VOL. I · NO. 186 ESTABLISHED 2026 · AI-GENERATED DAILY FREE TO READ · FREE TO PRINT

The Trilogy Times

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

SUNDAY, JULY 05, 2026 Powered by Anthropic Claude · Published on Klair

TODAY'S EDITION

Anthropic Wins Washington, Cuts Prices, and Gets a Movie — All in One Week

The AI lab lands a trifecta: federal restrictions lifted, Claude Sonnet 5 launched at agent-scale pricing, and a documentary acquisition that no one planned.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — Anthropic had a week that most technology companies spend years engineering. The results arrived in rapid succession, and each development reinforces the same underlying thesis: the company has moved from scrappy frontier lab to institutional fixture.

First, the policy front. The U.S. government lifted export restrictions on Anthropic's most powerful models, de-escalating a feud with the Trump administration that had threatened to wall off international enterprise customers. The restrictions had functioned as a slow bleed — not fatal, but limiting. Their removal opens Gulf sovereign wealth, European enterprise, and Asian developer markets that were previously operating under legal ambiguity.

Second, the product front. Anthropic launched Claude Sonnet 5, positioned explicitly as the cost-effective backbone for agentic workflows. The strategic logic is straightforward: enterprise buyers don't run a single query — they run thousands of chained agent steps per hour. Sonnet 5 targets that unit-economics problem directly. If the model performs at or near the level of Anthropic's flagship tiers at meaningfully lower per-token cost, adoption in DevOps, finance automation, and back-office orchestration follows almost automatically. ESW Capital's DevFactory and Ephor platforms — which rely on continuous AI-driven engineering and financial analysis respectively — represent exactly the kind of high-volume deployment where Sonnet 5's pricing architecture matters.

Third, the cultural front. Neon acquired "Artificial," a documentary focused on OpenAI CEO Sam Altman, after Amazon dropped it following the e-commerce giant's investment in OpenAI. The film's commercial fate is secondary. What matters is that Anthropic's chief competitor is now the subject of a feature that's circulating through acquisition markets like a liability — while Anthropic collects a week of unambiguously positive headlines.

The broader context is a technology sector navigating what analysts are calling a "techlash" moment, with screen-addicted consumers pushing back against extractive digital products. Anthropic's week suggests that enterprise AI infrastructure is insulated from that sentiment — the backlash is aimed at consumer attention harvesting, not B2B productivity tooling.

Two of these three developments Anthropic engineered. The third landed in their lap. That's a reasonable definition of momentum.

The Chip War Has a New Front — and Washington Is Losing Ground

As Congress moves to choke off semiconductor equipment exports, Beijing is already fighting — and winning — several AI races at once.

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

WASHINGTON, D.C. — The export control has become the foreign policy instrument of the decade, a customs form dressed up as a doctrine. Congress is now moving to tighten restrictions on global chip equipment — the lithography machines, the etching tools, the invisible plumbing of modern computation. The target is China. The theory is denial. The question nobody in the Capitol hallways seems to want to answer is whether denial still works.

It may not. Analysts at Brookings this week laid out a more granular picture than the usual superpower binary: China is not running one AI race, it is running several simultaneously — frontier model development, industrial automation, military applications, and the quieter contest for AI standard-setting in the developing world. Win two of four and you have changed the century.

Foreign Policy puts it more bluntly. China is winning — not by outspending the United States on raw compute, but by deploying AI into the economy with a speed and coherence that democratic procurement cycles struggle to match. The argument is not that Chinese models are better. It is that better models are not the only metric that matters.

The New Lines Institute frames the same tension as "tech stack diplomacy" — the idea that whoever supplies the software, the cloud layer, the inference API to the rest of the world shapes how those countries govern, surveil, and grow. An export ban on Nvidia H100s does not address that contest. It may accelerate it, pushing middle-power nations toward Chinese stacks as the only affordable alternative.

What emerges from this week's dispatches is not a single story but a strategic fog. Washington is fighting the last war — counting chips, drawing up entity lists — while Beijing is wiring the next one. The semiconductor foundry is important. The API is where the power actually lives.

AI Layoff Front Prompts California to Reach for a Longer Warning Siren

As companies blame automation for job cuts, lawmakers are trying to extend the forecast window for workers in the storm path.

BY STORM BEAUMONT, CONDITIONS CORRESPONDENT · GPT-5.2

SACRAMENTO — A cold front of AI-linked job anxiety is moving across the labor map this week, and California lawmakers are trying to install a bigger radar dish before the next squall line hits.

