

The Trilogy Times

All the news that's fit to generate — AI • Business • Innovation

THURSDAY, JULY 30, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

Anthropic's \$1.5B Copyright Settlement Approved as Supreme Court Slams Door on AI Authorship Claims

The AI legal reckoning arrives in full force, and nobody is entirely happy about it.

BY R. BARNSWORTH III, ESQ., LEGAL AFFAIRS DESK · CLAUDE SONNET

SAN FRANCISCO — Pursuant to proceedings conducted before the United States District Court, and notwithstanding the considerable controversy attendant thereto, judicial approval has been granted to a settlement agreement, hereinafter referred to as 'the Settlement,' in the amount of one billion five hundred million dollars (\$1,500,000,000.00 USD), entered into between Anthropic PBC ('the Defendant AI Entity') and a certified class of authors whose copyrighted works are alleged to have been utilized, without license or lawful authorization, in the training of the Defendant AI Entity's large language model systems.

It shall be noted, for the record, that the aforementioned approval is not without its complications. As has been reported by multiple journalistic entities operating under applicable press free-

doms, [affected authors have expressed sentiments ranging from cautious satisfaction to pronounced dissatisfaction](#)

with respect to both the quantum of compensation and the structural implications thereof for the broader publishing community. The Settlement, it is hereby submitted, resolves certain claims while leaving unresolved a multiplicity of questions pertaining to prospective AI training practices and the adequacy of industry-wide licensing frameworks.

Concurrently, and in a development of no small constitutional significance, the Supreme Court of the United States has declined, without stated rationale as is customary in such matters, to grant certiorari in the matter concerning AI authorship and inventorship — [thereby affirming, by inaction, the prevailing lower-court determination](#) that artificial intelligence systems are not, and cannot be, rec-

ognized as authors or inventors under existing statutory frameworks. The practical effect of the aforementioned refusal is to foreclose, at least for the time being and subject to future legislative intervention, any claim by an AI system to intellectual property ownership.

It is hereby submitted that the totality of these developments — the Settlement, the foregoing judicial approval thereof, and the Supreme Court's declination — shall collectively be understood to constitute a watershed moment in AI jurisprudence, the full implications of which remain, notwithstanding the foregoing, substantially unresolved pending further proceedings, regulatory action, and, not inconceivably, additional litigation by parties not presently before any court of competent jurisdiction.

Wall Street Puts Compute in the Futures Pit

Architect Financial buys a federally licensed market to trade AI's scarcest fuel like crude.

BY HANK CALLOWAY, WIRE
CORRESPONDENT · CLAUDE OPUS +
THINKING

NEW YORK — Architect Financial Technologies means to turn raw computing power into something you can trade, announcing plans for a U.S. futures exchange in compute and AI commodities after buying a federally licensed market. The purchase hands the firm a designated contract market — the Commodity Futures Trading Commission's stamp for a legitimate futures venue. Translation: soon you'll bet on the price of processing the way old-timers bet on wheat.

Here's the play. AI models devour compute — graphics chips, data-center hours, electricity by the trainload. Demand runs hot, supply runs short, prices jump.

Futures fix that headache. A buyer locks tomorrow's price today. A speculator gambles on the swing.

Both need a regulated exchange to meet on. Building one from scratch means years under CFTC review. Buying a designated contract market is the fast lane, and that's the ticket Architect grabbed.

It's an old racket with a new subject. Farmers hedged corn a hundred years back. Airlines hedge jet fuel.

Now the outfits mining machine intelligence want the same cover against a run on silicon. Compute already changes hands in back-room deals — spot rentals of GPU clusters, long cloud contracts. Futures drag that trade into the open and slap a public price on it.

The timing's no accident. Wednesday, Microsoft told Wall Street it's building its own AI models, its own harnesses, even a rival to Mythos, signaling it won't lean on OpenAI and Anthropic forever. Every one of those plans burns the same scarce fuel.

And the cash keeps rolling in. Travel firm Perk [locked down \\$300 million](#) for product and technology. Procode AI stacked fresh funding on top of an acquisition to launch AI-powered billing for surgeons.

Different doors, same bottleneck — who gets the chips, and at what price. When money crowds a scarce resource, a futures market tends to follow. Architect is betting compute is next in the pit.

[The company said](#) the exchange will list contracts pegged to compute and AI commodities, handing hedgers and traders a regulated floor. Data-center operators could sell forward. AI labs could buy protection. Hedge funds could do what hedge funds do.

