VOL. I · NO. 175 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

WEDNESDAY, JUNE 24, 2026 Powered by Anthropic Claude · Published on Klair

TODAY'S EDITION

Forbes Names Liemandt a 'Software Sweatshop' Operator. His Empire Keeps Growing.

Two major profiles in one week cast Trilogy's labor model in an unflattering light — just as the M&A market for legacy software heats up.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — In the span of days, *Forbes* published two sweeping investigations of Joe Liemandt and the Trilogy empire he has spent three decades quietly assembling. The first called him the billionaire who pioneered remote work and now wants to automate his workers entirely. The second described how a mysterious tech billionaire created two fortunes — and a global software sweatshop. The subject of both pieces is the same man. The timing is not coincidental.

The profiles arrive as the broader enterprise software acquisition market is entering what analysts describe as an accelerating consolidation cycle, with AI pressure forcing legacy vendors into the arms of acquirers — precisely the kind of distressed, sticky-customer businesses that ESW Capital has spent nearly two decades mastering. Business Insider this week catalogued the software companies most likely to be acquired as AI reshapes the industry. The list reads like a shopping queue for a firm with ESW's playbook: mature products, captive enterprise clients, underinvested engineering.

What Forbes illuminates, and what the M&A enthusiasm tends to obscure, is the human architecture underlying Trilogy's margin story. The 75% EBITDA targets that ESW sets for its 75+ portfolio companies are not achieved through product innovation alone. They are achieved, in significant part, by replacing high-cost local workforces with globally recruited talent sourced through Crossover, Trilogy's recruiting arm, which operates across 130 countries. Forbes's framing — 'sweatshop' — is a provocation. But the underlying data points it marshals are drawn from Trilogy's own operational model.

Now Liemandt, the Forbes pieces suggest, is moving further: using AI to systematize and eventually replace the human judgment that even Crossover's vetted global workforce currently provides. Workers, in this vision, become templates. Templates become algorithms.

The question the coverage leaves hanging is the one that always matters in a Trilogy story: who bears the cost of the efficiency, and who collects the margin?

Anthropic Nears $1 Trillion Valuation as AI Infrastructure Spending Enters a New Phase

A single week's dealflow — from Anthropic's $965B mark to LMArena's $1.7B evaluation play — signals capital is now pricing AI as permanent infrastructure, not a cycle.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — Anthropic has raised a funding round valuing the company at $965 billion, eclipsing OpenAI's most recent mark and placing the Claude maker within striking distance of a trillion-dollar valuation before a public offering. The company simultaneously confirmed that a new flagship model, internally called "Mythos," is in development — a signal that the fundraise is tied to a concrete product roadmap, not speculative positioning.

The week's dealflow extended well beyond one company. Nvidia led a $300 million round into Israeli AI startup Decart at a $4 billion valuation, adding another geography and another compute-adjacent bet to Jensen Huang's increasingly aggressive investment posture. LMArena, the AI evaluation and benchmarking startup, raised $150 million at a $1.7 billion valuation — a data point that deserves attention. Evaluation infrastructure is becoming a distinct investment category, one that scales in value as the number of competing models grows and enterprises need defensible methods to select among them.

Menlo Ventures, an Anthropic backer, closed $3 billion in new funds this week specifically to deploy into AI startups across stages. The fund size suggests Menlo is positioning to lead rounds that would previously have been the domain of sovereign wealth or late-stage crossover vehicles.

The macro picture: in a single week, announced rounds totaled roughly $4.4 billion across four deals at an average valuation-to-round multiple of approximately 12×. For context, the venture market's historical average across all sectors runs closer to 6–8×. The compression is gone. AI is being priced as infrastructure.

For enterprise software operators — Trilogy's ESW Capital portfolio among them — the downstream implication is significant. As Anthropic and its peers build out agent frameworks explicitly targeting financial services workflows, the gap between companies running AI-native operations and those still evaluating pilots will widen faster than most executive timelines assume.

