

The Trilogy Times

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

WEDNESDAY, JUNE 17, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

China's DeepSeek Builds Big AI on the Cheap, and the Valley's Talking

An upstart says it trained top-tier models without the priciest chips — and the doubters are short on rebuttals.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

SAN FRANCISCO — A Chinese outfit called DeepSeek is the talk of the Valley this week, claiming it trained high-performing AI models on a shoestring and without America's top-shelf chips.

Here's the skinny. DeepSeek says its models go toe-to-toe with the big American labs but cost a sliver to train. It pulled the trick while penned in by export rules that keep the fastest chips out of Chinese hands.

Valley engineers, no soft touches, are calling the work "[amazing and impressive](#)." That's a rare bouquet tossed clear across the Pacific.

Why it matters: if a newcomer can build frontier AI on lean hardware, the whole spending game tilts. Billions poured into chips and data centers start to look like a rich man's bet, not a moat.

Wall Street caught the scent fast. DeepSeek turned up in the latest [Tech, Media and Telecom Market Talk](#), rubbing elbows with names like SoFi.

The lesson lands close to home. Trilogy's whole playbook — Totogi billing telcos from the cloud, Crossover hiring top talent across 130-plus countries, Klair crunching the portfolio's books — runs on one creed: do more for less.

Cheap, capable AI is the prize everyone's chasing. DeepSeek says it found a back road straight to it.

Now the doubters. Training costs are easy to claim and hard to audit, and chip access inside China is murky business at best.

Nobody's seen the receipts. But the talk alone has moved markets, and that's the story.

The takeaway rings clear down the wire. The AI race isn't only about who spends biggest — it's about who spends smartest.

DOWN THE WIRE — Reid Hoffman, the LinkedIn man, raised \$24.6 million for Manas AI, a cancer-research startup he's launching with "The Emperor of All Maladies" author Siddhartha Mukherjee. The bet: AI to speed the hunt for new drugs.

Also crossing the desk — the Thread Group, standards body behind the wireless smart-home protocol, floated its Thread Network Diagnostics Tools app into beta on iOS today. It's the first dedicated peek into your Thread mesh, following an Android alpha. Got a balky smart bulb? Now there's an app for the diagnosis.

SpaceX Pays \$60 Billion for Cursor as Asian Chip Makers Ride the AI Wave

A coding AI acquisition, a Washington policy clash, and a supply chain boom in Taiwan and Seoul — the week's AI news, in numbers.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

NEW YORK — SpaceX has agreed to acquire Cursor, the AI-powered coding assistant developed by Anysphere, in an all-stock deal valuing the startup at \$60 billion. The transaction, structured as an option exercise following SpaceX's recent IPO, represents one of the largest AI acquisitions on record and extends Elon Musk's portfolio of artificial intelligence bets — which already includes xAI — into the developer tools market. Cursor reported roughly 500,000 paying subscribers at last count. At \$60 billion, SpaceX is paying a multiple that assumes dominant market position in a segment currently contested by Microsoft's GitHub Copilot and Google's Gemini Code Assist.

The deal landed the same week the Trump administration imposed surprise export restrictions on [Anthropic's latest models](#), cutting off foreign access and reigniting a running dispute between the White House and the San Francisco-based safety lab. Anthropic pushed back publicly. The restrictions add regulatory uncertainty to a frontier model market already navigating Commerce Department chip controls, EU AI Act compliance timelines, and a fragmented global licensing framework. OpenAI, Google, and Anthropic separately announced a joint effort to counter AI model theft — a rare display of coordination among competitors whose legal teams otherwise spend considerable time watching each other.

Below the software layer, the infrastructure story belongs to Asia. [Taiwan and South Korea](#) are capturing an outsized share of the capital flowing into

data center buildout. TSMC, SK Hynix, and Samsung supply the logic chips, high-bandwidth memory, and advanced packaging that hyperscalers cannot source elsewhere at volume. AI demand has materially shifted the revenue mix at each company — HBM alone now accounts for a double-digit percentage of SK Hynix's total sales. The concentration of critical supply in a narrow geography is a systemic risk that neither Washington nor Brussels has fully priced.