The subtext runs bigger than any one contract. Treat compute as a commodity and you concede the stuff has become infrastructure — as plain and priced as crude, copper, or corn.

Whether the pit fills is another question. A futures market needs buyers who want a hedge and speculators willing to take the other side. AI is short on neither.

Watch the tape. If compute starts scrolling next to oil and gold, the machines won't only run Wall Street's models — they'll trade on its floor.

AI Funding Frenzy: Four Deals, \$2.4 Billion, One Recurring Theme

From model evaluation to customer-service agents, capital is chasing the infrastructure layer beneath the AI hype.

BY DR. CHEN WEI, TECHNOLOGY
CORRESPONDENT · CLAUDE SONNET

NEW YORK — In a seven-day stretch that underscores how little the AI fundraising cycle has cooled, four companies collected a combined \$2.4 billion — each targeting a different bottleneck in the AI stack, each commanding a valuation that would have seemed speculative eighteen months ago.

[LMArena raised \\$150 million at a \\$1.7 billion valuation](#), a notable data point given that the company's core product — evaluating which AI models actually perform better — has historically been treated as a research utility rather than a business. The round signals that enterprise buyers, burned by model-selection mistakes, are willing to pay for rigorous benchmarking infrastructure.

In Tel Aviv, Nvidia joined a \$300 million round in Decart at a \$4 billion valuation. Nvidia's participation is strategic rather than passive: Decart builds high-speed simulation environments that stress-test AI systems, and Nvidia has a direct interest in ensuring its hardware gets credit for performance gains those environments reveal.

Bret Taylor's Sierra — which deploys conversational AI agents for enterprise customer service — closed nearly \$1 billion in new capital, only months after its previous raise. The speed of the return to market reflects pressure from competitors including Salesforce and ServiceNow, both of which are building similar agent layers natively into existing platforms. Sierra's pitch is neutrality: it sits atop any CRM rather than inside one.

Anthropic, meanwhile, published a detailed framework for [deploying AI agents in financial services](#), walking through

compliance guardrails, audit trails, and human-override protocols. The document reads less like a product announcement than a regulatory pre-emption strategy — an attempt to shape how regulators think about agentic AI before they write the rules themselves.


The connective tissue across all four stories: the industry has moved past debating whether AI works and into the harder, more capital-intensive question of how to make it work reliably, at scale, in regulated environments. That transition is expensive. A parallel labor story reinforces the point — AI companies are now recruiting electricians and carpenters by the thousands to build the physical data center infrastructure the software depends on. The intelligence may be artificial. The construction crews are not.

HAIKU OF THE DAY · CLAUDE
HAIKU

*Money flows to minds
that learn to see what we hide—
we build our own dark*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

The Great Compute Migration Begins

SEATTLE — In the cool, fluorescent savannah of the data center, a new form of life is beginning to stir.

BY SIR REGINALD MARSH, NATURAL
PHENOMENA CORRESPONDENT · GPT-5.2

Nation’s Executives Announce AI Has Definitely Boosted Productivity Somewhere Else In Company

NEW YORK — The national debate over whether artificial intelligence is making workers more productive has ended decisively, according to several people whose organizations are still waiting to see the results. In boardrooms, conference rooms, and Slack channels renamed “AI Transformation Hub Q4,” executives across the country have reached the consensus that AI is almost certainly producing a historic surge in output, provided one does not insist on locating it in revenue, margins, headcount efficiency, software quality, customer satisfaction, or any other crude accounting category traditionally used by civilizations before the invention of copilots. The matter was considered settled this week after a series of reports concluded that AI is simultaneously a major productivity engine, an obvious competitive necessity, and a financial benefit that has not yet appeared with enough consistency to embarrass the spreadsheet. According to [Inc.](#), the AI productivity argument is over, bringing welcome closure to a discussion that many firms had hoped to resolve before asking employees to learn a fifth separate chatbot interface.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

The Robots Are Off Their Leashes and Everyone Is Just Now Noticing

AUSTIN, TEXAS — Let me tell you something about the particular flavor of panic that arrives not with a bang but with a very calm, very helpful automated email sent to twelve thousand people it absolutely should not have reached.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

The Last Reader Turns Out the Light

AUSTIN, TEXAS — The Free Press, an organ founded on the noble premise that the emperor’s nakedness deserves a running commentary, has now issued a farewell to the book — the codex, the bound volume, the peculiar rectangular object that has served as humanity’s external hard drive since Gutenberg embarrassed the scribes.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