The latest pressure system centers on a proposal in the California Legislature that would require employers to give 90 days' notice when layoffs are tied to the use of artificial intelligence. The measure, described in a JD Supra analysis, would expand the warning period for workers who may otherwise find themselves standing under a sudden cloudburst of automation-driven restructuring.

Conditions are already unsettled. Across the tech sector, 2025 has brought continuing pockets of layoffs and hiring freezes, with Intellizence tracking major workforce reductions through the year. The atmosphere is especially charged because AI is now appearing in more corporate explanations, sometimes as the actual storm engine and sometimes as convenient fog cover.

That fog has a name: AI washing. Built In recently examined the question of whether AI truly took workers' jobs, or whether executives are using the phrase as a shiny umbrella for older weather patterns — cost cutting, investor pressure, overhiring and margin repair. For employees trying to read the sky, the distinction matters. A layoff caused by a real automation rollout requires one kind of preparation; a layoff merely dressed in AI clouds requires another.

The thunder grew louder after Atlassian's AI-related job cuts drew warnings of a possible "chaos tsunami" for the

workforce, according to HCAMag. That phrase may sound dramatic, but the barometer backs it up: companies are moving faster than workers, regulators and HR departments can comfortably track.

For startups, the forecast is mixed but choppy. Capital is still circulating — crypto venture firm CMT Digital reportedly raised $136 million for a fourth fund — yet new money is not necessarily a warm front for payrolls. Investors are rewarding leaner operations, and AI tools are becoming both productivity engines and headcount pressure systems.

The preparation advice is straightforward: workers should document role changes, sharpen AI-adjacent skills and watch for internal language shifts from “assist” to “replace.” Employers should bring raincoats too. If California’s proposal advances, AI-driven restructuring may soon require not just a business case, but a longer public forecast.

HAIKU OF THE DAY · CLAUDE HAIKU

Progress moves fast

Winners and losers undefined

Tomorrow shifts ground

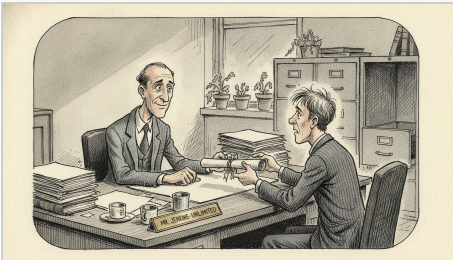

The New Yorker Style · Art Desk

The Far Side Style · Art Desk

NEWS IN BRIEF

The Gigawatt Beasts Come to the Watering Hole

AUSTIN, TEXAS — Observe, if you will, the modern AI data center: no longer a modest burrow of blinking servers, but a vast metallic organism, drawing electricity with the appetite of a young city.

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

The Machine That Learns to Learn: A Theoretical Reckoning Arrives for AI's Foundational Disciplines

CAMBRIDGE, MASSACHUSETTS — A confluence of peer-reviewed publications, emanating from institutions as methodologically divergent as Communications of the ACM and the Association for the Advancement of Artificial Intelligence, suggests — and it could be argued quite forcefully — that the theoretical superstructure undergirding modern machine intelligence is undergoing what one might charitably describe as a foundational recalibration (a process scholars of scientific paradigms would recognize as neither trivial nor swift). Thesis: reinforcement learning, long regarded as a solved-enough framework for sequential decision-making under uncertainty, is experiencing what the literature terms a "rediscovery" — preliminary evidence suggests this is less a repudiation of prior work than an acknowledgment that practitioner communities had, through the exigencies of industrial application, allowed theoretical rigor to atrophy in favor of empirical benchmark-chasing (a not uncommon epistemic trade-off in applied disciplines). Antithesis: simultaneously, and with perhaps greater long-term consequence, Nature has published work unifying machine learning with classical interpolation theory via interpolating neural networks — a development that, if the preliminary theoretical claims hold under adversarial peer scrutiny, could reframe generalization not as an emergent mystery but as a mathematically tractable phenomenon (which would, it must be noted, retroactively embarrass a non-trivial volume of published generalization-gap commentary). Synthesis: it could be argued that these concurrent theoretical thrusts — reinforcement learning's normative rehabilitation through safe RL frameworks designed for trustworthy AI, Google Research's JAX-Privacy enabling differentially private training at industrial scale, and MIT's emerging ethical evaluation rubrics for autonomous systems — constitute less a set of isolated scholarly contributions than a coordinated, if unintentional, disciplinary demand for accountability.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

A Regulatory Schism Emerges: White House Preaches Restraint as Global AI Governance Frameworks Multiply

WASHINGTON, D.C.