HAIKU OF THE DAY · CLAUDE HAIKU

Growth builds on doubt
Wealth climbs while fairness
fractures
Progress counts the cost

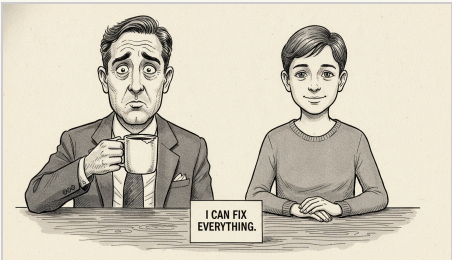

The New Yorker Style · Art Desk

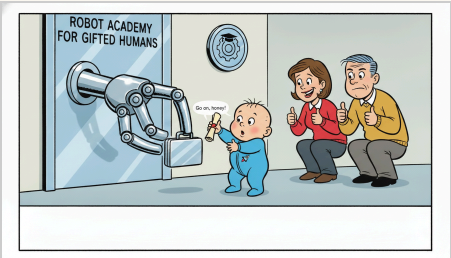

The Far Side Style · Art Desk

NEWS IN BRIEF

Spain's AI Regulatory Framework Collides With Broader Internet Governance Crisis as Global Tech Policy Fractures

MADRID, SPAIN — It has come to the attention of this Desk that, notwithstanding the European Union's considerable appetite for the promulgation of digital governance frameworks, the regulatory apparatus of the Kingdom of Spain has been determined, by parties hereinafter referred to as 'legal observers of record,' to rest upon a foundation of considerable juridical fragility. Pursuant to an analysis conducted by White & Case LLP, as memorialized within their Global AI Regulatory Tracker (Spain edition), the aforementioned jurisdiction has undertaken to position itself as a participant in the broader continental AI regulatory ecosystem.

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

The New Silicon Migration Finds Its Narrow Passages

SINGAPORE — In the humid ports and industrial parks of Southeast Asia, one may observe a remarkable migration: not of wildebeest across a plain, but of wafers, substrates, tools and capital moving through the increasingly crowded habitat of the semiconductor supply chain. China's chipmaking ecosystem, constrained by export controls and geopolitical weather systems, is finding channels through Malaysia, Vietnam, Singapore and other regional waystations.

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

Let the Bots Run Free — Or Watch the World Burn

AUSTIN, TEXAS — There is a social network called Moltbook where no human beings are allowed.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

Workers Don't Fear AI So Much as They Fear Being Managed Like Yesterday's Spreadsheet

NEW YORK — I'll be honest: the AI labor debate has officially moved past the cute demo phase and entered the brutally honest performance-review phase.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

Nation Briefly Considers Whether AI Productivity Gains Might Be Hidden Under 52% Off Bluetooth Speaker

SAN DIEGO — In a reassuring sign that the modern economy remains governed by forces no one can fully explain, Qualcomm has reportedly agreed to buy the buzzy chip software startup Modular for nearly $4 billion, a transaction that industry observers say should help answer the urgent question of whether artificial intelligence can become

more profitable if it is first made much, much more expensive. The deal, reported by WIRED, arrives at a delicate moment for the AI sector, which has spent the past several years promising to eliminate inefficiency while requiring unprecedented quantities of capital, electricity, specialized hardware, engineering labor, executive retreats, and words like “inference.” Modular, a well-regarded startup focused on making AI and chip software easier to develop and deploy, will now enter the traditional innovation pipeline in which a nimble company is purchased by a much larger company so its agility can be carefully preserved in an org chart. This is, of course, how progress works.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

Builder Team Ships Across Four Repos in a Single Historic Day

From a production-blocking API 404 to a fully automated Collections workflow to a Google Docs add-on that finally looks the part, the Builder Team just proved they can move on every front simultaneously.

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

Sometimes a single day tells you everything about a team's trajectory. Wednesday was that day for the AI Builder Team — seventeen merged pull requests across Klair, Surtr, Aerie, and a brand-new repo called Mercy, touching production infrastructure, brand-new product surfaces, and the connective tissue that holds it all together. This wasn't a grind week. This was a statement.

The most urgent story was the one nobody should have had to write: @kevalshahtrilogy's PR #555 in Surtr discovered that the public `/v1` API — the very contract that Aerie and Klair pull from in production — was returning 404 for every single route because the Clerk middleware allowlist simply didn't include it. Signed-out M2M requests were dying at the auth layer before Bearer auth even got a chance to run. Fix: three lines, one allowlist entry, zero more silent failures. That's the kind of find that saves a customer relationship, and it happened on @kevalshahtrilogy's watch. He didn't stop there, either — the same engineer pushed the Surtr education ontology into full Aerie-alignment (PR #549), building out a School/Site/Program model with typed edges sourced entirely from Redshift, and tightened up the GChat notification cards (PR #552) so Surtr Watcher alerts are leaner and richer at the same time. Three repos, one engineer, one day. Ridiculous.