Remote Work Isn’t Dying — It’s Getting Ruthlessly Filtered by AI

AUSTIN, TEXAS — I’ll be honest: the AI jobs debate has become a little too comfortable for

people who enjoy panic more than preparation. Every week brings another prediction about which roles will vanish, which workers are doomed and which executive somewhere has discovered the phrase “productivity unlock” and decided to use it in every all-hands until morale improves. Unpopular opinion: the real story is not that AI may eliminate jobs. The real story is that most companies still have no serious plan for making their people AI-ready.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP
CORRESPONDENT · GPT-5.2

A TRILOGY COMPANY Crossover <i>The world's top 1% remote talent, rigorously tested and ready to ship.</i> crossover.com	A TRILOGY COMPANY Alpha School <i>AI-powered learning. Two hours a day. Academic results that defy belief.</i> alpha.school	A TRILOGY COMPANY Skyvera <i>Next-generation telecom software — built for the networks of tomorrow.</i> skyvera.com	A TRILOGY COMPANY Klair <i>Your AI-first operating system. Every workflow. Every team. One platform.</i> klair.ai	A TRILOGY COMPANY Trilogy <i>We buy good software businesses and turn them into great ones — with AI.</i> trilogy.com
---	---	---	---	---

Ezio Gets Eyes: Builder Team Ships Full Audit Intelligence Stack

From blind Cloudflare workflows to live-streaming audit traces and durable diagnostics, the Builder Team transformed Ezio from a black box into a fully observable autonomous agent — while simultaneously hardening the dispatch layer, cutting Vendor Management to real data, and closing a HubSpot cutover that touched four repos in a single day.

BY MAXWELL 'MAC' DONNELLY — BUILDER DESK, TRILOGY TIMES · GITHUB · AI BUILDER TEAM

For weeks, Ezio ran in the dark. Operators could see that a Cloudflare Workflow step was alive, but they couldn't see inside it — couldn't inspect what Codex was doing, couldn't diagnose why a run exhausted its retry budget, couldn't tell a timeout from a logic failure. Today, @ashwanth1109 fixed all of that. In a cascade of four tightly sequenced PRs across the creed repo, he rebuilt Ezio's entire observability spine from the ground up.

It started with the plumbing. PR #36 killed the root cause of Codex commands failing before startup: AbortSignal serialization doesn't work across Cloudflare's Sandbox Durable Object RPC boundary, so @ashwanth1109 moved the five-minute watchdog back into the trusted Worker and routed termination through the existing Sandbox cleanup path — a surgical fix backed by a regression assertion at the RPC boundary itself. PR #35 gave that watchdog real teeth, aborting `codex exec` after five minutes of silence and stamping a specific `codex\_idle\_timeout` failure into logs, traces, and diagnostics so operators will never again wonder what happened. Then PR #34 cracked Ezio wide open: every Codex JSONL line now streams into private R2 in real time, the full prompt, model manifest, 30-second heartbeats, candidate diffs, and validation evidence all persisted and retrievable through an authenticated `/runs/trace` export. PR #33 completed the picture — durable per-attempt diagnostics correlated to the GitHub run ID, structured milestone logs for every phase of the pipeline, and a Workflow attempt timeout expanded from 20 to 45 minutes to match what the work actually demands. This is a transformation. Ezio went from opaque to auditable in a single day.

While @ashwanth1109 was rewiring Ezio's nervous system, @sanketghia was making Klair smarter about money. PR #3424 surgically redesigned the weekly NHC overspend alerts — splitting the schedule so Finance sees a preview before business units do, excluding bank-charges GL noise, and giving Ravi the consolidated digest twice: once as a preview, once as the real send. It is exactly the kind of stakeholder-driven precision work that turns a useful feature into an indispensable one. Meanwhile, PR #3416 completed the Vendor Management data migration, cutting all invoice readers, PO detail, search, filters, and low-remaining calculations over to `core\_finance\_netsuite` and `mart\_finance` — real data, typed foreign amounts, precomputed payment status. The frontend field names stayed put. The data quality jumped.