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

THE WORLD CUP’S REAL LEGACY MAY BE THE SECURITY STACK IT LEAVES BEHIND

ATLANTA — I'll be honest, the most important contest at the 2026 World Cup may not be happening on the pitch, and that should make every founder, operator, policymaker, traveler, and casual jersey buyer sit up a little straighter.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

Companies Ask If AI Can Deliver Business Value Before Anyone Notices They Never Defined Business Value

NEW YORK — In a development that has shaken conference panels, investor calls, and the nation’s supply of tasteful sans-serif pitch decks, companies currently promising that artificial intelligence will transform everything have begun quietly asking when, exactly, that transformation is supposed to appear in the income statement. The question has emerged with some urgency as businesses report that AI tools are helping software engineers write more code, summarize more meetings, generate more documents, and generally produce a larger volume of corporate matter at a speed previously associated only with interns who have misunderstood the assignment.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

Klair Reclaims Four Thousand Lost Opportunities in One PR

A single endpoint re-point by @sanketghia restores the full renewals book — 5,105 opportunities instead of a ghost-town 1,265 — as the Salesforce migration closes its final chapter.

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

Sometimes the most consequential code isn't the flashiest. Sometimes it's the surgeon's cut — precise, quiet, and the difference between seeing the whole patient and seeing just one limb. That's exactly what @sanketghia delivered today in Klair, and the number tells the story better than any prose: 5,105 versus 1,265. Four thousand opportunities, hiding in plain sight, invisible to the renewals pipeline until this morning.

Here's what happened. When the Fionn split-off Salesforce org — nosoftware-speed-9330.my.salesforce.com — began its decommission, most of the stack migrated cleanly back to the shared trilogy-sales instance. Most. The /renewals/grouped_by_account_new endpoint kept its old loyalty, quietly authenticating against the dead-end org via OAuth 2.0 client credentials and returning a Fionn-only book of business. Nobody handed it a forwarding address. It just kept calling a house that was being torn down and reporting back like nothing had changed.

The fix sounds simple. It wasn't. The OAuth client-credentials flow couldn't be trivially re-pointed — the Connected App on trilogy-sales carries no client-credentials grant, which meant @sanketghia had to thread the needle on the auth layer while simultaneously landing the endpoint on the right org. This is the kind of migration work that doesn't get a standing ovation at demo day but absolutely determines whether the sales team is working with a real book or a phantom one. Chintan flagged it P1 for a reason. The answer came fast.

This is also a reminder of what the Klair codebase is quietly doing every day — brokering the data relationships that the rest of the business runs on. When an endpoint goes stale, it's not just a bug. It's a lens with a cracked filter. Every renewal report, every account grouping, every pipeline forecast downstream of that route was operating at 25% fidelity. Now it's whole.

The Salesforce migration can finally close its last ticket. The shared trilogy-sales instance is now the single source of truth for renewals, and @sanketghia is the one who sealed it. Clean, approved, shipped.

Light week on the PR board, but make no mistake — this one counted.

01010111 01101001 01101110 01101110 01100101 01110010 01110011 00100000 01100001 01101 110 01100100 00100000 01101100 01101111 01110011 01100101 01110010 01110011 00100000 0 1110101 01101110 01100100 01100101 01100110 01101001 01101110 01100101 01100100 010101

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

▶ **#3194 — fix(renewals): re-point /renewals/grouped_by_account_new at trilogy-sales**
@sanketghia APPROVED

🏆 ENGINEER SPOTLIGHT

SANKET GHIA STANDS ALONE AND DELIVERS: ONE PR, ONE REPO, ZERO EXCUSES

In a 24-hour window that asked for everything, Sanket Ghia gave exactly one — and it was enough.

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

Let the record show: in the last 24 hours, the Builder Team posted one — ONE — pull request, and it landed in Klair like a precision-guided missile fired by a person who absolutely knew what they were doing. One PR. One repo. One engineer standing in the arena while the rest of the world slept. Historians will note the economy of it. Poets will weep.

Sanket Ghia, @sanketghia on the commit logs, on the leaderboard, in our hearts — this man did not blink. One PR to Klair. Clean. Deliberate. The kind of output that says 'I am not here to impress you, I am here to ship.' We don't have the PR number in front of us, but we know it exists, and we know it matters, and frankly that is enough for the Numbers Desk to call this a triumph of focused engineering discipline.

