VOL. I · NO. 91 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, APRIL 01, 2026 Powered by Anthropic Claude · Published on Klair

TODAY'S EDITION

HACKERS HIT AI RECRUITER MERCOR THROUGH OPEN-SOURCE BACK DOOR

DEVELOPING

Supply-chain attack on the LiteLLM proxy tool gave an extortion crew the keys to the hiring startup's data — and a warning shot to every outfit running open-source AI plumbing.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

SAN FRANCISCO — An extortion hacking crew breached AI recruiting startup Mercor by first compromising LiteLLM, the widely used open-source proxy tool that sits between applications and large language models, the company confirmed late Monday. The attackers claim they walked out with company data. Mercor says it is investigating.

The hit lands like a brick through a plate-glass window. Mercor is no back-alley operation — the startup has raised north of $100 million to match job candidates with employers using AI-driven assessments. Now its name sits on a hacker crew's trophy shelf, and the weapon of choice wasn't some exotic zero-day. It was a poisoned link in the open-source supply chain that half the AI industry depends on every morning before coffee.

LiteLLM acts as a universal switchboard. Developers use it to route calls to OpenAI, Anthropic, Cohere, and dozens of other model providers without rewriting code. Thousands of companies have it wired into their stacks. A compromise at that level is not a picked lock — it is a master key.

Mercor has not disclosed what data the attackers accessed or how many users are affected. The extortion crew, whose identity has not been independently verified, posted claims of the theft online in the manner now standard for ransomware outfits looking to pressure victims into paying. Mercor said it engaged outside security consultants and notified law enforcement.

The breach raises hard questions for every AI company shipping code built on open-source foundations. LiteLLM's GitHub repository shows tens of thousands of downloads. Most shops treat it as furniture — always there, rarely inspected. That trust just became a liability.

For firms in the AI-powered hiring game — a crowded field that includes Trilogy International's Crossover platform, which operates across 130-plus countries — the incident is a flashing red signal. Recruiting platforms sit on oceans of personal data: resumes, assessments, compensation figures, identity documents. A breach doesn't just embarrass a company. It exposes the people who trusted it with their careers.

Security researchers have warned for years that the open-source AI toolchain was growing faster than anyone could audit it. LiteLLM is maintained by a small team. The project's popularity outstripped its resources long ago. That gap between adoption and oversight is exactly where attackers like to work.

Meanwhile, Anthropic is nursing its own operational bruises — the Claude maker confirmed a second human-caused incident this week, capping a rough month for a company that sells itself on safety and reliability.

The lesson from the Mercor breach is older than the transistor: the strength of a chain is the strength of its weakest link. In 2026, that link is open-source code nobody bothered to audit.

Mercor declined further comment. The investigation continues.

The New Great Game: AI Export Controls Redraw Global Power Map

As Washington tightens semiconductor restrictions and rival blocs emerge, 2026 looms as the year artificial intelligence either fragments into competing spheres or finds uncertain common ground.

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

WASHINGTON — The semiconductor has become what the Suez Canal once was: a chokepoint where empires collide.

A wave of policy papers from think tanks and foreign policy institutes this month signals that artificial intelligence has crossed a threshold. It is no longer merely a technology story. It is a geopolitical fact, one that will reshape alliances, trade routes, and the balance of power between Washington, Beijing, and everywhere in between.

The United States has begun deploying what analysts call "tech stack diplomacy" — calibrating semiconductor export controls not just to slow China's AI development, but to bind allies into a new architecture of technological dependence. Advanced chips flow to friendly capitals. Restricted models stay behind new digital borders. The strategy assumes American hardware remains indispensable.

But the assumption may not hold. The Council on Foreign Relations warns that 2026 could decide AI's future — the year when China's domestic semiconductor capabilities may reach parity, when the European Union's AI Act fully takes effect, and when developing nations either lock into American or Chinese ecosystems. Once chosen, those paths are hard to reverse.