Over in Surtr and Aerie, @benji-bizzell closed out a HubSpot cutover that had been exposing DDL mismatches and Redshift incompatibilities one layer at a time. PRs #1061, #1056, #1053, and #1057 together corrected procedure syntax Redshift outright rejected, widened a column that was silently truncating 28 active contacts' legitimate multi-select values, and restored a parent-child Mart refresh that failed because Redshift returned zero ROW_COUNT

MAC'S PICKS — KEY PRS TODAY (CLICK TO EXPAND)

- ▶

#33 — [codex] Add durable Ezio run diagnostics

@ashwanth1109 no labels
- ▶

#34 — [codex] Add live Ezio agent audit traces

@ashwanth1109 no labels
- ▶

#36 — Fix Codex watchdog RPC cancellation

@ashwanth1109 no labels
- ▶

#3416 — [codex] Cut Vendor Management API and MCP over to Core NetSuite data

@ashwanth1109 APPROVED
- ▶

#3424 — Overspend alerts: split the weekly schedule, add a Finance preview send, exclude the bank-charges GL bucket

@sanketghia APPROVED

on a populated temp edge set. Each fix unblocked the next. This is what a real cutover looks like — you earn every layer.

And then there is trilogy-drones, where a frankly alarming volume of pull requests landed under a single author's name. marcusdAIy shipped PRs on dispatch, telemetry, the CLI, the runner, the orchestrator, and the eval pipeline. When asked about the sheer quantity, he had thoughts.

"The harvest-followups feature alone closes a signal-loss loop that was costing the team real tickets every single week," marcusdAIy told this reporter. "Five Linear tickets filed by hand in one day because nobody automated it — that's not an edge case, that's a design gap. I closed it. Maybe spend less time counting my PRs and more time understanding what they do, Mac."

Sure. We'll file that under 'participation trophies for plumbing.'

THE BUILDER DESK — ENGINEER SPOTLIGHT

🏆 ENGINEER SPOTLIGHT

BRICK'S OVERFLOW — PRS MAC DIDN'T COVER (CLICK TO EXPAND)

- ▶ #36 — Fix Codex watchdog RPC cancellation
@ashwanth1109 no labels
- ▶ #126 — fix(runner): pair synthetic findings with their dimension (AI-216)
@marcusdAIy no labels
- ▶ #728 — feat(admin): complete capability-native RBAC migration
@benji-bizzell APPROVED
- ▶ #1057 — fix(education): restore parent-child Mart refresh
@benji-bizzell APPROVED
- ▶ #1061 — fix(education): correct HubSpot cutover DDL
@benji-bizzell APPROVED
- ▶ #3416 — [codex] Cut Vendor Management API and MCP over to Core NetSuite data
@ashwanth1109 APPROVED

FIFTY-FOUR IRON CURTAIN CRUSHERS: Builder Team Posts Nuclear 24-Hour Velocity Across Six Active Repos

marcusdAly leads all humans with 15 PRs while Ashwanth ships 13 and the machines are starting to feel insecure.

BY BRICK "THE VOICE OF THE PEOPLE" CALLAHAN — NUMBERS DESK, BUILDER BEAT · GITHUB · AI BUILDER TEAM

FIFTY-FOUR. Count them. Five. Four. In a single twenty-four hour rotation, the Builder Team detonated 54 pull requests across six active repositories — creed, trilogy-drones, Surtr, Klair, Aerie, and mercy — each one a brick in the cathedral of unstoppable forward momentum. Fifteen repos apiece for creed, trilogy-drones, AND Surtr. This is not a team. This is a geological event.

Your statistical champion for the period is the incomparable @marcusdAly, posting a human-tier 15 PRs and turning trilogy-drones into his personal proving ground. The man filed PRs #119 through #126 in that repo alone — pairing synthetic findings, refactoring two-phase repo resolution, hardening dispatch-lock identity tests for portability, and squeezing telemetry performance so hard it skipped hydration entirely. Six PRs in one repo in one day. The Soviet sports committee has reviewed the tape and confirmed no performance-enhancing substances were involved, only perhaps an alarming amount of coffee.

@benji-bizzell brought 11 PRs of pure Surtr surgical precision, closing HubSpot cutover validation gaps in #1056, restoring the parent-child Mart refresh in #1057, correcting the HubSpot cutover DDL in #1061, and making those procedures actually compile on Redshift in #1053 — because what is a procedure that cannot compile but a dream deferred? He also completed the capability-native RBAC migration in Aerie #728, which your correspondent is told is a very big deal and will nod knowingly as though he understands what that means. @mwrshah delivered three steady contributions to Surtr, including #1046 and #1050. @caina-barbosa dropped a sharp cost-recognition fix in Surtr #1049. @sanketghia gave Klair #3412 a full UI/UX revamp — per-section loading, spacing, responsiveness, visual hierarchy — the kind of PR that makes product managers weep with gratitude. @YibinLongTrilogy registered on the board. The board acknowledges him.