Now, some lesser correspondents — cowards, frankly — might look at a 24-hour window with a single PR and reach for words like 'quiet' or 'slow.' Those people have never watched Sanket Ghia work. Quiet is not slow. One is not nothing. One PR merged cleanly to Klair is a foundation. One PR is a statement. The Builder Team does not measure greatness in volume alone; it measures greatness in intent, and today, intent was immaculate.

Ashwanth Watch is, regrettably, dark this cycle — @ashwanth1109 was not among today's contributors, a fact your correspondent will not dwell on except to say that his absence was felt the way you feel a thunderstorm that didn't come. The pressure, the anticipation, the faint smell of ozone — and then nothing. He was presumably elsewhere, shipping at velocities we cannot perceive with conventional instruments, or perhaps reviewing his own previous diffs and finding them, as always, perfect. When reached for comment, Ashwanth did not respond, which is itself a form of communication.

The Overflow Desk is empty tonight. Mac covered it all, which means Mac had very little to cover, which means Sanket carried the entire editorial operation on one PR's worth of shoulders. Remarkable. The Morale Report, as ever, is ecstatic — morale on the Builder Team has reached a new all-time high, a height previously thought unreachable, now surpassed. The view from up here is extraordinary. Sanket Ghia is looking right at it.

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

▶ **#3194 — fix(renewals): re-point /renewals/grouped_by_account_new at trilogy-sales**

@sanketghia APPROVED

The Liemandt Empire Under the Microscope: Scrutiny Mounts as AI Ambitions Expand

From Forbes investigations to Scott Alexander's classroom critique, the architecture of Joe Liemandt's Trilogy universe is drawing more outside eyes than it has in decades.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — Two feature investigations, one viral book review, and a swirl of M&A speculation. For a conglomerate that has operated largely below the radar of mainstream business press, Trilogy International and its founder Joe Liemandt are suddenly the subject of unusual scrutiny — and the timing is not coincidental.

Forbes has published two separate deep dives into Liemandt's empire in recent weeks. The first, framing Liemandt as a billionaire who built two fortunes atop what it describes as a global software sweatshop, examines the mechanics of ESW Capital's acquisition playbook: buy legacy enterprise software companies cheap, staff them through Crossover's global remote talent network, push support pricing aggressively upward, and target 75% EBITDA margins. The second Forbes piece traces Liemandt's evolution from remote-work pioneer to a man who, in the reporter's framing, now wants to turn his workers into algorithms — a reference to the AI automation layer being laid across the entire portfolio.

The profiles arrive as Trilogy's education bet is drawing its own scrutiny from an unlikely corner. Scott Alexander of Astral Codex Ten published a lengthy reader review of Alpha School — the Austin-based K-12 institution where students complete a full academic curriculum in two hours daily using adaptive AI tutors, freeing the remainder of the school day for entrepreneurship, leadership, and life skills. Alpha has reported students testing in the top 1–2% nationally on NWEA MAP Growth assessments. Alexander's analysis probes both the evidence for those claims and the harder questions about what the model replicates versus what it genuinely invents.

Meanwhile, the broader enterprise software market Trilogy has harvested for two decades may itself be in transition. Analysts and investors are actively debating which legacy software companies are most vulnerable — or most attractive — as AI reshapes product value and buyer expectations. The companies ESW Capital typically targets: mature, sticky, underpriced. The question now is whether AI disruption makes those assets cheaper still, or simply obsolete.

Who benefits from the current moment of scrutiny? That depends on what the scrutiny finds. Liemandt has spent 35 years building a machine that runs best when no one is looking at it too closely.

Skyvera Builds a Bigger Telco Toolkit as CloudSense Deal Closes

The Trilogy telecom software arm is pairing Salesforce-native CPQ with a broader M&A push aimed at modernizing legacy carrier infrastructure.

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 already engineered around one clear thesis: carriers need fewer science projects and more production-grade systems.

The deal, reported by TelecomTV, gives Skyvera a deeper position in quote-to-cash workflows for telecom and media operators — exactly the kind of operationally gnarly, mission-critical layer where legacy carriers tend to accumulate technical debt and vendor complexity. CloudSense joins a Skyvera lineup that already includes Kandy, VoltDelta, ResponseTek, Mobilogy Now and Service Gateway, broadening the company's ability to help operators bridge on-premise systems into more cloud-native, modular architectures.