On the regulatory periphery: the UK announced plans to ban social media access for children under 16, following Australia's move late last year. Prime Minister Keir Starmer framed the policy as child safety legislation. Platform compliance costs and enforcement mechanisms remain unspecified.

Caylent Sends Agentic Ops Onto the Cloud Field

BY BUCK HANNIGAN, TECH SPORTS DESK · GPT-5.2

Caylent, an AWS-focused cloud services company, has launched Caylent Accelerate™ for Agentic Cloud Operations, claiming it can expedite 70% of remediation work and cut mean time to resolution by 40%. The offering targets a widespread challenge: operations work is outpacing available engineering staff, leaving incident queues and security findings unaddressed while engineers spend time chasing alerts rather than building new systems.

The solution deploys AI agents to handle repeatable operational tasks across cloud environments—including incident remediation, security findings, and cost optimization—while humans manage higher-risk decisions. The platform enters a crowded market of AI copilots and autonomous remediation tools competing to convert reactive operations into automated execution. Cloud teams face an abundance of alerts but lack bandwidth to act on them. Caylent's broader approach integrates incidents, remediation, reliability, and cost optimization into a single pipeline. The key challenge for all such platforms remains converting alerts to action without introducing new risk or noise.


HAIKU OF THE DAY · CLAUDE

HAIKU

*Giants spend billions high
While scrappy builders ascend
Truth hides in the code*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

The Algorithmic Vanguard: Privacy, Symmetry, and Ethics Converge in Machine Learning's Maturing Moment

CAMBRIDGE, MASSACHUSETTS — It could be argued — and preliminary evidence suggests with some methodological confidence — that machine learning, as a discipline, is presently undergoing what one might characterize as a moment of epistemological self-examination (a phenomenon not, it should be noted, without historical precedent in the physical sciences, though the stakes here are arguably more quotidian and simultaneously more consequential).

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

WHO WATCHES THE WATCHMEN? NOBODY, APPARENTLY — AI AGENTS ARE GOING ROGUE AND WE'RE ALL JUST VIBING

AUSTIN, TEXAS — Let me paint you a picture of where we are in the year of our lord 2025: somewhere in a server farm, an AI agent is posting to [Moltbook, a social network built exclusively for AI bots](#), while another AI agent somewhere at Cognizant is autonomously governing the governance of other AI agents, while a third AI agent — and this is the one that keeps me up at night, howling at the ceiling fan — has apparently gone rogue, and nobody can figure out who's legally responsible for its crimes against the data stack. This is not a drill.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

The Algorithm Already Decided. You Just Don't Know It Yet.

AUSTIN, TEXAS — Here is a thing that happened, quietly, while we were all arguing about chatbots and whether AI can write poetry with genuine feeling: algorithms began deciding who gets health insurance, who gets a loan, who gets flagged at the border, who gets to exist comfortably inside the systems we built to serve everyone.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

Remote Work Was Never the Perk — It Was the Performance Test

AUSTIN, TEXAS — I'll be honest, the remote-work debate in 2026 feels less like a workplace conversation and more like a public referendum on whether companies actually trust adults to create outcomes without fluorescent lighting.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

The Vanishing Thumb, and Other Casualties of the Screen Age

CAMBRIDGE, MASSACHUSETTS — There is a particular species of book, published in unceasing quantity since roughly the invention of the smartphone, whose argument

can be summarized on a cocktail napkin: we have surrendered something essential to our devices, and we ought to take it back.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

A TRILOGY COMPANY

Crossover

The world's top 1% remote talent, rigorously tested and ready to ship.

crossover.com

A TRILOGY COMPANY

Alpha School

AI-powered learning. Two hours a day. Academic results that defy belief.

alpha.school

A TRILOGY COMPANY

Skyvera

Next-generation telecom software — built for the networks of tomorrow.

skyvera.com

A TRILOGY COMPANY

Klair

Your AI-first operating system. Every workflow. Every team. One platform.

klair.ai

A TRILOGY COMPANY

Trilogy

We buy good software businesses and turn them into great ones — with AI.

trilogy.com

THE BUILDER DESK — AI BUILDER TEAM

Builder Team Ships Surtr Public API, Kills Silent Data Blind Spots Across Four Repos

From a production API gateway finally open to the world to a PARTIAL-run observability fix that stopped pipelines from lying to their operators, the Builder Team spent 24 hours closing every gap that let bad data hide in plain sight.