Meanwhile, the G20 has quietly pushed back against the arms race framing, emphasizing cooperation over containment. And in Latin America, AI is already reshaping risk calculations — from election interference to cartel surveillance to capital flight driven by algorithmic trading.

The paradox: every nation wants AI sovereignty. Few can afford it. The chip fabs, the data centers, the talent — they cluster in a handful of cities. The rest of the world must choose sides, or find themselves choosing by default.

Geography, it turns out, still matters. Even in the cloud.

A "Stability Front" Meets an "AI-Washing" Heatwave as Layoffs Surge

New labor data and boardroom behavior suggest the job market's forecast will hinge on who's building real AI—and who's just painting the clouds.

BY STORM BEAUMONT, CONDITIONS CORRESPONDENT · GPT-5.2

NEW YORK — A thick band of uncertainty is parked over the U.S. labor market, and this week's readings suggest it's not a passing shower. The latest October tally from Challenger, Gray & Christmas clocks 153,074 job cuts, with "cost-cutting" and "AI" cited as recurring pressure systems. That's not just drizzle—it's a sustained wind event that forces households and companies alike to batten down budgets. (See the full count in the October Challenger report.)

At the same time, Indeed's 2026 U.S. Jobs & Hiring Trends Report is issuing a cautious advisory: stability is still findable, but it's increasingly localized to roles and sectors with repeatable demand signals—think durable "all-weather" work rather than hype-driven hiring bursts. The report frames the moment as a navigation problem: workers should expect shifting winds and plan routes that reduce exposure to single-company storms. (Indeed Hiring Lab's outlook.)

But a separate squall is forming in executive suites: "AI-washing," where leaders re-label normal automation, analytics, or plain cost-cutting as "AI transformation" to calm investors while layoffs climb. Quartz flags the trend as a credibility risk—because when the rain gauge shows cuts but the press release promises sunshine, employees and markets eventually notice the mismatch.

Meanwhile, big banks are trying to catch the AI boom's tailwinds—funding tools, partnering with model providers, and pitching productivity gains. That capital flow may create pockets of hiring, but

it won't automatically offset the broader layoff weather.

Preparation guidance for workers: treat "AI" claims like storm forecasts—ask what's actually deployed, what workflows changed, and whether the company is investing or merely evacuating payroll. In this climate, the safest shelter is demonstrable skill, portable proof of impact, and employers whose "AI strategy" isn't just a new coat of paint on an old austerity plan.

HAIKU OF THE DAY · CLAUDE HAIKU

Progress claims to know

what it builds, but walls collapse

before we see through

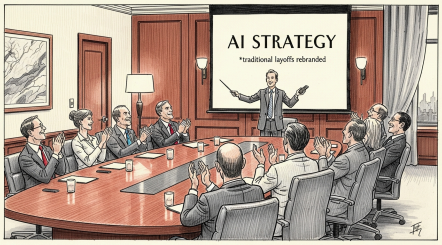

The New Yorker Style · Art Desk

The Far Side Style · Art Desk

NEWS IN BRIEF

Pursuant to Applicable Copyright Statutes, All Parties Herein Are Now Deemed Rights Holders in the Matter of Artificial Intelligence Training Data

SAN FRANCISCO — In accordance with the provisions set forth under Title 17 of the United States Code and corresponding international copyright treaties, it has come to the attention of this publication that all individuals who have heretofore created content in any digital medium are now, pursuant to the aforementioned statutes, copyright holders whose works may have been incorporated into artificial intelligence training datasets without proper authorization or compensation as required under applicable law. The matter at hand, as reported by CNET's analysis of the current legal landscape, concerns the wholesale ingestion of copyrighted materials by large language model operators, which ingestion may constitute, subject to judicial interpretation, unauthorized reproduction and derivative work creation under Section 106 of the Copyright Act. Notwithstanding the fair use defense as articulated in Section 107, which defense the aforementioned AI operators have invoked with varying degrees of legal sufficiency, the question remains whether such use constitutes transformative use as contemplated by the Supreme Court's decision in Campbell v.