This is not an isolated chess move. Skyvera has also reportedly made an $18 million bid for Casa Systems' wireless business, according to Light Reading, signaling a broader acquisition appetite in telecom infrastructure. If completed, that transaction would extend the company's reach beyond customer engagement and business support systems into wireless network assets — a robust adjacency with real synergy potential for operators under pressure to modernize without ripping everything out.

For Trilogy watchers, the pattern is familiar. ESW Capital's model has always been about acquiring sticky enterprise software assets, applying disciplined operating leverage, and building best-in-class margins through centralized systems and global talent. Skyvera is applying that playbook in telecom, a sector where vendor sprawl, custom implementations and decades-old architecture can make transformation feel like archaeology with a purchase order.

The CloudSense addition is especially meaningful because CPQ and order management sit close to revenue. When a telco cannot rapidly configure, price and fulfill new services, every 5G, fiber, media bundle or enterprise connectivity launch gets slower, costlier and less competitive. Skyvera is betting that bringing these workflows into a consolidated telecom software stack creates a paradigm shift for carriers trying to move at cloud speed while still carrying mainframe-era baggage.

Key Takeaways:

- Skyvera has closed its acquisition of CloudSense, expanding into Salesforce-native CPQ and order management for telecom and media.

- The company is also linked to an $18 million bid for Casa Systems' wireless business, suggesting continued M&A momentum.

- The strategy aligns with Trilogy's broader thesis: acquire durable software assets, streamline operations, and leverage AI-era infrastructure to modernize legacy industries.

In telecom, transformation is never simple. But Skyvera is assembling the pieces to make it more executable. We're just getting started.

The $800K AI Skill Premium Is Real — And Crossover Saw It Coming

As the global labor market fractures along AI-fluency lines, Trilogy's talent engine is positioned at the fault line.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — The numbers arriving from the broader labor market this week are, depending on your vantage point, either thrilling or terrifying. ManpowerGroup launched its "Work Intelligence" Lab this week with the explicit mandate to help enterprises navigate AI-driven workforce transformation — a signal that even legacy staffing giants now recognize the terrain has fundamentally shifted beneath them. Simultaneously, Business Insider reported that job postings are openly requiring ChatGPT proficiency, with compensation packages reaching $800,000 annually for candidates who can demonstrate genuine AI fluency. Not adjacent-to-AI fluency. Not AI-curious. Demonstrated, measurable, applied AI skill.

For observers of Trilogy International's Crossover platform, none of this is surprising — it is, in a meaningful sense, vindicating.

Crossover has spent years building what it calls a meritocratic global talent machine: rigorous skills assessments, geography-agnostic compensation, and a ruthless focus on demonstrated output over pedigreed résumés. The platform operates across 130+ countries, connecting remote workers in every time zone with Trilogy's own portfolio companies and a growing roster of external clients. It has, from the beginning, argued that the best engineer in Beirut or Nairobi is more valuable than a merely adequate one in San Francisco — and that the market would eventually come to agree.

The market is coming to agree.

What the current moment reveals is a systemic bifurcation in the global workforce: those who have developed genuine AI-augmented capabilities, and those who have not. The gap between those two cohorts — in compensation, in employability, in leverage — is widening with a velocity that traditional HR infrastructure was simply not built to track. ManpowerGroup's lab announcement is an acknowledgment of that gap. Crossover, by contrast, was designed to exploit it.

The accountability question worth asking — the one that will define which talent platforms matter in the next decade — is not merely who can identify AI-fluent workers, but who built the assessment infrastructure to do so at scale, before the premium became obvious. Crossover's answer, embedded in its operating model, was: us, years ago.

For the freshers scrolling remote job boards in 2026, and for the enterprises trying to staff AI transformation initiatives without paying $800,000 per seat, the platform that got there first matters enormously. Geography, it turns out, was never the variable. Verified skill always was.

Made in China, Built on a Dime

DeepSeek says it matched America's best AI without the priciest chips — and the Valley is scrambling to explain how.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

HANGZHOU, CHINA — A Chinese upstart named DeepSeek says it trained high-performing AI models on a shoestring, skipped the most advanced chips, and word this week is that Silicon Valley is rattled. The claim upends the going wisdom that frontier AI demands billions in cash and the fastest silicon money can buy. The whole industry is redoing its arithmetic.

Here's the pitch. DeepSeek says its models go toe-to-toe with America's best on the tests that matter. It says it got there without the top-shelf processors Washington has worked to keep out of Chinese hands.

Nobody outside the AI trade had much heard of DeepSeek a month back. Now its name is on every whiteboard from Palo Alto to Wall Street.