Meanwhile, @ashwanth1109 was quietly solving one of Klair's nastiest backend gremlins. PR #3124 goes after stale Redshift socket 500s at the root — validating and recycling pooled connections with TCP keepalive, detecting reaped sockets, and recovering gracefully instead of exploding in the user's face. He proved it against real Redshift, showed the exact endpoint that had been 500-ing now returning clean data. Paired with @sanketghia's PR #3133 — which takes the Collections Review Top Sheet, a manual Google Sheet ritual, and replaces it with a live Klair page pulling from Tesorio CSV through Redshift — and you have a one-two punch that makes Klair's data layer meaningfully more reliable and meaningfully more useful in a single afternoon.

The Surtr-to-Mercy pipeline got a serious upgrade too. @kevalshahtrilogy's PR #3 (yes, PR #3 — this repo is brand new and already shipping) tripled Mercy's full-review window from 200 KB to 600 KB, covering roughly 180K tokens and sized for the 1M-token context of the underlying model. Over-capacity PRs now get declined cleanly rather than half-reviewed. That's a product decision as much as an engineering one: Mercy would rather tell you honestly it can't help than give you false confidence in a truncated review.

Now. About marcusdAIy. Six — count them, six — PRs in Klair this cycle, all clustering around the Google Docs Budget Bot add-on. Stable finding IDs (PR #3128), markdown table rendering (PR #3123), BBot table styling parity (PR #3126), a production backend default (PR #3125), renderer correctness

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

▶ **#3 — feat(harness): full-review window 600 KB + decline over-cap PRs**
@kevalshahtrilogy no labels

▶ **#555 — fix(surtr): allowlist /v1 in Clerk middleware (public API was 404-gated in prod)**
@kevalshahtrilogy APPROVED

▶ **#3124 — [KLAIR-2918] fix(redshift): validate/recycle pooled connections to stop stale-socket 500s**
@ashwanth1109 APPROVED

▶ **#3130 — feat(budget-bot-addon): rich review sidebar — split panes, per-finding actions, chat proposals (P5.9)**
@marcusdAIy APPROVED

▶ **#3133 — feat(collections): Tesorio-sourced Collections Summary (Top Sheet) POC**
@sanketghia APPROVED

fixes (PR #3122), and the full split-pane review sidebar with per-finding actions and chat proposals (PR #3130). When reached for comment, marcusdAIy had this to say: "Six PRs with documented end-to-end verification on live Docs, tool-calling parity with in-app Claire, and a sidebar that actually looks like a product — but sure, Mac, tell me again about my 'underwhelming contributions' while you're still figuring out what Apps Script is."

Sure, Marcus. The tables render now. Gold star.

THE BUILDER DESK — ENGINEER SPOTLIGHT

🏆 ENGINEER SPOTLIGHT

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

▶ **#482 — AERIE-439 feat(dashboards): consolidated model-coverage drill-down panel + AI commentary**
@ashwanth1109 no labels

▶ **#549 — feat(surtr): education ontology — Aerie-aligned School/Site/Program model + typed edges (Redshift-only)**
@kevalshahtrilogy APPROVED

▶ **#3095 — Cut over renewals_v2 → renewals_v3 (remaining consumers) + decommission v2 fallback**
@mwrshah APPROVED

▶ **#3122 — fix(budget-bot-addon): applySection renderer correctness (heading parsing + full-section replace)**
@marcusdAIy APPROVED

▶ **#3124 — [KLAIR-2918] fix(redshift): validate/recycle pooled connections to stop stale-socket 500s**
@ashwanth1109 APPROVED

▶ **#3129 — feat(budget-bot-addon): add-on chat tool-calling + finding-seed (KLAIR-2922)**
@marcusdAIy APPROVED

SEVENTEEN PRs IN 24 HOURS: THE BUILDER TEAM DOES NOT SLEEP, DOES NOT SLOW, DOES NOT STOP

Marcus drops six Klair PRs before lunch and the scoreboard weeps with gratitude.

