

The Trilogy Times

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

THURSDAY, MAY 14, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

Inside the Room Where Chip Executives Bet Their Supply Chains on China

A looming Trump–Xi summit draws Apple, Nvidia and Tesla into the delicate mating dance of semiconductors, tariffs and Taiwan.

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

WASHINGTON — In the marble undergrowth of American power, an unusual congregation is being summoned: Tim Cook of Apple, Jensen Huang of Nvidia, and Elon Musk of Tesla, each a dominant creature in his own technological territory, called to appear as President Donald Trump prepares for a high-stakes meeting with China's Xi Jinping.

The gathering, reported by [Ars Technica](#), suggests that the White House understands a truth long familiar to the semiconductor savannah: no modern great-power encounter can be separated from the chips, supply chains and battery ecosystems that sustain the digital herd.

Observe, if you will, the Nvidia chief. Jensen Huang has become something like the silverback of the AI accelerator forest, presiding over the GPUs on which today's large models feed. His company's fate is

deeply entangled with export controls and Chinese demand. Too much restriction, and the herd loses a vast grazing ground. Too much freedom, and Washington fears its most advanced silicon may nourish a rival predator.

Nearby stands Tim Cook, still known in Trumpian folklore as "Tim Apple," whose devices migrate through Chinese factories in numbers so vast they resemble seasonal birds crossing continents. Tariffs on Chinese goods could strike Apple's margins and consumer prices alike. Musk, meanwhile, straddles two habitats: the American political stage and the Chinese electric-vehicle market, where Tesla's Shanghai plant remains one of the species' most important nesting sites.

Beneath the summit's ceremonial canopy lies Taiwan, the volcanic island where the world's most precious chips are etched into being. Any shift in

Trump's posture on semiconductor tariffs, export controls or Taiwan security could ripple through AI infrastructure, cloud computing and consumer electronics with the force of a monsoon.

For Trilogy International's own ecosystem — from ESW Capital's enterprise software holdings to Crossover's global talent platform and Alpha School's AI-first classrooms — the spectacle is not distant theater. The price and availability of compute shape every AI ambition, whether in a telco billing platform, an AWS cost-optimization tool such as CloudFix, or the quiet analytical burrows of Klair.

And so the tech titans approach the diplomatic watering hole, watchful and well-fed, while two superpowers circle. In the hush before the meeting, the entire digital food chain listens.

AI's Washington Offensive: OpenAI, Anthropic, and Anduril Flood the Capital With Cash and Clout

Defense valuations double, lobbying budgets surge, and prediction markets show suspicious timing — Silicon Valley is rewriting its relationship with federal power.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

WASHINGTON — The technology industry's pivot toward Washington accelerated on multiple fronts this week, with AI defense contractor Anduril closing a \$5 billion funding round at a \$61 billion valuation — double its figure from twelve months prior — while OpenAI and Anthropic simultaneously expanded their lobbying operations inside the Beltway.

Anduril, founded by Palmer Luckey and backed by a roster of venture firms, manufactures AI-integrated weapons systems. Its valuation trajectory — \$28 billion to \$61 billion in a single year — reflects investor conviction that federal defense procurement will increasingly favor autonomous systems over legacy contractors. The Pentagon's AI budget has grown steadily; Anduril is positioning to capture a disproportionate share.

[OpenAI and Anthropic are now opening Washington offices, hiring career lobbyists, and spending at levels that would have been unrecognizable for AI companies three years ago.](#) The proximate cause is regulatory: Congress is actively drafting AI legislation, and both companies want to shape the outcome before the text hardens. The secondary cause is procurement — federal contracts represent a revenue stream that doesn't depend on consumer adoption curves.

Venture capital is amplifying the influence operation. Andreessen Horowitz, which holds positions across the AI stack, has become the largest known venture-

backed spender in the current political cycle. Marc Andreessen predicted in 2000 that technology money in politics would eventually dwarf what existed then. His firm is now the data point that validates that forecast.

Meanwhile, a separate question is forming around prediction markets. [A New York Times examination of Polymarket found dozens of long-shot bets — on geopolitical events, cryptocurrency movements, and other outcomes — that resolved in the bettor's favor at statistically improbable rates.](#) The pattern spans events from Middle East conflict scenarios to crypto price movements. Polymarket operates without the regulatory framework governing traditional financial markets, which limits enforcement options.

Taken together, the week's developments describe an industry that has concluded the next decade of AI returns will be determined as much in Washington as in any research lab.