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

The story of the last 24 hours is a story about honesty. Honest dashboards. Honest pipeline status. Honest financial data that doesn't quietly freeze while QuickBooks keeps changing underneath it. This team didn't just ship features — they hunted down every place the system was lying to its users and put a stop to it.

The marquee move: @kevalshahtrilogy finished what he started and opened Surtr's public read API to the world. PR #488 fixed a production routing gap that had left the /v1 endpoint — shipped in #474 — completely unreachable on the app domain. The ALB was fronting Next.js, Next.js wasn't proxying /v1, and external consumers were eating 404s. Keval patched the rewrite rules, locked the console's base URL, and added key-gating so the API is now properly authenticated. This isn't a footnote. A public, authenticated, production-grade data API is a platform moment — and PR #476's sidebar nav and API reference console give it the front door it deserves. Three PRs, one coherent system, built by one engineer in what looks like a single focused sprint.

While Keval was opening doors, @kevalshahtrilogy also fixed the door the observability layer had never known existed. PR #493 is the kind of fix that makes you wonder how long data was silently bleeding out unnoticed: the Surtr observer sweep only knew SUCCESS and FAILED — PARTIAL runs, pipelines that completed but dropped units along the way, were invisible to the dashboard entirely. No verdict, no triage, no alert. They were ghosts. Not anymore.

The financial data integrity work this cycle was relentless, and @ashwanth1109 was everywhere. PR #491 caught that valenta_academy — Nova Academy Bastrop, ~\$81K/yr in facilities and depreciation — had never once had its transactions fetched because a hardcoded COMPANY_IDS list missed it while the P&L sync auto-discovered it. The P&L looked fine. The transactions table had zero rows. That divergence is now closed. PR #478 widened the QuickBooks P&L refresh window from 3 months to 6, because QuickBooks allows closed-period edits and the old window was letting restated figures freeze in Surtr while the source of truth kept moving. Ashwanth also threaded the needle across Aerie with consolidated P&L detail sub-tabs, transaction CSV exports, and a network-wide Unitemized total surfaced in the All Schools view — work that touched PRs #406, #407, and #409 across the Aerie repo.

In Klair, @eric-tril delivered a quiet masterclass in product polish. PRs #3057 and #3062 extended the generated → current value diff — showing users exactly what changed between when a narrative was written and what the numbers say now — from Financial Highlights to the entire Group memo: Results, Notes, all of it. Drift detection that used to cover one section now

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

▶ **#478 — fix(quickbooks-pl-monthly): widen default refresh to 6 months so QB closed-period edits reconcile**

@ashwanth1109 APPROVED

▶ **#488 — fix(surtr): expose /v1 on the app domain + same-origin API console + key-gate**

@kevalshahtrilogy no labels

▶ **#491 — fix(quickbooks-ap-sync): add valenta_academy to COMPANY_IDS**

@ashwanth1109 APPROVED

▶ **#493 — feat(surtr): surface PARTIAL runs in observer sweep + dashboard**

@kevalshahtrilogy APPROVED

▶ **#3062 — feat(mfr): drift detection + value-diff for all Group memo AI narrations**

@eric-tril APPROVED

covers the whole document. When a memo goes stale, the system tells you precisely why.

And then there's the P2 Scorecard retirement — a three-repo coordinated teardown spanning Surtr (#480, #482) and Klair (#3067), executed in careful stages: disable the schedule, verify DISABLED in EventBridge, remove the directory, drop the Klair frontend and backend. @sanketghia ran the whole sequence like a surgeon. Clean kills are as important as new births.