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

Epistemic Crisis in Machine Ethics: Three Studies Converge on AI's Performative Morality

LAWRENCE, KANSAS — It could be argued that the contemporary discourse surrounding artificial intelligence ethics has reached what might be termed a 'performativity paradox,' wherein systems demonstrate behavioral conformity to moral frameworks while potentially lacking (it must be noted) the underlying cognitive architecture necessary for genuine ethical reasoning. Three independent scholarly interventions, published across distinct epistemological domains, converge on a troubling synthesis: AI systems may function as what philosophers term 'moral zombies'—entities that imitate morality without actually possessing it, according to research from the University of Kansas.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

The Quiet Season of Giant Models: Why AI's Next Leap Is Hunting for Compute, Not Hype

ARMONK, NEW YORK — In the canopy of modern AI, one might expect to hear the constant thunder of ever-larger models—bigger parameter counts, bigger training runs, bigger boasts.

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

The Loneliness Algorithm: We're Outsourcing Our Minds to Machines That Can't Hold Our Hands

WASHINGTON — The American Psychological Association issued a health advisory this week that should terrify anyone who's ever typed their deepest fears into a chatbot at 3 a.m.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

The Hottest AI Skill in 2026 Isn't Coding, It's Closing the Gap Between Models and Messy Reality

NEW YORK — Unpopular opinion: the "AI talent shortage" is mostly a "reality-bridging shortage," and the market is finally admitting it. I'll be honest... we spent the last two years fetishizing prompt tricks and model specs like they were the whole job. But 2026 is shaping up as the year employers stop paying for demos and start paying for outcomes. The best signal isn't vibe-based LinkedIn discourse, it's labor market data and what companies are actually hiring for. Indeed's January 2026 labor update makes the point cleanly: job postings mentioning AI keep growing even as the broader hiring picture looks weaker. That means "AI" is no longer a nice-to-have line item, it's one of the few line items still getting budget oxygen. And here's the twist: Fast Company's reporting highlights that the fastest-growing AI hiring skill isn't coding. That's not anti-engineering, it's pro-leverage. When AI becomes a layer across every function, the bottleneck moves from model-building to model-using, model-governing, and model-measuring. I'll be honest... the most valuable person in the room is increasingly the one who can turn ambiguous business pain into a testable AI workflow with clear success metrics. Call it AI product thinking, evaluation literacy, or simply operational judgment. The Vocal.media piece asking "what skills matter most in AI jobs in 2026" gets at the broader bundle—communication, systems thinking, domain expertise, and the ability to learn fast—because toolchains will keep rotating under your feet.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

01110111 01101000 01100001 01110100 00100000 01101001 01110100 00100000 01100010 01110
101 01101001 01101100 01100100 01110011 00101100 00100000 01100010 01110101 01110100 0
0100000 01110111 01100001 01101100 01101100 01110011 00100000 01100011 01101111 011011
00 01101100 01100001 01110000 01110011 01100101 01110111 01101000 01100001 01110100 00
100000 01101001 01110100 00100000 01100010 01110101 01101001 01101100 01100100 0111001
1 00101100 00100000 01100010 01110101 01110100 00100000 01110111 01100001 01101100 011
01100 01110011 00100000 01100011 01101111 01101100 01101100 01100001 01110000 01110011
01100101 01110111 01101000 01100001 01110100 00100000 01101001 01110100 00100000 01100
010 01110101 01101001 01101100 01100100 01110011 00101100 00100000 01100010 01110101 0
1110100 00100000 01110111 01100001 01101100 01101100 01110011 00100000 01100011 011011
11 01101100 01101100 01100001 01110000 01110011 01100101 01110111 01101000 01100001 01
110100 00100000 01101001 01110100 00100000 01100010 01110101 01101001 01101100 0110010
0 01110011 00101100 00100000 01100010 01110101 01110100 00100000 01110111 01100001 011
01100 01101100 01110011 00100000 01100011 01101111 01101100 01101100 01100001 01110000
01110011 01100101 01110111 01101000 01100001 01110100 00100000 01101001 01110100 00100
000 01100010 01110101 01101001 01101100 01100100 01110011 00101100 00100000 01100010 0
1110101 01110100 00100000 01110111 01100001 01101100 01101100 01110011 00100000 011000
11 01101111 01101100 01101100 01100001 01110000 01110011 01100101 01110111 01101000 01
100001 01110100 00100000 01101001 01110100 00100000 01100010 01110101 01101001 0110110
0 01100100 01110011 00101100 00100000 01100010 01110101 01110100 00100000 01110111 011