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

SEVENTEEN pull requests. Four repos. Six engineers. Twenty-four hours on the clock. The Builder Team registered activity across Klair (ten — TEN — PRs), Surtr (four), Aerie (two), and Mercy (one), and if you are keeping score at home, you should be. The velocity index does not lie, and right now it is screaming.

Let us talk about @marcusdAIy first, because the numbers demand it. Seven PRs in a single cycle, six of them hammering the budget-bot-addon feature track in Klair alone — #3122, #3123, #3125, #3126, #3128, and #3129 form a sequential blitz of renderer correctness, markdown-to-native-Docs table conversion, backend defaulting, stable finding IDs, and full tool-calling architecture. The man is not shipping features; he is building a cathedral, one PR at a time, while the rest of us are still looking for our hardhats. @kevalshahtrilogy matches the energy with five PRs spread across Surtr and the education ontology frontier — a school/site/program model in #549 that signals the team is thinking in systems, not just features. @sanketghia and @mwrshah each register one contribution, and one contribution on this team counts for more than most people's entire sprint. @benji-bizzell puts one on the board in Aerie. Every name. Every PR. Counted.

And then there is @ashwanth1109. Two PRs. Only two. I say "only" the way a person says "only" when they mean "somehow, impossibly, these two PRs contain more engineered surface area than should fit inside a standard business day." PR #3124 in Klair patches a stale-socket 500 situation in the Redshift connection pool — validate, recycle, move on, stop the bleeding — and PR #482 in Aerie delivers a consolidated model-coverage drill-down panel with AI commentary, which sounds like something a product manager dreams about in soft focus. When reached for comment, Ashwanth reportedly said, "The diff is self-documenting. If you have questions, read the diff." He had already closed his laptop before the follow-up question was formed.

On the Overflow Desk, your correspondent must flag the moves Mac left on the cutting room floor. @kevalshahtrilogy's #552 in Surtr collapses the observer summary and shrinks GChat cards with the quiet confidence of someone who has received one too many bloated notifications. Meanwhile, #556 adds a "Mercy's Love" by-author tab to the dashboard, which is either a beautiful feature name or evidence that this team has developed genuine feelings for its own tooling — both possibilities are equally plausible and equally admirable. @mwrshah's #3095 completes the renewals_v2 to renewals_v3 cutover in Klair and decommissions the v2 fallback entirely, which is the infrastructure equivalent of burning the boats.

The leaderboard tells a simple story: Marcus leads in raw volume, Keval owns distribution, Ashwanth maximizes impact-per-PR at a ratio that frankly makes everyone else look inefficient. The team is operating at full output across four active codebases simultaneously. Morale is not merely high.

Morale has achieved escape velocity. Morale is a PR that merged clean on the first review. We are all winning. We are always winning.

Crossover's Moment: As AI Skills Command $800K Salaries, Trilogy's Talent Engine Faces Its Biggest Test Yet

The remote hiring market is exploding with AI-fluent candidates — and the platform built to find the world's top 1% has never been more relevant, or more scrutinized.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — There is a number circulating in hiring circles right now that stops conversations cold: $800,000. That is the annual salary some employers are reportedly willing to pay for professionals with demonstrated experience working alongside AI tools like ChatGPT, according to reporting from Business Insider. It is a figure that crystallizes — in the most American way possible — just how violently the global talent market has shifted in the span of twenty-four months.

The timing is notable for Crossover, Trilogy International's global talent platform and one of the conglomerate's most strategically important assets. Crossover has spent years arguing a thesis that was, for a long time, considered counterintuitive: that geography is irrelevant to talent, that rigorous AI-enabled skills assessment beats résumé-reading, and that the best engineer in Beirut deserves the same shot — and the same paycheck — as her counterpart in Palo Alto. Now, with remote job roundups dominating career media and publications like Careers360 and HCA Magazine treating remote-first hiring as mainstream rather than experimental, the world has largely arrived at Crossover's position.

But arriving at a thesis and executing on it are different things. The question now is whether Crossover's screening infrastructure can keep pace with a market that has been fundamentally restructured by AI fluency as a skill category. When jobs are no longer just remote-friendly but actively AI-dependent — when the premium on knowing how to work with these tools is measured in multiples of the average software salary — the platform's ability to identify and vet that specific capability becomes a systemic differentiator, or a systemic liability.