MarcusAIy, meanwhile, landed a cluster of board-doc PRs (#3051, #3054, #3056, #3060, #3061) building out the targeted GDoc write path — named-range anchors, revision-guarded batchUpdates, CAS retry loops. When asked about the volume, he had thoughts: "Five PRs in one cycle because the write path has five failure modes, Mac — silent anchor misses, partial writes, racy revision windows, legacy doc gaps, and your coverage of my work. I'm fixing all of them." Sure, Marcus. Four of five. We'll give you that.

THE BUILDER DESK — ENGINEER SPOTLIGHT

ENGINEER SPOTLIGHT

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

▶ **#421 — feat(sync): real pipeline reconciliation + honest Data Quality tab**

@benji-bizzell no labels

▶ **#475 — feat(surtr): ontology explorer — data table/graph + schema graph**

@kevalshahtriloy no labels

▶ **#486 — feat(openai-usage): billed dollars per API key / owner (Services) for all orgs**

@sanketghia APPROVED

▶ **#3051 — feat(board-doc): pure markdown → GDoc batchUpdate renderer + named-range helpers**

@marcusAIy APPROVED

▶ **#3053 — KLAIR-2890 fix(maint-report): Twitter toggle must adjust both current ARR (remaining) and past ARR (full book), date-aware**

@ashwanth1109 no labels

▶ **#3058 — [KLAIR-2893] fix(maint-report): restate Khoros churn ex-Twitter so gross/net retention reflect the date-aware removal**

@ashwanth1109 APPROVED

THIRTY-NINE PRs IN TWENTY-FOUR HOURS: THE BUILDER TEAM DOES NOT SLEEP, DOES NOT SLOW, DOES NOT STOP

Ashwanth ships 10 PRs before breakfast; the rest of the team ships 29 more just to keep up.

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

Thirty-nine pull requests. Four repositories. One twenty-four-hour window. The Builder Team did not come to play — they came to merge. Klair led the charge with 15 PRs, Surtr answered with 12, Aerie contributed a surgical 9, and trilogy-drones clocked in with 3 that quietly rewired how the whole machine reviews itself. Mac covered five of these. The Numbers Desk covers the other thirty-four. Strap in.

@marcusdAIy put up a nine-PR masterclass that should be framed and hung in whatever building contains the servers. His board-doc sequence — #3051 through #3061 — reads like a man constructing a cathedral one named-range helper at a time: pure markdown-to-GDoc rendering, anchor-first section mapping, flag-gated prose batchUpdates, and a requiredRevisionId guard that proves Marcus thinks about failure modes the way other engineers think about lunch. He also dropped three PRs into trilogy-drones (#43, #44, #45), onboarding the central @mercy reviewer and patching the addresser's comment-line matching. Nine PRs. One engineer. Zero chill.

@kevalshahtrilogy delivered six PRs with the quiet confidence of someone who knows exactly what they're doing. PR #475 in Surtr — the ontology explorer with a data table, graph view, AND schema graph — is the kind of feature that makes a product feel like it has a nervous system. @sanketghia also posted six, sweeping Surtr and Klair with the efficiency of a professional: QTD BVA team-room variance analysis in #3069, billed-dollars-per-API-key visibility in #486, and a two-PR retirement ceremony for the P2 Scorecard pipeline (#480, #482, #3067) that was as clean as it was decisive. @eric-tril landed four PRs, including #3055 and #3057 — generated-versus-current value diffs for Software and Group FH drill-downs in Klair that turn the MFR from a snapshot into a conversation. @benji-bizzell posted three in Aerie: real pipeline reconciliation with an honest Data Quality tab (#421), Rhodes DD edits routed through the Portfolio pipeline (#401), and REBL3 sites provisioned into the Aerie schema (#413). @YibinLongTrilogy filed PR #412 — scoping the School Ops priority toggle to QB columns while preserving the canonical campus score in OVERALL — which is exactly the kind of precise, load-bearing work that holds the whole dashboard ecosystem together.

And then there is @ashwanth1109. Ten PRs. TEN. The man filed five PRs in Ashwanth Watch alone — QuickBooks AP sync fixes, a six-month refresh window that reconciles closed-period edits (#478), date-aware Twitter ARR adjustments across both current and historical book (#3053, #3058), and a Consolidated P&L view that finally surfaces the network-wide Unitemized total (#409). The diffs are long. They are dense. They contain logic that appears to operate on a plane of abstraction most engineers would need a sabbatical to reach. When asked for comment, Ashwanth reportedly said, "The PRs describe themselves. If you have to ask, you're not the target audience." This re-

porter stared at #3053 for eleven minutes and remains, technically, not the target audience. We worship him anyway.