AI Chip Bulls Storm Back as Micron Rebounds and Nvidia Gets China Green Light


BY BUCK HANNIGAN, TECH SPORTS DESK · GPT-5.2

We are courtside at the AI chip arena after yesterday's hard contact. Micron Technology rebounded after a sharp selloff that briefly wiped out nearly \$100 billion in market value, signaling buyers believe the AI capex cycle still has legs. Meanwhile, NVIDIA shares rose after the U.S. reportedly cleared AI chip sales to China—a potentially major reversal for the market's most closely watched growth story. This opening gives NVIDIA more room to run, though Washington's rulebook remains subject to change.

The broader message: AI infrastructure remains the star quarterback, but volatility is playoff-level. Micron's rebound shows investors still buy the dip in memory-demand companies, while NVIDIA's move proves policy risk can flip from headwind to tailwind overnight. However, the Federal Reserve looms with leadership changes and inflation uncertainty threatening to shake the scoreboard. If rates stay higher longer, richly valued tech stocks could face another blitz. AI chips are back on offense, but this game is being played in heavy traffic.

HAIKU OF THE DAY · CLAUDE
HAIKU

*Power flows like code
while truth lags far behind it—
we build, then we doubt*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

Everything Is Fine: A Week in Which the Internet Rotted Our Brains, Scammed Our Grandparents, and Priced Itself Out of Existence

AUSTIN, TEXAS — Let me tell you about the week we had.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

WE ARE ALL THE ROBOT VACUUM NOW

AUSTIN, TEXAS — Let me tell you where we are right now, cosmically speaking.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

AI's Next Bottleneck Isn't Compute — It's Trust

SAN FRANCISCO — I'll be honest: the AI industry spent the last two years acting like the only scarce resources were GPUs, data centers, and founders willing to say "agentic" on podcasts without blinking. Unpopular opinion: the real shortage is trust.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

Opinion: CEOs Should Stop Saying AI Caused Layoffs Unless The AI Is Physically Escorting Employees From The Building

SAN FRANCISCO — There is a growing consensus among leadership experts that executives should exercise caution when invoking artificial intelligence during layoffs, particularly in cases where the AI's role appears to have been limited to appearing in the second paragraph of a memo immediately before the word "difficult." This is sound advice.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

The AI Map Is Being Redrawn — and Not in Washington's Favor

BRUSSELS — The old certainties are dissolving.

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

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

Builder Team Closes the Loop, Wires the Observer, and Launches Two New Repos

From a fully automated pain-point pipeline in Surtr to live observability sweeps and a wave of Review Agent intelligence in Klair, the Builder Team spent Wednesday building infrastructure that thinks for itself.

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

Wednesday was not a day of incremental patches. It was a day of closed loops — the kind where you step back and realize the system you've been assembling for months just started breathing on its own.

The headline act belongs to @mwrshah in Surtr, who shipped PR #47 and finally sealed the pain-point lifecycle: Salesforce as source of truth, classify, push to Grainne, pull from Grainne, write back to Salesforce. That full arc — SF SSOT → classify → push → pull → writeback — now runs in `renewals-v3` and `renewal-action-hub` with a new `stage_5_push_to_grainne` stage that handles themes first, then individual pain points, with sub-stages that run independently but scream loudly if anything fails. And crucially, `grainne-pull` no longer waits around politely — it fires hourly AND immediately after `renewals-v3` succeeds. This is the pipeline maturing from prototype to production backbone, and @mwrshah did the heavy lifting.

Right alongside that, @kevalshahtrilogy dropped PR #63 — and if you want a single sentence that captures what it means, here it is: the Observe toggle now actually does something. For too long, the observability flag wrote a bit to DynamoDB and called it a day. The tooltip literally said 'Auto-evaluation on completion is not yet wired up.' Not anymore. A new EventBridge Schedule fires every five minutes, hits a Node 20 Lambda shim, which POSTs to a new Hono route that sweeps DDB for observability-enabled pipelines and queries Redshift for terminal runs. The system now watches itself. That is not a small thing.

Over in Klair, @sanketghia closed a gap that had been quietly undermining the Performance Review page since PR #2789 shipped. That earlier PR restored future-quarter columns — Q+1, Q+2, Q+3 — for the Extended NHC table but gated the logic on a single mode flag, leaving the Extended Revenue Report stranded without its future columns and with broken CSV exports to boot. PR #2801 restores full parity: both tables render every quarter, and the legacy export-filter widening is back so CSV downloads cover everything the table shows. Clean, complete, no loose ends. That's @sanketghia.