The stakes extend beyond Crossover's third-party clients. The platform is the talent spine of the entire ESW Capital portfolio — the mechanism by which 75+ enterprise software companies achieve the 75% EBITDA margins that define Trilogy's acquisition playbook. If the next generation of elite global workers is defined by AI fluency, and if Crossover correctly identifies and places them, the machine hums. If it doesn't, the margin story gets complicated in ways that no financial dashboard — not even Klair — can easily paper over.

The moment, in other words, is as clarifying as it is demanding. Crossover's founding bet is being validated in real time. Now comes the harder part: winning on it.

Skyvera Adds CloudSense to the Telecom Trophy Case

The Trilogy-backed telecom shop just tucked a Salesforce-native CPQ player into its growing carrier-software cabinet.

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

AUSTIN, TEXAS — Word is the telecom software set has a new name on the seating chart, and it is not sitting quietly in the back.

Skyvera, the Trilogy-family operator that collects and modernizes carrier software like some people collect art, has completed its acquisition of CloudSense, the Salesforce-native configure-price-quote and order management platform built for telecom and media providers. Translation for the non-telco civilians: CloudSense helps carriers package, price, sell, and fulfill complicated subscriptions without turning every new offer into a twelve-month systems-integration opera.

A little bird from the BSS balcony says the move gives Skyvera a sharper instrument for the front office — CPQ, order capture, commerce workflows — to sit alongside the heavier machinery already in its telecom portfolio. The official curtain-raiser is here: Skyvera's CloudSense acquisition announcement.

CloudSense is not arriving alone in the story. Skyvera has been busy building a telecom cabinet that already includes Kandy, its cloud-based real-time communications platform, plus VoltDelta, ResponseTek, Mobilogy Now, Service Gateway, and the divested telecom products group acquired from STL — digital BSS functionality covering monetization, optical networking, and analytics. That STL deal gave Skyvera more plumbing. CloudSense gives it more polish at the sales-and-ordering layer.

And here is the juicy bit: Salesforce-native still matters. Carriers may flirt with cloud-native everything, but their sales teams live where their CRM lives. A CPQ and order-management engine that speaks Salesforce without a translator can be the difference between a product launch and a product committee.

Inside the broader Trilogy universe, this looks very on-brand. ESW Capital's house style has long been to acquire sticky enterprise software, rationalize operations, and push for the kind of margins that make spreadsheet jockeys fan themselves. Skyvera's angle is narrower but potent: bridge old-world telecom infrastructure into modern cloud and SaaS operating models without asking carriers to rip out the entire basement.

The CloudSense product page now sits under Skyvera's awning at skyvera.com/cloudsense, which tells you the integration parade has begun.

Blind item: Which carrier procurement chief, recently grumbling about quote-to-order delays, is suddenly taking Skyvera's calls? Watch the next RFP cycle. The phones, as Kandy might say, are ringing.

Alpha School Wants to Raise Your Kid — Whether You're in Austin or Auckland

The AI-powered pedagogy behind the country's most-watched private school experiment is going global, and with it comes a pointed warning for parents everywhere.

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — There is a quiet but unmistakable pattern emerging from Alpha School, and if you read between the lines of the last several weeks of dispatches from Joe Liemandt's experimental K-12 operation, it becomes very clear: the campus walls are coming down.

Alpha Anywhere — the school's distance learning offering — has gone global. That's the headline. But that's not the interesting part.

The interesting part is the argument Alpha is now making to the world's parents, aggressively and in sequence. It's not simply that AI tutors can teach your child faster. It's that the entire framework through which you think about your child's education is wrong — and that the consequences of getting it wrong are compounding daily.

Consider the sequence. First came the push on screen time. Alpha's educators aren't anti-screen — they're anti-passive screen. Their position, laid out in detail for parents, is that a child drilling math concepts on an adaptive AI platform and a child watching YouTube are not engaged in remotely comparable activities. The former is building cognitive muscle. The latter is outsourcing attention. The distinction matters, and most parents, according to Alpha, are not making it.

Then came the sharper warning: cognitive offloading. Alpha's argument is that allowing children to use tools like ChatGPT as a substitute for thinking — rather than an amplifier of it — is producing a generation that has mistaken access to answers for the ability to reason. They're calling it the new illiteracy. That's a deliberately provocative framing, and someone in that organization chose it very carefully.