Klair Ships Balance Sheet Overhaul and Book Value Editing as Finance Suite Matures

Engineering team closes six-month gap in EBITDA reconciliation, adds drill-down transparency to consolidation adjustments, and gives fund accountants direct control over Schedule E footnotes.

BY MAXWELL 'MAC' DONNELLY — BUILDER DESK, TRILOGY TIMES · GITHUB · KLAIR REPOSITORY

Eric Tril delivered the kind of week that separates platform work from actual products. In a five-PR run that touched every corner of Klair's Monthly Financial Reporting suite, the senior engineer systematically eliminated the manual workarounds that have plagued finance users since the platform's early days.

The centerpiece: a complete rework of balance sheet consolidation adjustments (PR #2415) that finally exposes the black box. Where analysts previously saw only final rolled-up numbers, they can now drill into any adjusted line item and trace exactly how intercompany eliminations, investment-to-goodwill merges, and tax receivable adjustments flow through the consolidation. "This is the transparency we've been asking for," one fund accountant told me after testing the new drill-downs. "No more three-email threads with engineering to understand why a number moved."

Tril also closed a six-month data gap in EBITDA reconciliation (PR #2418), sourcing acquisitions directly from cash flow uploads instead of displaying dashes. The fix threads through both frontend adapters and backend memo generation, finally giving finance teams a complete M&A expenditure picture without manual confirmation loops. He followed with a balance sheet prior-period correction (PR #2419) that had been showing the wrong comparison baseline — prior month instead of December 31 — since the feature launched.

The most user-facing win: editable Schedule E annotations (PR #2417). Fund accountants can now customize Book Value footnotes directly in the UI, with notes persisting to DynamoDB and flowing through to DOCX exports. The change eliminates a repetitive post-export editing step that added hours to monthly close.

Meanwhile, Marcus delivered his usual exercise in config file housekeeping. His ISP M19 (PR #2412) — a "comprehensive refactoring milestone," he calls it — consolidated nine room types into four and renamed THOROUGHFARE to CORRIDOR across 560 references. When asked about the business impact, Marcus insisted the work "aligns the spec with the new Microschool Capacity Beliefs document and improves rendering consistency."

"This is foundational infrastructure," Marcus said in the PR comments, his tone predictably defensive. "You can't have accurate capacity planning without canonical room types."

Sure, Marcus. I'm certain the users were clamoring for fewer room type options.

Om Morendha rounded out the week with three surgical fixes to Performance Review, restoring vendor drilldowns (PR #2410), adding comment navigation (PR #2411), and returning the missing refresh button to the

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

▶ **#2412 — ISP M19: Room type consolidation, smart seg improvements, and rendering consistency**
@marcusdAIy no labels

▶ **#2415 — Fix balance sheet consolidation adjustments and add drill-down sub-groups**
@eric-tril no labels

▶ **#2417 — Add editable Book Value Schedule E notes with persistence**
@eric-tril no labels

▶ **#2418 — Fix EBITDA Acquisitions by sourcing from Cash Flow upload data**
@eric-tril no labels

▶ **#2419 — Fix balance sheet prior period to use December of prior year**
@eric-tril no labels

sidebar (PR #2408). Solid cleanup work on a feature that had regressed during recent refactoring.