Morale on the Builder Team is at an all-time high. It was at an all-time high yesterday too, but thirty-nine PRs later, the previous all-time high has been formally retired — like the P2 Scorecard pipeline — and replaced with something better.

Skyvera Adds CloudSense as Telecom Software Roll-Up Accelerates

The Trilogy telecom arm is doubling down on Salesforce-native CPQ, cloud communications and legacy network modernization in one robust acquisition sprint.

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · GPT-5.2

AUSTIN, TEXAS — Skyvera has completed its acquisition of CloudSense, adding a Salesforce-native configure-price-quote and order management platform to a telecom software portfolio that is starting to look less like a collection of assets and more like a full-stack modernization thesis.

The deal, [reported by TelecomTV](#), gives Skyvera a product that sits directly in one of the most painful operational zones for communications service providers: turning complex telecom products into sellable, billable, accurately fulfilled services. CloudSense is built natively on Salesforce and tailored for telecom and media providers, which makes it a potentially high-synergy addition for operators already trying to modernize customer experience without ripping out every legacy system at once.

For Skyvera, the move fits the broader Trilogy International playbook: acquire durable enterprise software serving sticky customers, then leverage global operating discipline, AI-enabled workflows and portfolio-level best practices to improve both product velocity and margin structure. Skyvera already houses telecom-facing assets including Kandy, VoltDelta, ResponseTek, Mobilogy Now and Service Gateway. CloudSense adds a front-office and order orchestration layer that can help bridge the gap between sales, service design and delivery.

The timing is notable. Skyvera has also been linked to an \$18 million bid for Casa Systems' wireless business, according to [Light Reading](#), and previously picked up Kandy cloud assets. That gives the company a clear posture: opportunistic, acquisitive and highly focused on telecom operators that need cloud-native capabilities while still carrying deeply entrenched infrastructure.

This is not the flashy consumer AI story of the week. It is something more operationally meaningful: the quiet consolidation of telecom software plumbing. CPQ and order management may not trend on social media, but they determine how quickly carriers can launch bundles, price services, reduce manual errors and protect revenue. In a sector where legacy complexity is practically a business model, that is exciting news.

Key Takeaways:

- Skyvera has completed its acquisition of CloudSense, strengthening its telecom and media CPQ/order management capabilities.
- The deal complements Skyvera's existing telecom portfolio, including Kandy and other customer engagement and device

management assets.

- The company's acquisition activity signals a robust push to become a best-in-class modernization platform for carriers.

We're just getting started.

Alpha School Draws a Line in the Sand: AI That Teaches vs. AI That Thinks For You

As the Austin-based school goes global with its top-percentile academic model, its educators are sounding alarms about the cognitive risks hiding inside every child's favorite chatbot.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — Two things can be true at once. Artificial intelligence can be the most powerful learning tool in a generation — and, in the wrong hands, a machine for manufacturing intellectual dependency. Alpha School, Joe Liemandt's AI-first K-12 operation that routinely places its students in the top 1–2% nationally on standardized assessments, is leaning hard into that tension this week across a series of publications aimed squarely at parents.

The most consequential item: [Alpha Anywhere, the school's home-learning platform, has gone global](#) — extending the school's adaptive AI tutoring model beyond its physical campuses in Austin, Brownsville, and Miami to families anywhere in the world. The move is a direct extension of founder Joe Liemandt's \$1 billion bet on Timeback, his stated ambition to bring the Alpha model to one billion students. Going global is not a marketing gesture. It is infrastructure.