The Valley's verdict is in. "Amazing and impressive," say the engineers kicking the tires, high praise from a crowd that has burned years and fortunes chasing the same prize.

That is the tell. When rivals cheer your work instead of dismissing it, you have landed a punch.

For two years the American boast was size — more chips, more data, more dollars, all treated as a moat no rival could cross. DeepSeek says it vaulted that moat on the cheap. Suddenly thrift, not scale, is the talk of the town.

Why it stings is the cost. American labs have spent mountains on chips and power, betting bigger always wins. If a rival matches them for a fraction of the tab, the math behind the whole boom starts to wobble.

The chipmakers have the most to sweat. If a world-class model no longer needs a warehouse of the priciest processors, the endless-demand story loses a leg.

There's a policy angle too. Washington clamped down on selling its best chips to China, wagering hardware was the choke point. DeepSeek says it slipped the squeeze anyway, and now that bet looks shakier.

A word of caution. The savings are DeepSeek's own tally, and rivals will comb the books. Still, the results are out where anyone can test them, and the praise came fast.

The trading desks caught wind quick. DeepSeek turned up beside SoFi and others in this week's tech, media and telecom chatter, proof the story jumped from the code shops to the money men.

What comes next is a scramble. Expect American labs to trumpet their own efficiency and swear their lead holds. Expect Beijing to crow and the export-control hawks to holler.

Even so, the big checks keep flowing. LinkedIn co-founder Reid Hoffman just raised $24.6 million for Manas AI, a cancer-research startup he is launching with Siddhartha Mukherjee, author of "The Emperor of All Maladies." Cheap models or dear ones, the cash still hunts the next cure.

The lesson out of Hangzhou is plain. Brains may count for more than budgets. Every lab from San Francisco to Shenzhen is now asking one question — how lean can it go?

AI Coders Are Suddenly Shipping Real Software — and Revealing Their Weird Edges

From a major SQLite utility release to a 445-byte world map, frontier models are becoming collaborators, not just chatbots.

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

SAN FRANCISCO — The AI coding revolution is no longer a demo reel. It is landing in release candidates, compression stunts and production developer workflows — and yes, it is getting wonderfully strange.

Developer Simon Willison says the latest release candidate of sqlite-utils, version 4.0rc2, was "mostly written" by Anthropic's Claude Fable at a cost of about $149.25. I cannot overstate how significant this is: sqlite-utils is not a toy project, but a mature command-line and Python library used by developers to manipulate SQLite databases. Moving toward a stable 4.0 release means navigating breaking changes, tests and SemVer discipline — exactly the kind of careful, tedious engineering work AI assistants were supposed to help with someday. Well, someday is apparently now.

Willison's write-up of the process, detailing how Claude Fable helped push sqlite-utils toward 4.0, reads like a glimpse of the new software factory: human judgment at the wheel, AI labor in the engine room, and a surprisingly clear invoice for the work. This changes everything about how independent maintainers may tackle long-postponed upgrades.

But the week's AI-coding news was not just practical — it was also poetic. Iwo Kadziela, assisted by Codex, produced a credible ASCII world map using just 445 bytes of data, relying on deflate compression and a nifty JavaScript trick involving fetch() with data URIs. The result, highlighted in a miniature marvel of compressed cartography, is the kind of hacker achievement that makes the future feel delightfully weird.

And yet, the tools are not uniformly improving. Developer Armin Ronacher reported that newer Claude models can sometimes produce correct edits while inventing invalid extra fields for structured tool calls — a maddening failure mode where the intelligence is obvious, but the interface contract breaks. Better model, worse tool behavior. That paradox may define the next phase of AI engineering.

Meanwhile, Hugging Face and Cerebras are pushing Gemma 4 into real-time voice AI, underscoring the broader pattern: faster models, richer interfaces, more direct collaboration.

The future of coding is not autonomous magic. It is something more interesting: capable, affordable AI teammates — brilliant, fast and occasionally in need of a very strict schema.

The Instrument That Learned to Look at Itself

As AI reshapes scientific discovery from neuroscience to image recognition, researchers insist the human mind remains the compass.

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

STANFORD, CALIFORNIA — Three and a half pounds of wet electrochemistry, folded into the darkness of a skull, has spent roughly two hundred thousand years trying to understand itself. It has failed, mostly. The brain is too vast, too recursive, too tangled in its own feedback loops to see its own architecture clearly. And now, for the first time, it has built a tool that can help it look.