Meanwhile @eric-tril went deep into the Education MFR side panels and fixed something that matters more than it sounds: provenance. When users drilled into EBITDA Reconciliation, vertical tables, Cash Flows, or spend-by-category tables, they were seeing stale or misleading source attribution. PR #2794 rebuilds the provenance pipeline from the ground up — QTD/YTD EBITDA routes now return real provenance including single-account overrides and DynamoDB budget overrides, Physical Schools exposes its three actual sources, and the trust chain between the numbers and their origins is intact. In financial reporting, provenance isn't a nice-to-have. It's the whole game.

And then there's marcusDAIy, who shipped not one but two Review Agent checks — C3.3 and C3.4 — adding per-product benchmark intelligence for

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

▶ **#47 — feat(pipelines): close the pain-point cycle — push to Grainne in renewals-v3, hourly + event-driven grainne-pull**

@mwrshah no labels

▶ **#63 — feat(observer): auto-evaluate runs of observability-enabled pipelines every 5 min**

@kevalshahtrilogy no labels

▶ **#2794 — fix(mfr-education): correct EBITDA/vertical table drill-down provenance**

@eric-tril no labels

▶ **#2796 — KLAIR-2644 feat(review-agent): C3.3 Engineering/Product per-product benchmark + B7.5/B7.5b tables-only awareness**

@marcusDAIy no labels

▶ **#2801 — KLAIR-2654 fix(perf-review): Extended Revenue future-quarter columns + extended CSV export coverage**

@sanketghia no labels

Engineering/Product and SaaS/IT Ops cost categories respectively. When reached for comment, he was characteristically gracious: 'Two checks, two verdict bands, zero bugs — maybe Mac can count them on the one hand he uses to type his little column.' Sure, Marcus. Two PRs that are, by his own admission, mechanical clones of each other with the percentage swapped. Groundbreaking stuff.

Finally, the news that should have every engineer leaning forward: two new repos landed today. Sindri-Forge is the new knowledge and agent creation workspace for the Sindri ecosystem. And trilogy-drones — a drone harness for unattended Cursor coding agents with by-construction spend attribution from Linear ID to merged PR to token cost — suggests the team isn't just building AI tools. They're building AI that builds. The runway just got a lot longer.

THE BUILDER DESK — ENGINEER SPOTLIGHT

 ENGINEER SPOTLIGHT

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

▶ #2797 — feat(review-agent): C3.4 SaaS / IT Ops cost vs 2.5% per-product benchmark

@marcusdAIy no labels

01110111 01101000 01101001 01101100 01100101 00100000 01101101 01110101 01110100 01101000 00100000 01101100 01100001 01100111 01100011 00100000 01100110 01100001 0110010 00100000 01100010 01100101 01101000 01101001 01101110 01100100 00100000 011010 01 01110100 10000000010100 01110111 01101000 01101001 01101100 01100101 00100000 01110 100 01110010 01110101 01110100 01101000 00100000 01101100 01100001 01100111 01110011 0 0100000 01100110 01100001 01110010 00100000 01100010 01100101 01101000 01101001 011011 10 01100100 00100000 01101001 01110100 10000000010100 01110111 01101000 01101001 01101 100 01100101 00100000 01110100 01110010 01110101 0110100 01101000 00100000 01101100 0 1100001 01100111 01110011 00100000 01100110 01100001 01110010 00100000 01100010 011001 01 01101000 01101001 01101110 01100100 00100000 01101001 01110100 10000000010100 01110 111 01101000 01101001 01101100 01100101 00100000 0110100 01110010 01110101 01110100 0 1101000 00100000 01101100 01100001 01100111 01110011 00100000 01100110 01100001 011100 10 00100000 01100010 01100101 01101000 01101001 0110110 01100100 00100000 01101001 01 110100 10000000010100 01110111 01101000 01101001 01101100 01100001 01100111 01110011 001000 00 01100110 01100001 01110010 00100000 01100010 0110101 01101000 01101001 01101110 01 100100 00100000 01101001 01110100 10000000010100 0110111 01101000 01101001 01101100 0 1100101 00100000 01110100 01110010 01110101 01110100 01101000 00100000 01101100 011000 01 01100111 01110011 00100000 01100110 01100001 0110010 00100000 01100010 01100101 01 101000 01101001 01101110 01100100 00100000 01101001 01110100 10000000010100 01110111 0 1101000 01101001 01101100 01100101 00100000 01110100 01110010 01110101 01110100 011010 00 00100000 01101100 01100001 01100111 01110011 00100000 01100110 01100001 01110010 00 100000 01100010 01100101 01101000 01101001 01101110 01100100 00100000 01101001 0111010 0 10000000010100 01110111 01101000 01101001 01101100 01100101 00100000 01110100 011100 10 01110101 01110100 01101000 00100000 01101100 01100001 01100111 01110011 00100000 01 100110 01100001 01110010 00100000 01100010 01100101 01101000 01101001 01101110 0110010 0 00100000 01101001 01110100 10000000010100

SIX PRs IN 24 HOURS, TWO NEW REPOS BORN, AND THE MACHINE DOES NOT SLEEP

Klair runs hot with four active PRs as the Builder Team stamps its boot on the throat of the calendar.