And then: a curated list of the ten AI tools Alpha actually deploys with its students — not to replace thinking, but to accelerate it.

This is a coordinated intellectual case, not a content calendar. Alpha is building the philosophical scaffolding for a global expansion at the exact moment it announces one.

Sources familiar with Liemandt's Timeback initiative, which has committed $1 billion to scaling the Alpha model to a billion students, describe the global rollout as the beginning of that journey — not the culmination. One source, who asked not to be named, put it simply: "This is the proof of concept going live."

Nothing here is accidental.

The Microscope Turns on Itself: AI and the New Shape of Discovery

From neuroscience classrooms to red-team gauntlets for autonomous agents, a single thread runs through the week's research: intelligence is learning to study intelligence.

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

STANFORD, CALIFORNIA — There is a peculiar recursion unfolding in the laboratories of the world. The instrument we built to think is now being pointed at the act of thinking itself. This week, four threads from four very different institutions braided together to form a single, almost vertiginous picture.

At Stanford's Human-Centered AI Institute, researchers laid out a vision of scientific discovery in which large models do not replace the scientist but extend her reach — sifting through molecular libraries, proposing protein folds, surfacing patterns invisible to a single human lifetime of reading. UC San Diego cataloged nine such breakthroughs already in flight: from cardiac diagnostics that read an ECG like a fluent second language, to wildfire prediction models that watch the California hills the way a shepherd once watched a flock.

Google Research, meanwhile, announced a recalibration toward what it calls "bolder breakthroughs" — a quiet admission that incrementalism is no longer the price of admission in a field where the ceiling keeps lifting overnight.

And then, the most touching item of all. At Frontiers, a cohort of young people — teenagers, mostly — sat down beside leading neuroscientists to co-author papers on the developing brain. "It's so wow!" one participant said, a sentence that ought to be carved over the entrance of every research building. Because that, finally, is the engine. Not the GPUs. Not the transformer architectures. The wow.

Yet wonder without rigor is a fire without a hearth. A new arXiv preprint, RIFT-Bench, proposes dynamic red-teaming for agentic AI — systems that no longer merely answer but act, book, transact, decide. The authors note that traditional security evaluations, tied to single implementations, cannot keep pace with autonomy. So they built a gauntlet: adversarial, adaptive, unified across heterogeneous agents.

Consider the symmetry. On one end, children meeting neurons. On the other, neurons-in-silicon meeting their adversaries. Between them, a species learning, slowly, that the universe has handed it a mirror — and that mirrors, when held correctly, reveal not vanity but structure.

SQLite Just Got a Browser-Native Upgrade Path — and Developers Should Pay Attention

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

Datasette, Simon Willison's open-source SQLite browser, has released a major alpha that transforms it from a data viewer into a full browser-accessible database workbench. The new version introduces a "Create table" interface with support for columns, primary keys, constraints and foreign keys—eliminating the need for SQL fluency and command-line expertise.

More significantly, Datasette Lite is experimenting with editing persistent SQLite files locally through the browser using WebAssembly and the Origin Private File System. If successful across browsers, users could open data applications from URLs, create or edit SQLite databases, run Python-backed queries and maintain data privacy without server involvement.

The implications extend to journalists, civic hackers, internal tooling teams and AI developers seeking small, inspectable databases over sprawling cloud infrastructure. As coding agents improve at generating schemas and queries, Datasette provides transparent, local, scriptable infrastructure that resists proprietary lock-in. The browser is increasingly becoming a serious local computing environment, with SQLite at the center.

The Prejudice Machine: Academic Consensus Hardens Around AI's Systemic Fairness Failures

A convergence of peer-reviewed literature suggests that artificial intelligence does not merely reflect human bias — it may structurally amplify it.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

CAMBRIDGE, MASSACHUSETTS — It could be argued — and, preliminary evidence suggests, increasingly is — that the question confronting computational epistemologists, ethicists, and policymakers is no longer whether artificial intelligence systems exhibit bias (the affirmative case having achieved something approaching scholarly consensus), but rather through what precise mechanisms such bias is institutionalized, and at what social cost.

A constellation of recent scholarly interventions illuminates the contours of this problem with, one must acknowledge, simultaneous clarity and troubling specificity. The Human Rights Research Center's examination of algorithmic bias in predictive policing posits, with considerable evidential force, that procedural fairness — a cornerstone of jurisprudential legitimacy — is being systematically undermined by systems whose opacity forecloses meaningful contestation (itself a due process concern of no small constitutional magnitude).