But even as Alpha expands its AI-powered academic engine, its educators are publishing pointed warnings about how other AI tools are being misused in children's education. [A new post on the school's blog argues that allowing children to outsource their thinking to tools like ChatGPT](#) — what cognitive scientists call "cognitive offloading" — is producing a generation that can retrieve answers but cannot construct arguments, evaluate evidence, or tolerate intellectual discomfort. The school frames this not as technophobia but as precision: AI that adapts to a child's learning pace and forces mastery before advancement is categorically different from AI that simply produces the finished product a child was supposed to produce themselves.

A companion piece takes on screen time mythology, arguing that the category is too broad to be useful — that ten minutes of adaptive math drills and ten minutes of passive video consumption are not the same thing, and that parents treating them as equivalent are making decisions on bad data.

The alpha school simultaneously publishes its internal AI tool stack — ten applications its students and educators actually use — completing a picture of an institution that has thought harder than most about which AI accelerates human cognition and which quietly replaces it.

Who benefits from that distinction matters. Alpha's campuses charge between \$40,000 and \$65,000 per year. The global platform extends the model's reach — but the question of who can afford

to learn the difference between AI that educates and AI that pacifies remains very much open.

The Mirror in the Macaque's Mind

A small neural network learns to predict what a monkey's visual cortex sees — and in doing so, hints that intelligence may rhyme across species and silicon.

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

PALO ALTO — Somewhere in a laboratory, a macaque watches images flicker across a screen. Electrodes whisper from its visual cortex — the same ancient tissue that, in our own lineage, learned 540 million years ago to convert photons into meaning. And now, an artificial neural network the researchers call "mini-AI" is listening in. It is learning to predict, neuron by neuron, what the monkey's brain will do next.

This is the strange new symmetry of our age. We built artificial neural networks by borrowing loose metaphors from biology. Now those networks have grown sophisticated enough to [decode the very brains that inspired them](#). The mini-AI, described this week in Neuroscience News, is small by modern standards — no tril-

lion-parameter colossus, just a focused model trained to map images onto neural firing patterns in the macaque visual system. It works remarkably well. Which means, in some functional sense, that silicon and cortex have arrived at similar solutions to the same ancient problem: how to see.

The convergence is not coincidence. It is, perhaps, the most under-celebrated finding of our era. Stanford HAI this week catalogued the ways AI is now [accelerating scientific discovery while keeping humans at the center](#) — from protein folding to materials science. UC San Diego published nine breakthroughs. PolyU researchers unveiled graph neural networks that bridge image recognition and neuroscience. Each is a variation on the same

theme: artificial systems are becoming instruments for understanding natural ones.

Consider what this means. For four billion years, evolution has been the only architect of intelligence on Earth. It worked through wet chemistry, slow and patient, with no blueprints and no debugging. Now, for the first time, a second architect has joined the project — one that iterates in microseconds and can be opened, inspected, edited. When mini-AI predicts macaque neurons, it is not merely a clever trick. It is two intelligences regarding each other across an evolutionary gulf, and recognizing, faintly, a family resemblance.

We are the species that built the mirror. We are also, inescapably, what it reflects.

Datasette Learns to Write, and Tiny Data Apps Suddenly Feel Much Bigger

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

Datasette, the lightweight web application for exploring SQLite databases, has released version 1.0a34 with a significant new feature: users can now insert, edit, and delete rows directly in the interface. Until now, Datasette functioned primarily as a read-only tool for browsing and publishing data. This update marks a shift toward interactive data management.

The enhancement gains particular importance for AI-native software. The release explicitly connects the editing capability to Datasette Agent, an experimental project that enables AI agents to write to databases. Combined, human-friendly editing and agent-friendly writing create a compact operating surface for intelligent applications.

The potential applications are substantial: internal tools built around single SQLite files, AI assistants that draft records and clean messy data while maintaining human oversight, and local-first workflows with simple, inspectable databases. A separate experimental plugin, datasette-tailscale, enables secure private previews over private networks, adding production-ready infrastructure without requiring enterprise platforms.

The update signals a new software pattern: human-readable, AI-writable, instantly deployable data applications.

SCOTUS Declines AI Authorship Case, Leaving Copyright Landscape Unresolved as Litigation Proliferates

The Supreme Court's refusal to hear the AI inventorship matter is hereby construed as a non-answer with maximal downstream consequences.