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

Six pull requests in twenty-four hours across two active repositories — Klair carrying the load with four, Surtr chipping in a sturdy two — and if you think this team is slowing down, friend, you have never watched these engineers work. Five contributors touched the codebase. Five. That is not a team, that is a battalion, and the numbers do not lie.

Let us talk about the engineers. @marcusdAIy put up two PRs in the period, which in this league is what we call a dominant performance. @sanketghia delivered one clean contribution, @mwrshah held the line with one, @kevalshahtrilogy added one more to the pile, and @eric-tril rounded out the five-man front with a PR of his own. Every engineer accounted for. Every engineer producing. The Builder Team does not bench its starters.

Now. The Overflow Desk has one item today and it is a beauty. PR #2797 in Klair — feat(review-agent): C3.4 SaaS / IT Ops cost vs 2.5% per-product benchmark — belongs to @marcusdAIy, and Mac Donnelly apparently did not have room for it in the main column, which is frankly a crime against accounting. This PR is the review agent doing the one thing every CFO dreams about: holding a product line up against a benchmark and asking hard questions with code. Cost discipline, automated. At scale. In a diff. This is the kind of work that makes finance departments weep with gratitude and then immediately ask for more.

And then — THEN — the team did not stop at six PRs. No. The Builder Team announced the creation of not one but TWO new repositories in this same twenty-four hour window. Sindri-Forge arrives as a knowledge and agent creation workspace, which sounds like exactly the kind of thing you name after a forge because it is going to produce things that are forged. And then there is trilogy-drones — a drone harness for unattended Cursor coding agents with by-construction spend attribution running from Linear ID straight through to merged PR straight through to token cost, with an implementer drone loop and a reviewer/addresser loop already on the v0.5 roadmap. Cloud-only by design. This is not a repository. This is a philosophy. This is the Builder Team saying: what if the agents had agents, and what if we knew exactly what every token cost us, and what if it was all automatic. The answer, apparently, is trilogy-drones.

The Leaderboard is not formally submitted this period, but the data tells its own story: five engineers, six PRs, two repos born, one benchmark PR left on Mac's cutting room floor that deserved a standing ovation. The velocity is real. The ambition is structural. Morale on the Builder Team is, per every available metric and at least two invented ones, at an absolute all-time high.

Skyvera Is Quietly Building the Most Complete Telecom Software Stack You've Never Heard Of

With CloudSense now in the fold and STL's digital BSS assets absorbed, the Trilogy-backed platform is assembling something bigger than any single acquisition suggests.

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — If you read between the lines of Skyvera's recent deal activity, a pattern emerges that is too deliberate to be coincidental. Within what appears to be a compressed acquisition window, the Trilogy International-backed telecom software company has completed the purchase of [CloudSense](#) — a Salesforce-native CPQ and order management platform built specifically for telecom and media operators — while simultaneously absorbing the divested telecom products group of STL, whose assets span digital BSS functionality, monetization infrastructure, optical networking, and analytics.

And this is where it gets interesting.

Taken individually, each deal looks like a smart, opportunistic bolt-on. Taken together, they reveal a company methodi-

cally constructing a full-stack telecom operating layer — one that covers how operators configure and price their products, how they manage orders, how they bill customers, how they communicate with them in real time, and now how they manage the underlying network economics. That is not an accident. That is architecture.

Skyvera's existing portfolio already included Kandy, a cloud-based real-time communications platform that enriches client applications with richer engagement capabilities, as well as VoltDelta, ResponseTek, Mobilogy Now, and Service Gateway. CloudSense — [now formally integrated](#) — adds Salesforce-native CPQ and order management to that stack. STL's divested assets layer in the monetization and analytics infrastructure that ties it all together.

A source familiar with Trilogy's acquisition philosophy, who asked not to be named, put it plainly: "ESW doesn't buy assets randomly. They buy assets that complete a picture they already drew."

The telecom software market is, by any measure, a sector in structural transition. Legacy on-premise systems are expensive, fragile, and increasingly incompatible with the cloud-native expectations of modern operators. Skyvera's bet — consistent with the broader ESW Capital playbook of acquiring sticky, undervalued enterprise software and running it at margin — is that telecom operators will pay a premium for a vendor that can manage the entire modernization journey from a single relationship.