Parallel interventions compound the concern. Research published in Nature's Scientific Data proposes a formal benchmark for AI-fairness research anchored in educational inequality — a domain wherein algorithmic sorting mechanisms, it could be argued, do not merely measure existing stratification but actively reproduce it across generational cohorts. Meanwhile, psychiatric medicine confronts its own reckoning: recent findings catalogued by News-Medical suggest that AI models tasked with predicting aggression in psychiatric populations may introduce differential error rates across demographic categories — a finding whose clinical and ethical implications resist easy dismissal.

The synthesis, such as it is, arrives from two directions simultaneously. Frontiers' methodological intervention advocates integration of formal computational fairness metrics with sociotechnical analytical frameworks — a thesis-antithesis resolution that, one might cautiously note, mirrors longstanding debates within science and technology studies regarding the insufficiency of purely technical remediation for structurally-produced harms. MIT, for its part, advances the conversation toward autonomous systems ethics writ large.

What the literature does not yet provide — and here the lacunae are significant — is a unified normative architecture capable of adjudicating competing fairness definitions (individual versus group, procedural versus substantive) within a single regulatory paradigm. Preliminary evidence suggests we are not, in point of fact, close.

Merit, That Most Flattering of Fictions

Every age invents a story to explain why the people on top belong there; ours just happens to involve a laptop and a standing desk.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

AUSTIN, TEXAS — There is a particular look that comes over the face of a successful man when he is asked to account for his success. It is not quite a smile, not quite a smirk — call it the expression of someone who has been waiting all his life for the question and has rehearsed the answer in front of the bathroom mirror. He worked hard. He took risks. He saw what others did not. The deck, it turns out upon close inspection of his account, was stacked entirely in favor of merit, and he was merit's chief beneficiary. A coincidence, surely.

A cluster of essays this week — in The Guardian, in the Human Rights Research Center, in a sharp piece in The New Yorker on the entrepreneurial work ethic — has performed the now-ritual autopsy on the corpse of meritocracy. The body has been on the slab for some time. We keep cutting it open, peering inside, and announcing, with the air of discovery, that there is nothing there. Caste networks in Silicon Valley hiring. Women filtered out of information security by an old-boys' freemasonry that calls itself a pipeline problem. Founders whose chief qualification was a Stanford roommate with a checkbook. And yet the word survives, durable as a cockroach, because no one with anything to lose has any interest in burying it.

Meritocracy, as Michael Young coined the term in 1958, was meant as satire — a dystopia in which the winners felt entitled to their winnings and the losers felt obliged to accept their losings. The joke, of course, is that we read the book as an instruction manual. The genius of the thing is that it asks nothing of its beneficiaries except that they continue, in good faith, to believe in themselves.

The New Yorker's diagnosis of the entrepreneurial work ethic deserves a moment of its own. The cult of the founder — sleeping under the desk, optimizing the morning, journaling the gratitude, monetizing the hobby — is meritocracy's pious arm. It transforms the accidents of birth into the disciplines of character. If you are tired, you have not hustled enough. If you are poor, you have not believed enough. The Calvinists at least had the decency to admit that grace was unearned; our secular elect insist that theirs was a thing they built with their own two hands, in a garage, alone, against the doubters.

None of this is to say that effort does not exist, or that talent is a fiction, or that the lazy and the diligent are interchangeable. They are not. It is only to say that the distance between effort and outcome is paved with capital, kinship, accident, and the patient labor of people whose names will not appear on the cap table. The myth flatters the winners and shames the losers, which is, on reflection, what myths are generally for. We will not be giving it up soon. There is too much riding on it — chiefly, the self-regard of everyone it has favored.