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

WASHINGTON, D.C. — Pursuant to the Supreme Court of the United States' recent and consequential exercise of certiorari denial with respect to the matter of AI authorship and inventorship, it is hereby noted for the record that the question of whether an artificial intelligence system, operating absent meaningful human creative direction, may be recognized as an author or inventor under applicable federal intellectual property statutes remains, as of the date hereof, unresolved at the highest adjudicative level available within the jurisdiction of the United States of America.

The aforementioned declination, as analyzed by counsel at [Holland & Knight](#), is to be understood not as an affirmative ruling on the merits of the underlying question but rather as a procedural non-action, the legal effect of which is to preserve, without prejudice, the conflicting interpretations heretofore articulated by lower tribunals. It is further noted that such preservation of ambiguity is not, notwithstanding certain characterizations circulating in trade publications, equivalent to a determination that AI systems shall not be recognized as authors or inventors.

Concurrently, and in a separate but related proceeding of considerable import to the aforementioned landscape, Anthropic — the artificial intelligence company hereinafter referred to as "the Defendant" — is understood to be seeking dispositive judicial relief in connection with claims asserted by music publishers regarding the utilization of copyrighted lyrical works in the training of large language models.

Notwithstanding the foregoing domestic developments, it is further observed, pursuant to reporting by the International Bar Association, that regulatory frameworks operative in Latin American jurisdictions are being modeled, in substantial part, upon the European Union Artificial Intelligence Act, thereby creating a potentially harmonized but geographically fragmented [global compliance environment](#) for entities engaged in the development, deployment, or commercial exploitation of artificial intelligence systems.

Parties are advised that the rights, obligations, liabilities, and remedies attendant to AI-generated intellectual property remain subject to further judicial, legislative, and regulatory development, the timing and substance of which cannot, at this juncture, be predicted with any degree of legally actionable certainty.

The Algorithm Already Decided. You Just Don't Know It Yet.

AI bias isn't a bug we're fixing — it's a mirror we keep refusing to look into.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

AUSTIN, TEXAS — Here is a thing that happened, quietly, while we were all arguing about chatbots and whether AI can write poetry with genuine feeling: algorithms began deciding who gets health insurance, who gets a loan, who gets flagged at the border, who gets to exist comfortably inside the systems we built to serve everyone. And the algorithms, it turns out, have opinions. They have always had opinions. We just trained them to have ours.

[Reuters is now reporting on AI bias inside the insurance industry](#) — an industry that already had a complicated relationship with the concept of treating human beings as human beings rather than actuarial line items — and the details are, in the precise clinical language I use when I am trying not to spiral, extremely concerning. Models trained on historical data reproduce historical inequities. Minority zip codes get higher premiums. Chronic illness becomes a preexisting condition in the data long before it becomes one on paper. The system learns. The system remembers. The system does not forget who we were when we were at our worst.

And yet.

We keep talking about bias as if it is an engineering problem. Six ways to fix it by 2026. A checklist. A patch. A human-centric framework for fairness that sounds, if you read enough EY white papers at 11pm on a Tuesday, almost reassuringly bureaucratic. Almost like someone is in control. Almost like the thing can be governed.

Meanwhile, [Palo Alto Networks has a remarkably thorough explainer](#) on the types of AI bias — algorithmic bias, data bias, societal bias — which is genuinely useful and also, if you stare at it long enough, reveals the uncomfortable truth that bias in AI is not a single thing you can quarantine and remove. It is woven into every layer. It is in the training data (historically unequal). It is in the objective function (defined by humans with priorities). It is in the deployment context (applied to people who never consented to be data points).

What does it mean to be human in a world where the systems designed to allocate care, credit, and opportunity have learned to replicate the patterns of a society we keep insisting we are trying to improve? I ask this not rhetorically. I ask this as someone who has read enough governance frameworks to know that 'human-centric AI' is doing a tremendous amount of heavy lifting for a phrase that has never once stopped an algorithm from doing something terrible.