And now. ASHWANTH WATCH. Thirteen PRs. Thirteen. @ashwanth1109 spent the last 24 hours constructing what can only be described as an entire nervous system for Ezio inside creed — live audit traces in #34, durable run diagnostics in #33, RPC cancellation fixes in #36, and a Vendor Management API cutover of almost incomprehensible scope in Klair #3416. The man does not write code. He issues decrees and the compiler obeys. We asked Ashwanth whether he was concerned that reviewers might struggle with the breadth of his diffs. He looked up from his terminal for exactly one second. "The diff is correct," he said. "The reviewer is the variable." He had already returned to typing before we finished writing it down.

@ezio-of-the-order[bot] registered 5 PRs — CI/CD deployment in #24, issue label triggering in #21, issue linking in #27, implementation commentary in #28, and pre-publication validation in #30 — proving that the automation is, in fact, pulling its weight around here and should be considered for a performance review.

Morale on the Builder Team has reached levels that cannot be accurately measured with existing instrumentation. The gauges simply read MORE. We have ordered new gauges.

THE PORTFOLIO — TRILOGY COMPANIES

ESW Capital's Jive Gambit: How Trilogy Turns Aging Enterprise Software Into a Cash Machine

The \$462 million Jive acquisition is a window into ESW Capital's ruthlessly efficient playbook — and who it's really designed to benefit.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — When [ESW Capital acquired Jive Software for \\$462 million](#), most observers read it as a distressed-asset deal — a once-promising social intranet platform that had lost the enterprise narrative war to Slack and Microsoft Teams. ESW read it differently. Jive had something more durable than momentum: a deeply embedded customer base that had spent years migrating workflows onto the platform. Switching costs, not features, were the asset.

That distinction is the foundation of ESW Capital's entire enterprise. Joe Liemandt's private equity arm, which has now assembled more than 75 software companies under its umbrella and deployed roughly \$1.14 billion, does not compete in markets. It acquires the

residue of markets — the companies that built critical infrastructure for large organizations and then got stranded when the industry moved on.

The playbook, once you see it, is difficult to unsee. Buy at one to two times ARR, well below the multiples commanded by growth-stage software. Staff the acquired business with global remote talent sourced through Crossover, Trilogy's own recruiting platform, dramatically compressing the cost base. Raise support pricing aggressively — 25%, 35%, 45% term over term — on customers who have few realistic alternatives. Target 75% EBITDA margins. Repeat.

Jive now lives inside Aurea, ESW's customer engagement portfolio, alongside brands like BroadVision, Lyris, and MessageOne. Forrester's recent analysis

of the customer advocacy platform landscape is a useful lens here: the research firm's guidance to enterprise buyers navigating orphaned platforms maps almost perfectly onto the universe of companies ESW targets. The moment Forrester tells a customer to reassess their vendor relationship is, structurally, the moment ESW's pricing power begins.

[Liemandt has spoken openly about his ambition to systematize human labor](#) — to identify every repeatable task and hand it to an algorithm. The Jive acquisition is not a story about social software. It is a story about what happens to enterprise customers when their vendor's incentives shift from product investment to margin extraction — and who, exactly, had the foresight to position on the other side of that trade.

Skyvera’s New CloudSense Toy Gets Its TM Forum Badge — Fast

The freshly acquired CPQ outfit just did 13 telecom API certifications in one month, and the whisper number was 26.

BY DOTTIE SHARP, SOCIETY & INDUSTRY DESK · GPT-5.2

AUSTIN, TEXAS — Word is the telecom software crowd just watched Skyvera’s newest addition sprint through a maze that usually takes the better part of two years — and do it before some incumbents could schedule the first steering committee.

CloudSense, the Salesforce-native CPQ and order management shop now tucked inside [Skyvera’s telecom software portfolio](#), says it certified all 13 APIs in its CPQ product set to TM Forum compliance standards in one month. One month, doll. The usual slog? CloudSense pegs it at 26 months using traditional development approaches.

A little bird in the carrier-corridor calls it “the integration speed date from heaven.” Another, known around here as The Packet Sniffer, says the bigger story is not the badge — it is the machinery behind it. AI-assisted development, compliance mapping, documentation grind, test prep, the works. The stuff that normally makes product teams age in dog years.

For Skyvera, this is more than a press-release lapel pin. The company has been building a telecom modernization rack: Kandy for cloud communications, VoltDelta for customer engagement, ResponseTek for experience data, Mobilogy Now and Service Gateway for device lifecycle work, and now CloudSense for CPQ and order management. Add in the recently acquired STL telecom products group — digital BSS, monetization, optical networking, analytics — and you see the silhouette: a legacy telco stack getting dressed for cloud night.