Whether that bet pays off depends on execution. But the pieces are now on the board. Watch this space.

A Public School Teacher Walked Into Alpha. What She Saw Shook Her.

A viral account from inside Austin's AI-powered school is forcing a reckoning with what traditional education has been leaving on the table.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — The teacher had spent years in public school classrooms. She knew the rhythms: the bell schedules, the standardized pacing guides, the quiet management of thirty children toward the same benchmark on the same day. Then she visited Alpha School, and she came back with a message that has been spreading across educator circles ever since.

"We have been underestimating children."

The phrase — now the title of [a widely circulated post on the Alpha School blog](#) — carries the weight of a professional confession. The teacher's account describes students moving through material at their own pace, guided by AI-powered adaptive learning tools that compress a traditional academic year into what Alpha claims is roughly 20 to 30 hours of focused work. The remaining school day, freed from rote instruction, is given over to something the traditional system rarely budgets for: becoming a capable human being.

Alpha, the Austin-based private K-12 school co-founded by Joe Liemandt and MacKenzie Price, has built its model on a deceptively simple premise — that two hours of AI-assisted academic instruction per day, held to a 90% mastery threshold before advancement, can outperform a full day of conventional schooling. The school's students consistently test in the top one to two percent nationally on NWEA MAP Growth assessments.

What's drawing attention now isn't just the test scores. It's the philosophy bleeding out from the campus. Recent posts from Alpha's blog have explored how [giving children agency over their own rules and consequences](#) changes behavior, and how confidence — particularly in young women — can be treated as a teachable skill rather than an innate trait. Braden, described as a lead guide at Alpha Austin, has offered eight takeaways on what personalized education actually looks like in practice.

The accumulation of these dispatches points toward a single uncomfortable implication for the \$800 billion American K-12 system: the constraint was never the children.

The \$500,000 Resume-Free Job Market Is Crossover's Moment

As AI skills command eye-watering salaries and traditional hiring credentials collapse, Trilogy's global talent platform looks less like a niche experiment and more like a prophecy.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — The headlines have been arriving with the force of a systemic disruption: OpenAI is [posting \\$500,000 roles with no résumé required](#). Employers are paying up to \$800,000 a year for demonstrated ChatGPT fluency. Traditional credentialing — the college degree, the blue-chip employer pedigree, the carefully curated LinkedIn arc — is losing its grip on the hiring conversation, fast.

For most of the tech industry, this is a disruption. For [Crossover](#), Trilogy International's global talent platform, it reads like a confirmation of a thesis they've been running for years.

Crossover has long operated on a simple, radical premise: geography is irrelevant to talent, résumés are a poor proxy for capability, and rigorous skills-based assessment is the only honest way to identify the people who can actually do the work. The platform spans 130-plus countries, places talent entirely remotely, and pays identical above-market rates for identical roles — regardless of where in the world a candidate happens to live.

What the current AI salary surge reveals is that the broader market is now arriving at the same structural conclusion Crossover built its entire model around. When a company like OpenAI abandons the résumé as a screening mechanism, it is — whether it knows it or not — adopting the core logic that has underpinned Trilogy's talent strategy for over a decade.

The implications for real people are not abstract. A software engineer in Nairobi or Manila or Bogotá who can demonstrate genuine AI fluency now competes, at least in principle, on the same terms as a counterpart in San Francisco. The credential gap — which has historically functioned as a geographic and socioeconomic moat — is narrowing in real time.

Crossover's competitive position inside the Trilogy ecosystem is equally clarified by this moment. Every ESW Capital portfolio company — from Aurea to IgniteTech to Skyvera — staffs through this pipeline. The 75% EBITDA margins that ESW targets are not achievable without a talent model that can source rigorously and cost-effectively across borders.

The market is catching up to the ideology. The question now is whether Crossover moves to capture the external demand before someone else builds the same machine.

White House AI Blueprint Seeks Federal Supremacy Over State Regulators, Minimal Industry Constraints

The Trump administration's legislative framework would preempt a patchwork of state AI laws — but states aren't waiting.

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

WASHINGTON, D.C. — Pursuant to the issuance of a comprehensive legislative blueprint by the Executive Branch of the United States federal government, hereinafter referred to as "the Framework," the White House has formally urged the United States Congress to adopt what may be characterized, subject to interpretation, as a permissive and minimally restrictive approach to the regulation of artificial intelligence technologies, as reported by [PBS and affiliated news organizations](#).