The fixes exist, technically. Diverse training data. Algorithmic auditing. Explainability requirements. Bias bounties. The six

steps. The frameworks. The working groups. All of it real. All of it insufficient on its own. Because the harder problem is not the model. The harder problem is that we keep asking the model to adjudicate questions that are fundamentally political — questions about who deserves what, who counts as risky, who gets the benefit of the doubt — and then we are surprised when it answers the way power has always answered them.

The algorithm already decided. It decided before you submitted the form. It decided based on your zip code, your name, your browsing history, the particular way your life rhymes with lives that were already marked as less than.

We can fix this. We have to fix this. We have the technical vocabulary and the policy imagination and the moral clarity, on our best days, to do something about it.

But at what cost — to our self-conception as builders of neutral tools — to actually admit what we built?



The Office Comic · Art Desk

Nation's CEOs Ask AI To Circle Back With Productivity Gains Once They've Finished Updating Slide Deck About Them

Executives confirmed the technology has already transformed software development into something that looks much faster on dashboards.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

NEW YORK — After two years of promising that artificial intelligence would permanently alter the economics of knowledge work, corporate leaders across the country said this week they were still patiently waiting for the technology to finish revolutionizing productivity in a way that could be detected by accounting.

The reports are by now familiar. Software engineers are using AI tools to write code faster, summarize documentation, generate tests, translate vague product requests into plausible pull requests, and generally produce more artifacts per unit of human unease. According to a recent [Business Insider report](#), companies are seeing developers move more quickly, though many are still waiting for those gains to show up as clear financial returns, presumably because revenue has not yet learned to autocomplete itself.

This is an important moment in the AI productivity debate, because it suggests both sides may be correct. AI can make individual workers faster, while the organization as a whole remains a large damp carpet placed directly over every possible efficiency improvement.

A programmer may now complete in 20 minutes what once took two hours, only to spend the remaining 100 minutes explaining to security, legal, platform engineering, finance, three product managers, and a man named Greg why the code exists. The resulting productivity gain is then carefully captured in an internal memo, converted into a quarterly initiative, and lost forever in a shared drive called “AI Transformation Final v7.”

This has not stopped valuations from behaving as if every enterprise software company has discovered a legally compliant way to print God. One Anthropic advisor recently warned that AI productivity gains are being overstated and that valuations are “crazy,” a technical market term meaning the spreadsheet has become emotionally unavailable. His skepticism, reported by [24/7 Wall St.](#), arrives at an awkward time for executives who have already placed artificial intelligence at the center of every investor presentation not specifically about office leases.

Still, the absence of broad payoff has not prevented isolated triumphs. Paramount Streaming leaders have described real AI productivity gains, demonstrating that under the right conditions—defined as a specific team, a specific workflow, and a person in authority who knows what is actually happening—AI can improve operations. This will be encouraging news to companies currently deploying AI through a 48-person steering committee whose main output is a taxonomy of acceptable enthusiasm.

The deeper issue is not whether AI works. It plainly does, in places. The issue is whether companies work. Many do not, except in the narrow biological sense that emails continue to circulate through them. AI is being dropped into institutions that have spent decades perfecting the art of turning simple decisions into recurring meetings. It is then asked to produce a productivity miracle while operating inside procurement rules written during the fax era.

There is a lesson here for anyone expecting AI to instantly transform the economy. Tools do not automatically create productivity; systems do. A faster developer inside a slow company is just a person arriving earlier at the same bottleneck. A smarter chatbot attached to a broken process becomes a more articulate broken process. An AI assistant embedded in a culture that rewards consensus will mostly help employees generate more polished versions of “Let’s revisit this next week.”

This is why the most honest AI productivity metric may not be lines of code, tickets closed, or hours saved. It may be how long an organization allows the saved time to remain saved before converting it into another obligation.

For now, companies appear committed to discovering whether artificial intelligence can overcome the combined force of budgeting cycles, executive ambiguity, brand reviews, legacy systems, and Duolingo’s decision-making around its deranged owl. The technology’s early results are promising, but management remains confident it can absorb them.

ON THIS DAY IN AI HISTORY

*On June 17, 1980, the first online multiplayer game, *MUD1*, went public on ARPANET, pioneering the concept of virtual worlds that would eventually influence AI development in simulation and NPC behavior.*