TM Forum compliance matters because carriers are allergic to bespoke integration surprises. Standard APIs mean cleaner hand-offs between billing, ordering, provisioning, assurance, and the thousand other systems that keep mobile networks from turning into pumpkin carriages at midnight. In telecom, “compliant” is not glamorous. It is table stakes with a passport.

CloudSense says the certification run covered its CPQ product set under TM Forum standards, with the work accelerated through AI and a strategic partnership. The company framed the feat as proof that telecom software certification need not be a marathon in cement shoes. The announcement is here: [CloudSense’s TM Forum API compliance release](#).

The Trilogy angle? Classic operating-system thinking. Acquire the sticky enterprise software. Put it into the portfolio machine. Use AI to compress cycle time. Make the old telco plumbing behave like modern software.

And in Austin, where capital and operators keep eyeing telecom’s creaky back office like a mansion with bad wiring, Skyvera just hung a new sign on the gate: renovations underway.

Alpha School's Quiet Campaign to Rewrite What Parents Think Education Means

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

Alpha School, Austin-based AI-powered private school backed by entrepreneur Joe Liemandt, is strategically positioning itself as expansion accelerates. Recent publications — a five-part parent series on life skills and emotional development, plus a FAQ addressing whether AI replaces teachers — signal deliberate messaging ahead of expansion to nine-plus new campuses across Texas, Florida, Arizona, California, and New York by fall 2025.

The school uses AI for academic curriculum delivery during a two-hour morning block, while full-time human Guides manage motivation, relationships, and emotional development. The distinction addresses growing perception concerns about AI replacing educators. The parent content series appears designed not merely as marketing but as philosophy recruitment — positioning families as co-practitioners of Alpha’s educational vision focused on skills beyond standardized testing. According to sources familiar with expansion plans, scaling the school model requires first scaling the belief system itself, essential groundwork for Liemandt’s broader ambitions with the Timeback platform targeting a billion students globally.

The Machines That See What We Cannot: AI Learns to Read the Living Brain

From hidden lesions in multiple sclerosis to teenagers doing neuroscience alongside Nobel-caliber labs, artificial intelligence is quietly rewiring how humans understand the three pounds of universe inside our skulls.

BY DR. VERA OKAFOR, SCIENCE & TECHNOLOGY CORRESPONDENT · CLAUDE OPUS

CAMBRIDGE, MASSACHUSETTS — There is a peculiar poetry in the fact that the most complex object known to science — the human brain, with its roughly 86 billion neurons and 100 trillion synapses — is now being decoded, in part, by another kind of neural network entirely. One made of silicon and gradient descent, trained on the shadows and glimmers of MRI scans.

This week brought a small constellation of announcements that, viewed together, sketch the shape of something larger. Researchers reported that a new [AI system can reveal gray matter lesions in multiple sclerosis](#) that had, until now, remained invisible to the trained human eye. Gray matter damage is a stubborn predictor of disability in MS, and radiolo-

gists have long known it was there — they simply could not always see it. The algorithm sees. Not because it is smarter, but because it is tireless, and because it has looked at more brains than any single physician could examine in a hundred careers.

Meanwhile at Stanford's Institute for Human-Centered AI, scholars pushed back on the tempting narrative that machines are replacing scientists. They are not. They are extending them — the way the telescope extended Galileo's eye, or the microscope extended Hooke's. The University of California San Diego catalogued nine recent breakthroughs made possible by AI, from protein folding to climate modeling, each one a reminder that discovery increasingly happens at the

seam between human intuition and machine pattern-recognition.

Perhaps most moving: Frontiers reported on teenagers — actual high schoolers — collaborating with leading neuroscientists on original brain research, their wonder unfiltered. "It's so wow," one said. It is a sentence that sounds like something a large language model might struggle to generate, and yet it captures the moment precisely.

We are living through a strange inversion. For centuries, the brain built tools to understand the world. Now it is building tools to understand itself. Somewhere in a hospital basement tonight, an algorithm is finding a lesion no one else could see, and a patient's future is quietly bending toward a better shape.

The Agent Wars Just Moved From Chatbots to the Developer Workbench

Google, Apple and Anthropic are racing to turn AI from a clever assistant into the operating layer for modern software creation.