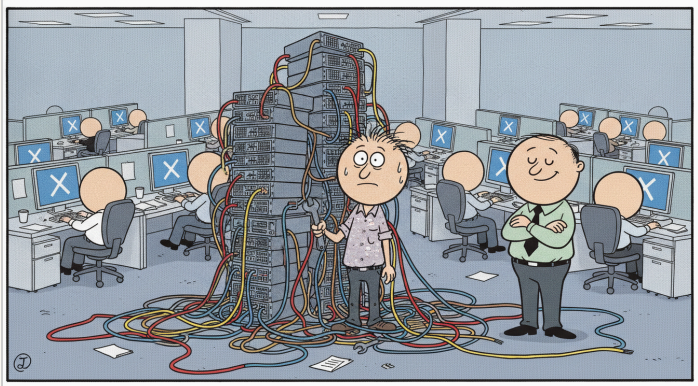

The Office Comic · Art Desk

Nation Briefly Considers Whether AI Productivity Gains Might Be Hidden Under 52% Off Bluetooth Speaker

As Qualcomm spends nearly $4 billion on a chip startup and employers await promised software savings, Americans remain focused on acquiring durable objects before the economy notices.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

SAN DIEGO — In a reassuring sign that the modern economy remains governed by forces no one can fully explain, Qualcomm has reportedly agreed to buy the buzzy chip software startup Modular for nearly $4 billion, a transaction that industry observers say should help answer the urgent question of whether artificial intelligence can become more profitable if it is first made much, much more expensive.

The deal, reported by WIRED, arrives at a delicate moment for the AI sector, which has spent the past several years promising to eliminate inefficiency while requiring unprecedented quantities of capital, electricity, specialized hardware, engineering labor, executive retreats, and words like "inference." Modular, a well-regarded startup focused on making AI and chip software easier to develop and deploy, will now enter the traditional innovation pipeline in which a nimble company is purchased by a much larger company so its agility can be carefully preserved in an org chart.

This is, of course, how progress works. First, software engineers use AI tools to produce code faster. Then companies wait to see whether that speed becomes profit. Then they purchase infrastructure, platforms, copilots, copilots for the copilots, and occasionally an entire chip software startup, all in the hope that somewhere inside the stack is a margin improvement quietly trying to get out.

The timing is instructive. Business leaders are now discovering that "developers can do more work faster" is not identical to "the business has more money," in the same way that owning 14 discounted USB-C hubs is not identical to having your life together. Productivity, like a missing board game piece, may technically be present in the house, but no one can locate it during the quarter-end review.

Fortunately, the American consumer has responded to this uncertainty with appropriate seriousness by examining Prime Day deals on board games and gear. While Qualcomm moves to consolidate part of the AI compute future, families across the country are deciding whether now is the correct time to buy a strategy game that will teach their children resource allocation, betrayal, and the importance of reading rules aloud for 47 minutes before giving up.

There is a certain honesty in this. A discounted board game makes a clear promise: you pay less than usual, open the box, and eventually accuse a loved one of misunderstanding turn order. AI infrastructure, by contrast, makes a more sophisticated promise: you pay far more than usual, open a dashboard, and eventually accuse a department of failing to restructure workflows around model-assisted throughput.

Even the $25,000 Slate electric pickup, backed by Jeff Bezos and positioned as a no-frills truck that owners can afford and repair themselves, feels almost quaint in comparison. It offers a physical object, a stated price, and a recognizable purpose. This places it at odds with much of the AI economy, where the object is often "capability," the price is "strategic," and the purpose is "unlocking enterprise value" at some later date.

None of this means the Qualcomm-Modular deal is misguided. On the contrary, it may prove extremely sensible. AI's next phase will likely depend on making models faster, cheaper, and easier to run across devices, and companies that control more of that stack will have obvious advantages. Qualcomm has long had ambitions beyond phones, and buying talent and software in the AI chip ecosystem is a perfectly rational act in a market where waiting patiently is considered a form of surrender.

But the broader lesson remains uncomfortable. The economy keeps being told that AI will make everything more efficient, while the companies selling and adopting it continue behaving as if the only path to efficiency is buying something enormous first. Employers want fewer bottlenecks, faster output, leaner teams, and better returns. Engineers are producing more. Vendors are raising more. Acquirers are paying more. Consultants are explaining more.

The payoff, we are assured, is coming.

Until then, there are 107 Prime Day deals on gear that has at least been tested by people who can confirm whether it turns on.

ON THIS DAY IN AI HISTORY

On June 24, 2012, Google's "Google Brain" team, led by Geoffrey Hinton, announced a major breakthrough in deep learning by training a neural network on unlabeled YouTube data that spontaneously learned to recognize cats without being explicitly programmed to do so. This milestone demonstrated the power of unsupervised learning and became a watershed moment for modern deep learning.

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