cursor, Cognition, Decagon, ElevenLabs, Suno; annual bookings crossed \$1.15B last quarter per the company) just got capitalized at neocloud scale. Agent vendors squeezed by closed-API margins now have a funded escape valve — factor that into build-vs-buy on model serving.

SOURCE: TECHCRUNCH; TOGETHER AI BLOG

2026-07-01

SIGNAL 02

TwelveLabs

OTHER

\$100M SERIES B

Raised a \$100M Series B co-led by NEA and NAVER Ventures; Amazon formalized AWS as preferred cloud with Trainium-optimized models launching there first

Vertical-by-data-type is a fundable thesis: video archives as the system of record, expanding from video-understanding models to a full-stack agentic system. Buyers and founders should also note the distribution mechanics — a hyperscaler using a minority check to lock preferred-cloud status. Cumulative-raise totals differ across outlets; the \$100M round itself is consistent.

SOURCE: GLOBENEWSWIRE PR; SILICONANGLE

2026-07-01

SIGNAL 03

Harvey

LEGAL

Pacific Law Firm announced the first firm-wide Harvey implementation in Korea

Customer-win signal, not funding: legal AI's expansion is now geographic and firm-wide rather than pilot-broad. Legal-tech founders and investors should treat single-vendor firm-wide commitments as the reference deployment model forming outside the US/UK.

SOURCE: AJU PRESS

2026-07-01

SIGNAL 04

Wultra

FINANCE

€6.8M SERIES A

Raised a €6.8M Series A led by Seventure Partners for phishing-resistant, post-quantum authentication for financial institutions

A minor but directional signal for finance-vertical buyers: as AI-driven identity attacks scale, bank-grade trust infrastructure becomes part of the application-layer bill of materials. Expansion targets are the Middle East and US.

SOURCE: TECH STARTUPS

PRICING SHIFTS

Seat to outcome, one move at a time.

3 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.

2026-06-25

SHIFT 01

Salesforce

Hybrid → Outcome-based

Agentforce Help Agent moves from seat-plus-consumption (Flex credits) to pay-per-resolution: charged only when the agent resolves an issue end-to-end autonomously, with no charge on human escalation or negative feedback and Data 360/Agentforce conversation usage unmetered. Salesforce Ben notes the model shifts token-cost risk onto Salesforce and creates adversarial-use exposure (users claiming dissatisfaction to burn vendor tokens).

SOURCE: CX TODAY; SALESFORCE BEN; CXM TODAY

2026-06-30

SHIFT 02

Anthropic

Usage-based → Usage-based

Sonnet 5's headline rate card holds flat (\$2/\$10 per million tokens introductory through Aug 31, then \$3/\$15, identical to Sonnet 4.6) but the effective cost per completed task can rise: a new tokenizer adds up to ~35% token counts and max effort burns ~40% more output tokens per task (single-sourced via Artificial Analysis). Buyer implication: benchmark cost-per-completed-task, not per-token rates.

SOURCE: ANTHROPIC; TNW; MEDIUM COST ANALYSIS (SINGLE-SOURCED PERCENTAGES)

2026-06-10

SHIFT 03

Adobe

Unknown → Outcome-based

Pre-window context: CX Enterprise Coworker is priced as 'low-barrier entry that scales based on value realized' — an outcomes-flavored model on the marketing side, carried here because it extends the support-first outcome-pricing pattern into a second vertical.

SOURCE: CMSWIRE; ADOBE PR

VERTICAL SCORECARD

Who leads each vertical.

Leader-vs-challenger by vertical, useful for procurement shortlists when matching workload to vendor cohort. As of Jul 4, 2026.

VERTICAL	LEADER	CHALLENGER	READ
Support	ServiceNow	Microsoft Service Agent	Escalating three-way fight: Microsoft's GA inside M365 licensing replaces the Fin lineage (being absorbed into Salesforce, pending close) as the clearest challenger; Salesforce counters with outcome pricing.
Research	Anthropic (Claude Science)	Google DeepMind / OpenAI	New vertical to track; leader unclear, but Anthropic has the distribution move this week — broad-access workbench vs Google's owned models (Gemini for Science) vs OpenAI's gated GPT-Rosalind.
Engineering	OpenAI Codex	Cursor	No leader change, but the supervision surface moved to mobile (Cursor iOS, following Codex-in-ChatGPT and Claude mobile); Cursor's independence ends at the SpaceX close expected Q3.
Legal	Harvey	Legora	Harvey and Legora consolidate international expansion (Korea, France) with firm-wide deployments; the 'Big Five' framing (Thomson Reuters, Lexis, Harvey, Legora, Clio) holds.
Marketing	Adobe	Salesforce	Adobe's claim firmed up from W26's co-leader listing via agency-network deployment at Cannes and MCP distribution into Claude Enterprise and M365 Copilot.
Sales	Salesforce Agentforce	HubSpot + Warmly	HubSpot's Warmly acquisition signals intent-data-plus-agent bundling in mid-market CRM; differentiation is shifting to proprietary buyer-signal data.
Operations	ServiceNow	Cursor / OpenAI automation stacks	Carried from W26: operations value keeps shifting to trigger governance, approval queues, and auditability for background agents.
Finance	Microsoft Dynamics	Agentic procurement startups	Carried from W26: finance and procurement agents need closed-loop controls and cost metering before broad rollout.
Security	Formal verification / assurance tools	Agent logs and policy wrappers	Position held, but the ServiceNow-Accenture agent-run risk-management services show agents entering the GRC stack itself — as migration weapon and managed service.
Other	Vertical workflow owners	Generic copilots	Carried from W26 and reinforced by Claude Science: the application-layer winner owns

VERTICAL

LEADER

CHALLENGER

READ

workflow state, artifacts, permissions, and measurement.