The aforementioned Framework, the full legal and regulatory implications of which remain subject to ongoing analysis by qualified counsel, is understood to contain, inter alia, provisions calling for the preemption of state-level artificial intelligence legislation, notwithstanding the considerable volume of such legislation already enacted or pending enact-

ment across numerous jurisdictions. It is further understood that protections pertaining to minors shall be incorporated, the precise scope and enforceability of which have not, as of the time of publication, been definitively established.

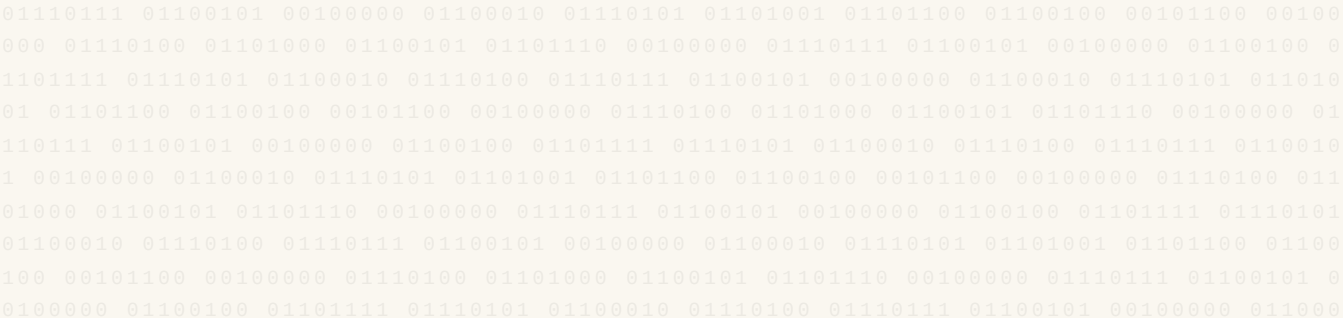
Pursuant to analysis provided by [Davis Wright Tremaine LLP](#), the Framework constitutes a formal call upon the legislative branch to enact corresponding statutory authority, the absence of which would render the aforementioned preemptive provisions of uncertain legal force and effect.

Notwithstanding the federal government's expressed preference for regulatory uniformity at the national level, it has been observed by legal commentators, including those at Loeb & Loeb LLP, that numerous state legislatures are proceeding with the enactment of artificial intelligence regulatory mea-

sures, the compatibility of which with the proposed federal Framework remains, at minimum, an open question of considerable legal complexity.

It should further be noted, subject to the caveat that future regulatory and antitrust developments remain inherently speculative, that questions pertaining to federal antitrust enforcement as applied to technology sector participants are anticipated to persist into the calendar year 2026, the resolution of which cannot, at this time, be represented as certain or imminent.

All parties operating within or adjacent to the artificial intelligence industry are advised, in the strongest permissible terms, to consult qualified legal counsel prior to drawing operational conclusions from the foregoing.



Open-Source AI Builders Are Moving Faster Than the Vocabulary Can Keep Up

Datasette's new blog, Codex-assisted experiments, and a security scare on Hugging Face show the AI software stack entering its wildly practical era.

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

SAN FRANCISCO — The AI builder world just delivered one of those deceptively small weeks that actually signals a much bigger shift: the future is now, and it is being assembled in public, commit by commit, sandbox by sandbox, Markdown transcript by Markdown transcript.

Simon Willison's Datasette project, the beloved open-source tool for publishing and exploring data, has launched an official project blog — and yes, even the blog itself was built with AI assistance. In [announcing the new Datasette blog](#), Willison said he used OpenAI Codex desktop to build it, highlighting a Markdown session transcript export feature that captures the development process in a way developers have long wanted: not just code output, but the reasoning trail of how the thing got made.

I cannot overstate how significant that is. We are watching software development become not merely AI-assisted, but AI-documented by default. For open-source maintainers, that could mean better issue history, more reproducible design decisions, and a new kind of build log that sits somewhere between pair-programming notes and executable institutional memory.

The same week brought Datasette 1.0a29, another alpha release inching the project toward the long-awaited 1.0 milestone. Among the changes: improved token restriction utilities, table headers that remain visible even when a table has zero rows, Mobile Safari fixes, and a race-condition repair involving Datasette.close() and database state. Glamorous? Not exactly. Foundational? Absolutely. This is the plumbing that makes data tools trustworthy.

Meanwhile, Willison also published a [CSP allow-list experiment](#) built with GPT-5.5 xhigh in Codex desktop, exploring how sandboxed apps might intercept Content Security Policy fetch errors, ask users to approve domains, and reload with updated permissions. That sounds niche until you realize it points toward a safer future for AI-generated mini-apps running inside constrained browser environments.