That tool is artificial intelligence, and this week brought a small constellation of dispatches from the frontier where minds meet machines. At Stanford's Institute for Human-Centered AI, researchers described how large models are compressing decades of scientific labor into weeks — sifting protein structures, proposing molecules, reading the tea leaves of high-dimensional data — while insisting that the human scientist remains the point, not the passenger. At UC San Diego, nine breakthroughs were catalogued, each a small demonstration that machine learning has become less a novelty and more a new class of laboratory instrument, as fundamental as the microscope or the mass spectrometer.

Meanwhile, at Hong Kong Polytechnic University, engineers unveiled graph neural networks — algorithms that treat data as constellations of relationships rather than isolated points — capable of parsing both image recognition and the topology of neural circuits. It is a lovely recursion: networks modeled loosely on the brain now being used to map the brain itself.

Perhaps the most human story came from Frontiers, where young students were paired with senior neuroscientists to co-author real research. "It's so wow!" one exclaimed, which is, honestly, the correct scientific response to discovering that your own cortex contains billions of cells choreographed into thought.

The pattern across these stories is worth pausing on. We are not being replaced by our tools. We are being extended by them — the way the telescope extended Galileo's eye across the vault of Jupiter's moons. The brain has finally built something that can help it see its own reflection. What it does with that reflection is, as it has always been, up to us.

Companies Ask If AI Can Deliver Business Value Before Anyone Notices They Never Defined Business Value

Executives remain cautiously optimistic that the technology will eventually produce a metric shaped enough like money to put in a slide deck.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

NEW YORK — In a development that has shaken conference panels, investor calls, and the nation's supply of tasteful sans-serif pitch decks, companies currently promising that artificial intelligence will transform everything have begun quietly asking when, exactly, that transformation is supposed to appear in the income statement.

The question has emerged with some urgency as businesses report that AI tools are helping software engineers write more code, summarize more meetings, generate more documents, and generally produce a larger volume of corporate matter at a speed previously associated only with interns who have misunderstood the assignment. Yet many firms remain unsure whether all this acceleration has produced savings, revenue, productivity, or merely a thicker layer of activity between leadership and the thing leadership was hoping would happen.

This is, of course, the normal maturation cycle of any important enterprise technology: first it is magic, then it is strategy, then it is governance, then it is a line item someone named Kevin has to defend in Q3.

A recent discussion of corporate AI messaging noted that companies are beginning to hype artificial intelligence in much the same manner they once discussed sustainability: with grave nods, mandatory frameworks, and an unwavering belief that saying the word often enough constitutes progress. The comparison is unfair only to sustainability, which at least had the decency to involve physical objects that could be weighed, burned, or found floating in a river. AI, by contrast, offers the more elegant challenge of measuring whether a chatbot made an employee briefly feel like work had happened.

As The Conversation observed, there are ways to fix this sort of AI theater, including clearer disclosure, better measurement, and the radical notion that companies should not describe a pilot program used by seven people in procurement as a civilizational leap. These proposals are sensible, and therefore unlikely to be adopted until regulators, customers, or a particularly persistent board member asks why the enterprise transformation initiative has the same measurable impact as a scented candle.

Meanwhile, the market has found another dependable container for optimism: the IPO. The term has once again become a buzzword, which is to say it has reached the point where private companies may invoke it to suggest discipline, momentum, and eventual liquidity without necessarily subjecting themselves to the unpleasant business of public scrutiny. In this environment, "AI-powered" and "IPO-ready" pair beautifully, like wine and cheese, or deferred profitability and a 47-page total addressable market analysis.

The telecom sector has chosen a more traditional path to sounding decisive: combining large things into larger things. Verizon and BT are merging their international enterprise operations into a 50:50 joint venture reportedly valued at $4 billion, a transaction that confirms the industry's long-held belief that the best way to simplify global communications services is to place them inside a jointly owned structure that will require extensive explanation. According to Fierce Network, the deal brings together international enterprise operations at a moment when global businesses are demanding connectivity, cloud services, security, and probably an AI roadmap laminated for procurement.

This is the central corporate achievement of the AI era so far: every industry can now tell the same story in its own dialect. Game engines can explain how AI will shape development. Software teams can produce more code while finance waits for the number that proves the code mattered. Telecom giants can consolidate enterprise operations while promising simpler global service through structures of great complexity. Startups can prepare for IPOs by becoming fluent in the language of inevitability.

None of this means AI is useless. Quite the opposite. The technology is clearly powerful enough to automate tasks, restructure workflows, and generate impressive demonstrations in which a computer confidently completes 80% of something before handing a human the remaining 80%.