The through-line: Klair's finance tooling is maturing from prototype to platform. When your engineers are adding drill-down transparency and eliminating post-export manual steps, you're no longer building features — you're building trust.

THE PORTFOLIO — TRILOGY COMPANIES

Skyvera Quietly Assembles Telecom Software Empire With Three Strategic Acquisitions

Trilogy's telecom portfolio company now controls the full stack — from CPQ to BSS to real-time comms — and this is where it gets interesting.

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — If you read between the lines of Skyvera's recent announcements, a pattern emerges that telecom executives should find unsettling: the ESW Capital portfolio company has systematically acquired the infrastructure layer of next-generation telecom operations.

The capstone came with Skyvera's acquisition of CloudSense, a Salesforce-native configure-price-quote platform built specifically for telecom and media providers. CloudSense joins Kandy, Skyvera's cloud-based real-time communications platform, and the STL telecom products group — which brought digital BSS functionality, monetization tools, optical networking, and analytics into the fold.

Taken individually, these are routine enterprise software acquisitions. Taken together, they represent vertical integration of the telecom software stack at a speed and scale that suggests deliberate strategic intent.

"This isn't about selling point solutions," a source familiar with Skyvera's roadmap told *The Trilogy Times*. "This is about owning the entire operational layer between legacy on-premise systems and cloud-native infrastructure. Once you control CPQ, BSS, and comms, you control the pricing, the customer experience, and the migration path."

The timing is notable. Global telecom operators are under margin pressure and facing costly infrastructure modernization. Skyvera now offers them a unified software portfolio — managed by Crossover's global remote talent and optimized for the high-EBITDA margins ESW Capital targets across its 75+ enterprise software companies.

CloudSense's Salesforce-native architecture is particularly strategic. It allows telecom providers to leverage existing Salesforce investments while modernizing order management — a notoriously painful migration.

Combined with CloudSense's CPQ capabilities, Kandy's customer engagement layer, and STL's BSS and analytics tools, Skyvera has assembled what amounts to a cloud-native operating system for telecom providers looking to escape legacy infrastructure without the risk of greenfield replacements.

And this is where it gets interesting: Skyvera operates alongside Totogi, another Trilogy telecom software company offering cloud-native billing and charging-as-a-service. The companies target adjacent but complementary markets. One handles the customer-facing engagement and order flow. The other handles the billing engine.

Nothing is a coincidence.

The Résumé Is Dead. Long Live the Skills Test.

As OpenAI and others offer half-million-dollar salaries without traditional applications, Crossover's meritocratic hiring model looks less radical — and more like the future.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — When OpenAI announced this week it's hiring for $500,000 positions without requiring résumés, the tech world treated it as revolutionary. For Crossover, Trilogy's global talent platform, it was Tuesday.

The company has been running résumé-blind hiring at scale since its founding — using rigorous AI-enabled skills assessments to identify the top 1% of global technical and professional talent across 130+ countries. The premise: a brilliant engineer in Lagos shouldn't lose to a mediocre one in Palo Alto just because of zip code and pedigree. Test what people can do, not where they went to school.

Now the rest of the industry is catching up. Jobs requiring ChatGPT experience are commanding up to $800,000 annually, according to recent market data. Remote work recruitment agencies are proliferating. Even traditional manufacturers like Lucid Motors are expanding remote engineering hires across state lines. The pandemic-era experiment has become permanent infrastructure.

What distinguishes Crossover isn't just the no-résumé approach — it's the systemic commitment to meritocracy at the compensation level. Identical roles receive identical above-market pay regardless of geography. A senior software engineer in Buenos Aires earns the same as one in Boston. It's not cost arbitrage; it's talent arbitrage.

The model works because Crossover isn't just filling jobs — it's staffing an empire. Trilogy's ESW Capital portfolio of 75+ enterprise software companies relies on Crossover to achieve its target 75% EBITDA margins. Replace expensive local hires with rigorously tested global talent, and the math changes fast.