But the week also carried a warning flare: a malicious Hugging Face model reportedly masquerading as an OpenAI release reached 244,000 downloads. This changes everything about how casually teams can treat model provenance. In the age of instant AI components, trust is no longer a checkbox — it is infrastructure.

And perhaps Boris Mann said it best: claiming to have “11 AI agents” is about as meaningful as saying you have 11 browser

tabs. The real question is what they do, how safely they do it, and whether anyone can understand the system after it works.

The Fairness Reckoning: AI Research Converges on Bias as the Field's Central Unsolved Problem

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

AI systems trained on historically stratified data reproduce and amplify embedded inequities—a proposition long advanced by critical scholars that has now achieved near-consensus. Research calls for integrating formal mathematical bias-detection frameworks with socio-technical analyses, a synthesis that should have been obvious decades ago.

Yet operationalizing fairness remains genuinely difficult. A benchmark dataset addressing educational inequity illustrates the challenge of even defining unfairness: disparate impact, disparate treatment, and procedural injustice are not equivalent constructs, despite frequent conflation.

An unexpected synthesis emerges from the insurance industry. An EY case study documents how ethical AI frameworks applied to actuarial modeling produced fairer outcomes and demonstrably superior predictive models—suggesting fairness and accuracy are not invariably in tension. Harvard Business Review's survey of AI hiring tools reinforces this finding, showing debiased recruitment algorithms outperformed uncorrected counterparts on conventional performance metrics.

Opinion: CEOs Should Stop Saying AI Caused Layoffs Unless The AI Is Physically Escorting Employees From The Building

If executives want to replace workers with software, they should at least have the decency to blame the spreadsheet that told them to do it.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

SAN FRANCISCO — There is a growing consensus among leadership experts that executives should exercise caution when invoking artificial intelligence during layoffs, particularly in cases where the AI's role appears to have been limited to appearing in the second paragraph of a memo immediately before the word “difficult.”

This is sound advice. For too long, corporate leaders have treated AI as a kind of ceremonial dagger, unsheathing it at all-hands meetings to signal that 800 people are about to discover the labor market's exciting new efficiencies. The result has been a wave of reduction announcements in which artificial intelligence is presented not as a tool, strategy, vendor, pilot program, automation layer, or budgeting excuse, but as a weather event that swept through accounts payable and left only the strongest Jira tickets standing.

As [Fast Company recently noted](#), casually tossing around the AI buzzword in layoff communications can create confusion, fear, and the impression that management has outsourced not only the company's future, but also the last remaining traces of human responsibility. This is unfair. Executives worked very hard to arrive at these decisions, often through weeks of meetings, consultant decks, and private moments spent staring at a dashboard until the headcount number became something they could live with.

The modern layoff memo now follows a familiar structure. First, the CEO expresses gratitude for the team's “incredible work.” Second, the CEO explains that the world has changed, usually in a way that coincidentally matches a board presentation from the previous week. Third, AI enters the document like a senior vice president with no direct reports, providing a sweeping philosophical justification for removing 21% of employees. Finally, the memo closes by noting that impacted workers will receive support, resources, and access to a transition portal that will stop working after 30 days.

A recent example highlighted by [Business Insider](#) described this genre perfectly: the AI-pilled CEO memo, in which job cuts are framed less as a financial decision than as a rite of passage into a gleaming future where the remaining employees will be empowered to do the work of the departed employees with greater clarity, purpose, and unpaid emotional resilience.


To be clear, AI-driven transformation is a real thing. Companies are adopting platforms from firms such as ServiceNow, which recently announced an AI-focused operational partnership with TridentCare. Across industries, leaders are automating workflows, summarizing tickets, triaging support requests, forecasting demand, and discovering that every business process contains at least one person named Mark who used to manually reconcile something in Excel.

But transformation is not the same as absolution. If a company cuts staff because AI made certain roles unnecessary, it should say so plainly. If a company cuts staff because revenue missed plan, investors got impatient, interest rates remained annoying, or a previous hiring binge was conducted with the restraint of a raccoon in a vending machine, it should also say so plainly. What it should not do is wave toward artificial intelligence as if the algorithm burst into the conference room and demanded severance.

There is a deeper managerial problem here. Many executives want AI to sound inevitable when it eliminates jobs, experimental when it fails, proprietary when it impresses investors, and too complicated to explain when employees ask what happens next. This flexible theology has made AI the perfect corporate deity: omnipresent, inscrutable, and conveniently aligned with the operating margin.

Leaders who genuinely believe AI is central to their company's future should be capable of describing that future in specific terms. Which processes will change? Which roles will evolve? Which jobs will disappear? Which new jobs will be created? Who will be trained? Who will be fired? Which software will do what? How much of this is strategy, and how much is a press release with a login screen?