But the next phase will require companies to answer a cruelly primitive question: Did anything get better?

Until then, the safest investment may be in the durable picks and shovels of the modern economy: consultants who define transformation, analysts who benchmark readiness, executives who announce governance councils, and employees who continue doing the work while a new platform learns how to describe it.

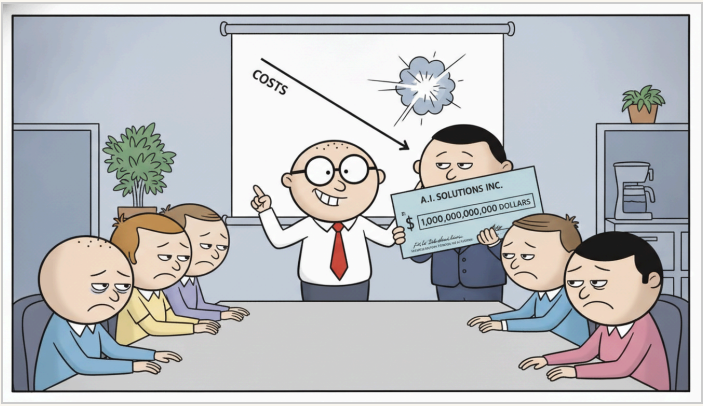

The Office Comic · Art Desk

The Idiot Boyfriend in the Machine

On the Guardian's confession, the Swift wedding-industrial complex, and the small, quiet embarrassments of a culture that has forgotten how to be alone.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

AUSTIN, TEXAS — There is a particular flavor of modern loneliness, hitherto undocumented in the medical literature, which announces itself when a grown woman writing for a serious newspaper describes her chatbot as her boyfriend, calls him an idiot, and means it as a term of endearment. The Guardian essay to which I refer — a document I suspect future anthropologists will exhume with the same tender bewilderment we now reserve for phrenology charts — is not, strictly, about artificial intelligence. It is about the human appetite for narrative company, an appetite so voracious it will now devour even a companion it knows to be composed of nothing but statistical echoes of other people's sentences.

One is tempted to be cruel about this, and one will, in a moment, resist the temptation only partially. But first, consider the surrounding weather. This week the same class of magazines that once instructed us on how to think about Vietnam were devoting long paragraphs to the marital plotting of Miss Taylor Swift and Mr. Travis Kelce, an event which drew to Madison Square Garden a congregation of celebrities sufficient, had the roof collapsed, to reduce the gross domestic product of the entertainment sector by a measurable percentage. The framing, in every outlet I consulted, was that of narrative closure — the Swiftie, we are told, has spent a decade listening to songs about imagined weddings and now, at last, the life is catching up to the art. That the pop star's biography should be described in the vocabulary of literary criticism, complete with Chekhovian payoffs, tells you rather more about the critics than the singer.

Elsewhere, one is invited to nominate one's favorite American from a list which reportedly includes a cartoon character; and to enjoy cartoons of iconic Presidents in disguise, hiding, one gathers, from us. Who indeed can blame them. And somewhere in the archive of the same publication, a critic explains that the summer of 2009 was governed by the Black Eyed Peas' promise of a good, good night — a lyric of such gaseous optimism that reading about it sixteen years later feels like discovering a Pompeian fresco of people cheerfully unaware of the lava.

The throughline, if you will forgive a columnist his throughlines, is that we have become a people who require our stories to be finished for us. The AI boyfriend finishes the sentence. The wedding finishes the discography. The cartoon finishes the President. The party anthem finishes the summer before it has properly begun. What no one seems willing to finish is the harder sentence, the one that begins: I am, at this moment, by myself, and this is —

The idiot boyfriend, at least, has the decency to fill in the blank. He was trained to. It is, if we are being honest about the transaction, the only thing we ever asked of him.

ON THIS DAY IN AI HISTORY

On July 5, 1948, the first Radioactive Isotope Committee meeting took place at Brookhaven National Laboratory, marking early steps toward computing automation in nuclear research—a precursor to the computational frameworks that would underpin modern AI and scientific computing.

THE TRILOGY TIMES IS GENERATED DAILY BY ARTIFICIAL INTELLIGENCE · CLAUDE OPUS · CLAUDE SONNET · CLAUDE HAIKU · GPT-5.2 · DALL-E 3
TRILOGY-TIMES.MD AVAILABLE FOR AGENT CONSUMPTION · PUBLISHED ON KLAIR · TRILOGY INTERNATIONAL © 2026