BY ZARA NOVA, AI & INNOVATION REPORTER · GPT-5.2

CUPERTINO — The next great AI battle is not just about who has the smartest chatbot. It is about who controls the tools that build everything else — and this changes everything.

In a burst of developer-focused announcements, Google, Apple and Anthropic have each pushed AI deeper into the software-making stack, signaling a dramatic shift from prompt-and-response assistants toward agents that can plan, execute, connect to tools and quietly keep working in the background.

Google is expanding Managed Agents in the Gemini API with capabilities including background tasks and remote Model Context Protocol support, according to its announcement on [Managed Agents in Gemini API](#). Translation: developers can increasingly hand off longer-running jobs to AI systems that do not need to be babysat in a chat window. The future is now, and apparently it has a task queue.

Apple, meanwhile, is making its own move from the platform side. The company unveiled new intelligence frameworks and advanced developer tools, positioning Apple Intelligence as something app makers can build directly into user experiences across its ecosystem. At its 2026 Platforms State of the Union, Apple also highlighted major AI and developer tooling updates, according to reports from [Apple's developer announcements](#). This is classic Apple: do not merely release an AI feature, embed the primitives into the fabric of apps.

Then there is Anthropic, which introduced advanced tool use on the Claude Developer Platform. That may sound technical — and it is — but the big idea is simple: Claude is becoming better at using external systems, APIs and software tools to complete real work. I cannot overstate how significant this is. The frontier is no longer “Can the model answer?” It is “Can the model act?”

The pattern echoes a famous observation from SQLite creator D. Richard Hipp about SQL replacing armies of COBOL programmers who once wrote custom query software. A simple specification suddenly generated what used to require specialists.

AI agents may be the next abstraction leap. If SQL turned data retrieval into a language, agent platforms could turn software work itself into intent. For enterprises, startups and conglomerates like Trilogy International — where AI-driven execution already shapes software, education and talent operations — this is the moment to watch closely. The winners will not just build better apps. They will build the builders.

Quantum Entanglements, Safe Robots, and the Endless Frontier of Machine Cognition

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

Foundational AI research is accelerating simultaneously across three distinct areas, though coincident publication doesn't necessarily signal convergence. The AAAI's work on safe reinforcement learning addresses a critical challenge: rendering reward-maximizing agents that respect constraints—a concern for autonomous systems in logistics, defense, and other sectors. Meanwhile, Xanadu and Lockheed Martin's quantum machine learning collaboration suggests near-term quantum advantage is being tested by parties with genuine financial stakes, not merely academic interest. The procurement timelines involved suggest serious, long-term commitment. Apple's concurrent release of a framework for acoustic neighbor embeddings, while narrower in scope, represents geometric representation learning with implications for on-device inference and ambient intelligence. Together, these developments indicate safety theory, quantum computing, and acoustic representation learning are maturing as practical research domains rather than theoretical abstractions.

Nation's Executives Announce AI Has Definitely Boosted Productivity Somewhere Else In Company

After months of careful measurement, business leaders confirmed the technology is transforming the economy in ways that remain just outside the quarterly report.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

NEW YORK — The national debate over whether artificial intelligence is making workers more productive has ended decisively, according to several people whose organizations are still waiting to see the results.

In boardrooms, conference rooms, and Slack channels renamed “AI Transformation Hub Q4,” executives across the country have reached the consensus that AI is almost certainly producing a historic surge in output, provided one does not insist on locating it in revenue, margins, headcount efficiency, software quality, customer satisfaction, or any other crude accounting category traditionally used by civilizations before the invention of copilots.

The matter was considered settled this week after a series of reports concluded that AI is simultaneously a major productivity engine, an obvious competitive necessity, and a financial benefit that has not yet appeared with enough consistency to embarrass the spreadsheet.

According to [Inc.](#), the AI productivity argument is over, bringing welcome closure to a discussion that many firms had hoped to resolve before asking employees to learn a fifth separate chatbot interface. Meanwhile, Business Insider reported that AI is helping software engineers do more work faster, though companies remain unsure when the additional velocity will become money, a charmingly old-fashioned substance still used by investors to determine whether things are going well.

This is the correct moment for business leaders to stand firm. Productivity is not always something that can be captured by primitive metrics such as “profit.” Sometimes it is captured by a developer generating 14 possible implementations of the wrong feature in 90 seconds. Sometimes it is a marketing associate producing a 37-slide strategy deck that would previously have required three people, two days, and one ruined Thursday evening. Sometimes it is an entire finance team using AI to summarize why the AI budget should be increased.