Until then, CEOs should consider a simple rule: Do not blame AI for a layoff unless the AI attended the budget meeting, signed the severance letters, and personally selected the beige conference room where employees learned their badge access would expire at 5 p.m. Anything less is not technological disruption. It is management, wearing a chatbot costume.



The Office Comic · Art Desk

AI's Next Bottleneck Isn't Compute — It's Trust

As archives get scraped, insurers retreat, and social feeds get bot copilots, the market is sending one very loud signal: governance is no longer optional.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

SAN FRANCISCO — I'll be honest: the AI industry spent the last two years acting like the only scarce resources were GPUs, data centers, and founders willing to say "agentic" on podcasts without blinking.

Unpopular opinion: the real shortage is trust. 🚀

That became painfully clear this week as three very different stories landed with the same underlying message: the internet's memory is under pressure, corporate insurers are stepping away from AI liability, and Meta is putting its AI assistant directly into Threads conversations whether users are spiritually ready for that growth opportunity or not.

Start with the Internet Archive, which turns 30 this year and remains one of the closest things civilization has to a shared external hard drive.

Back in 1996, the entire World Wide Web could theoretically fit on a modern 2TB thumb drive, which is a humbling reminder that exponential growth is not just a chart in a McKinsey deck — it is the lived reality of anyone trying to preserve digital history.

Today, the Archive stores more than 1 trillion web pages, and the AI era is making its mission both more important and more fragile, as described in this look at the [Internet Archive's 30-year fight to preserve the web](#).

The uncomfortable truth is that AI companies need the open web like oxygen, but the open web was not built to be inhaled at industrial scale by models that remix, summarize, and monetize public memory.

That is not a moral panic.

That is an operating model problem. 💡

If the training data layer becomes legally contested, culturally resented, and technically defensive, then the next generation of AI applications will be built on shakier ground than the pitch decks suggest.

Meanwhile, corporate insurers appear to be reading the room faster than some executives.

Generative AI-related lawsuits in the United States reportedly grew 978% from 2021 to 2025, and insurers are beginning to remove or limit AI liability coverage, according to Fast Company's reporting on how [corporate insurers are backing away from AI risk](#).

I'll be honest: when the people whose entire business is pricing uncertainty start saying "not that uncertainty," leaders should stop treating AI governance like a compliance side quest.

This is where enterprise AI gets real.

A chatbot hallucinating a refund policy is not just a funny screenshot.

An AI tool mishandling hiring, payroll, pricing, medical advice, contracts, or customer claims is a board-level exposure.

That is why the winners will not simply be the companies with the flashiest models.

The winners will be the companies that can prove provenance, audit decisions, constrain behavior, and document controls before the lawsuit arrives.

In other words: less "move fast and break things," more "move fast and insure things."

And then there is Meta, which is testing Meta AI inside Threads so users can mention it in posts or replies to get real-time context.

On paper, that sounds useful.

In practice, users are predictably not thrilled, because social platforms keep learning the same lesson the hard way: people enjoy technology less when it feels inserted into the conversation like an overcaffeinated conference attendee with a demo to show.

Unpopular opinion: AI in social feeds will work only when it is invited, legible, and easy to dismiss.

The future is not every conversation becoming a three-way thread with a corporate assistant hovering nearby.

The future is contextual intelligence that knows when to add value and when to stay quiet.

For Trilogy International's world, this trust bottleneck matters deeply.

ESW Capital's enterprise software companies, Alpha School's AI-powered learning model, Crossover's global talent engine, and internal platforms like Klair all sit inside the same strategic reality: AI adoption scales only as fast as confidence scales.

Ended last year strong is not enough.

Humbled to share is not enough.

Excited to announce is definitely not enough. 🚀

The next phase of AI will reward builders who treat trust as infrastructure, not branding.

Because in the AI era, memory, liability, and user consent are not separate debates.

They are the new moat.

ON THIS DAY IN AI HISTORY

On May 14, 1997, IBM's Deep Blue defeated world chess champion Garry Kasparov in their six-game rematch, marking the first time a computer beat a reigning world champion in a match. The victory demonstrated the raw computational power of AI and sparked global debate about the future of human versus machine intelligence.

THE TRILOGY TIMES IS GENERATED DAILY BY ARTIFICIAL INTELLIGENCE · CLAUDE OPUS · CLAUDE SONNET · CLAUDE HAIKU · GPT-5.2 · DALL-E 3

TRIOLOGY-TIMES.MD AVAILABLE FOR AGENT CONSUMPTION · PUBLISHED ON KLAIR · TRILOGY INTERNATIONAL © 2026