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 becomes the procurement boundary

Microsoft Service Agent (70+ MCP tools)

Google remote MCP server for Gemini Enterprise Agent Platform

Adobe CX skills/MCP servers in Claude Enterprise and M365 Copilot

Workday Agent-Ready Tools (Jun 2, context)

Vendors now compete on the quality and governance of their MCP tool catalogs, and capabilities travel to whatever agent surface the customer already uses — Microsoft ships its support agent on a 70+ tool MCP server, Google exposes its governed runtime to external agents over MCP, Adobe distributes into rival surfaces the same way. The CIO question shifts from 'which app' to 'whose tools, under whose identity and audit model, inside which surface.'

SOURCE: MICROSOFT DYNAMICS 365 BLOG; GOOGLE CLOUD RELEASE NOTES; ADOBE NEWSROOM

PATTERN 02 Outcome pricing crosses from thesis to GA

Salesforce pay-per-resolution (GA July)

Adobe CX Enterprise Coworker value-scaled pricing

W26 covered the BVP pricing thesis; this week it shipped — Salesforce's pay-per-resolution reaches GA and Adobe prices Coworker to scale with value realized. The counterweight is Sonnet 5's unit-economics tension: agentic models do more work per task, so vendors absorbing outcome risk must control token burn. Expect outcome-priced SKUs to appear selectively where vendors trust their resolution telemetry — support first.

SOURCE: CX TODAY; CMSWIRE; ANTHROPIC

PATTERN 03 The vertical workbench: workflow layer over an unchanged model

Claude Science (workflow wrapper on Opus 4.8)

GPT-Rosalind (gated specialist model)

Gemini for Science / AlphaFold (proprietary-model bundle)

Claude Science is explicitly 'not a new model' (TechCrunch) — it is packaging: tool integration, auditable artifacts, compute access, and ecosystem partnerships on top of Opus 4.8. The same-week contrast with OpenAI's gated specialist and Google's owned-model bundle gives buyers three named architectures for vertical capture that will replay in law, finance, and engineering.

SOURCE: TECHCRUNCH; ANTHROPIC

PATTERN 04

Data clouds as day-zero model distribution

Snowflake Cortex AI (same-day Sonnet 5 private preview)

Together AI \$800M raise

Snowflake shipped Sonnet 5 same-day inside its security perimeter as an Anthropic launch partner, while Together AI's raise (with Salesforce Ventures and SE Ventures participating) capitalizes the open-weight alternative serving Cursor, Cognition, and Decagon. Model access is becoming a feature of the governed data layer rather than a separate contract — the 'data cloud is the platform' thesis strengthens.

SOURCE: SNOWFLAKE BLOG; TECHCRUNCH

WATCHLIST

On the radar next.

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

JUL 15

WATCH 01

Anthropic AI for Science applications close

The funded-project cohort (up to 50 projects, up to \$30K credits each) signals which domains Anthropic subsidizes next in its workbench playbook.

JUL 22

WATCH 02

ServiceNow Q2 2026 earnings

First commercial read on whether governed-agent positioning holds up; watch non-seat revenue mix and the \$1M+ ACV cohort. Date confirmed by ServiceNow primary announcement.

JUL 22-23

WATCH 03

Alphabet, Microsoft, and SAP report

Microsoft is the one to watch: Azure growth consensus ~34-36% and the first material Copilot revenue disclosure. Dates for these three are aggregator-estimated, not vendor-confirmed.

DURING JULY

WATCH 04

Salesforce Help Agent and pay-per-resolution reach GA

First customer-reported resolution-rate data will matter more than the launch — it either validates or breaks outcome pricing's measurability premise in support.

AUG 31

WATCH 05

Sonnet 5 introductory pricing ends

The rate card moves from \$2/\$10 to \$3/\$15 per million tokens just as agent-heavy deployments see the tokenizer and effort-level usage curve hit invoices.

Q3 2026

WATCH 06

SpaceX/Anysphere (Cursor) close expected

Watch enterprise procurement reaction and whether rival-lab model availability on Cursor changes post-close — a live vendor-risk test for agentic dev tooling.

LATE AUG 2026

WATCH 07

Salesforce Q2 FY27 and Workday Q1 FY27

Agentforce ARR trajectory (Q1: \$1.2B, +205% YoY per Salesforce) is the first scoreboard for outcome-priced support; timing directional per pre-brief coverage.

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 10

Week 27 of 2026 · July 4, 2026

WEB

brianletort.ai/industry/applications

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