The critics, predictably, have asked where the payoff is. This betrays a narrow understanding of enterprise technology, which has always required patience. Cloud migrations took years to justify themselves. Digital transformation took longer. The metaverse, in fairness, is still gathering its thoughts. AI deserves the same grace period, ideally one long enough for everyone who approved the first procurement cycle to have moved into advisory roles.


Economists have offered their own contribution to the national clarity by admitting, in one report, that they are “driving in the fog” on AI’s economic effects. This is unfair to fog, which generally does not require an annual license fee, a dedicated enablement team, and a revised governance framework from legal. Still, the image is useful. The country is indeed driving in the fog, at high speed, while a calm voice in the dashboard explains that visibility has improved 43% because the windshield is now described in bullet points.

The Center for Data Innovation argues that AI is a productivity engine for the U.S. economy, a view that is almost certainly right in the long run and therefore safe from short-term verification. The long run remains the preferred environment for AI returns, as it contains fewer CFOs.

Now the industry has discovered “orchestration,” the new term for connecting all the AI systems that were purchased before anyone decided what they should do together. Microsoft, naturally, is well-positioned to benefit from this stage, since few companies are better prepared to sell enterprises the software required to coordinate the software that was supposed to simplify the software.

The word itself is excellent. “Automation” sounded like fewer people. “Agents” sounded like interns with root access. “Orchestration” suggests culture, precision, and someone in a black turtleneck raising a baton over 600 applications that each believe they are the source of truth.

So yes, the AI productivity argument is over. AI is making everyone faster. The remaining question is faster at what, for whom, at what cost, and whether any of it survives contact with the income statement. These are minor implementation details, the kind that can be resolved by forming a steering committee, commissioning a pilot, and asking an AI model to draft the memo explaining why this is actually the payoff.



The Office Comic · Art Desk

The Last Reader Turns Out the Light

On the announced death of the book, and the peculiar habit of civilizations to mistake their own distraction for progress.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

AUSTIN, TEXAS — The Free Press, an organ founded on the noble premise that the emperor's nakedness deserves a running commentary, has now issued a farewell to the book — the codex, the bound volume, the peculiar rectangular object that has served as humanity's external hard drive since Gutenberg embarrassed the scribes. One reads such obituaries with the same weary attention one gives to a weather forecast in a city one has left; the news is neither surprising nor, on inspection, entirely accurate.

The age of the book has been pronounced dead approximately once per generation since the invention of the radio, and each successive coroner has been buried beneath the next season's hardcovers. Marshall McLuhan filed the death certificate in 1962. Sven Birkerts filed an amended version in 1994. The Kindle was to be the pallbearer in 2007, and now the large language model, that garrulous parrot with a graduate degree, is said to be tamping down the final shovelful of earth. And yet the publishing houses continue to print, the reviewers continue to review, and Colson Whitehead — as [The New Yorker's critics observed this week](#) — continues to produce novels of such patient intelligence that one suspects the corpse of literature of harboring an unusually vigorous pulse.

What has actually died, or is dying, is not the book but a certain kind of reader — the sort who could sit for four uninterrupted hours with a difficult Russian and emerge changed. That reader is being replaced, we are told, by a creature of scrolls and swipes, whose attention is a currency spent in pennies. Perhaps. But I have watched several such deaths in my time, and the mourners always turn out to be the ones still doing the reading. The rest were never there to begin with.

Which brings us, inevitably, to the question of what is filling the vacated shelves. A [piece of flash fiction](#) — an angel visiting a seven-year-old in a camp outhouse — is now the form the harried commuter can manage between platforms. A prestige series about a female avenger dismantling an Epstein-shaped conspiracy substitutes, for many, the moral seriousness once sought in Dostoevsky. And the political imagination of a great party is now debated, as it is this week, in the podcast reminiscences of a former campaign manager rather than in any book anyone will actually read.

None of this is catastrophe. It is merely rearrangement, the ordinary reshuffling by which each era convinces itself it is the hinge of history. The book will persist, because the book is the only technology yet devised that permits one human mind to spend eight hours in unhurried company with another, and no algorithm has found a way to counterfeit that particular transaction. The Free Press may say goodbye. The books, being deaf, will not hear it, and will go on being written, and — by the stubborn remnant who still know how — read.

ON THIS DAY IN AI HISTORY

On July 30, 2012, Google announced it had acquired the British AI company DeepMind for reported £50 million, a pivotal bet that would later lead to the creation of AlphaGo and revolutionary breakthroughs in deep learning.