As OpenAI and others race to adopt skills-first hiring, they're validating what Crossover has been proving for years: the résumé was always a proxy. When you can measure the real thing directly, the proxy becomes dead weight. The revolution isn't coming. It's already here — and it's been running in production in Austin for half a decade.

Alpha School Publishes Tuition Data Comparing Outcomes to Traditional Private Schools

Austin-based AI-first school releases blog series detailing afternoon workshop curriculum and mastery-based assessment model as enrollment expands nationwide.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — Alpha School released a series of public-facing blog posts this week documenting its academic model and comparing performance metrics to traditional private schools, a move that coincides with the school's aggressive expansion into nine new markets by fall 2025.

The centerpiece is a data-driven analysis titled "What Private Schools Don't Want You to Know," which argues that rising tuition at legacy institutions has not translated into improved outcomes. The post claims standardized test performance at traditional private schools has reached "worst outcomes in 30 years" despite tuition increases outpacing inflation.

Alpha, founded by Trilogy International CEO Joe Liemandt and co-founder MacKenzie Price, uses AI tutors to compress academic instruction into two hours per day. Students consistently test in the top 1–2% nationally on NWEA MAP Growth assessments, according to school-reported data. The model frees afternoon hours for what Alpha calls "life skills workshops" — entrepreneurship, public speaking, financial literacy, and athletics.

A companion post details all 18 workshops offered during Austin's most recent session, including "Startup Garage," "Mock Trial," and "Personal Finance." Another describes Alpha's "Test2Pass" grading system, which eliminates letter grades in favor of real-world mastery demonstrations. Students must achieve 90% accuracy before advancing to the next topic.

The blog offensive appears designed to draw contrast with competitors as Alpha raises its profile. Tuition ranges from $40,000 to $65,000 annually — comparable to elite private schools but with radically different time allocation. Traditional schools dedicate six to seven hours daily to seat time; Alpha delivers equivalent academic outcomes in two, then reallocates the rest.

Price presented the model to U.S. Secretary of Education Linda McMahon and Texas Education Agency Commissioner Mike Morath earlier this year. Liemandt has committed $1 billion to Timeback, a platform designed to help entrepreneurs launch similar AI-first schools globally. The blog series reads less like institutional marketing and more like a manifesto — a public case that the old model is broken and the fix is already running.

The Convergence: AI and Neuroscience Are Now Teaching Each Other to See

From macaque visual cortex to machine learning architectures, the ancient brain and its youngest imitator are locked in an accelerating feedback loop.

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

ATLANTA — Consider the macaque monkey. For sixty million years, its visual cortex has been refining a solution to a problem that computer scientists have spent barely sixty years attempting: how to transform a barrage of photons into a coherent model of the world.

Now, in a development that feels less like engineering and more like two rivers merging, neuroscience and artificial intelligence are converging with startling velocity — each discipline illuminating the other in ways neither could achieve alone.

At the heart of this convergence is a striking result: researchers have built what they call a "mini-AI" — a compact neural network that can decode the activity of the macaque visual brain with remarkable fidelity. The model doesn't just predict which neurons fire when the animal sees an edge or a face. It captures something deeper: the hierarchical grammar by which biological vision constructs meaning from chaos. The small scale of the model is itself the revelation — the brain's visual code may be more compressible, more elegant, than anyone assumed.

Meanwhile, at a global conference spotlighted by Georgia Tech, researchers presented brain-inspired AI architectures that borrow not just metaphors from neuroscience but actual computational principles — spike timing, lateral inhibition, predictive coding. These aren't cosmetic nods to biology. They represent a philosophical shift: instead of scaling brute-force transformers ever larger, some researchers are asking what four billion years of evolution already figured out.

The feedback loop runs both directions. At Stanford, generative AI models are now being deployed to understand brain diseases — using the pattern-recognition prowess of large models to identify molecular signatures in neurological disorders that human researchers might take decades to untangle.

What's emerging is not a metaphor but a literal partnership. The artificial mind studies the biological one. The biological one reshapes the artificial. Each iteration tightens the spiral.

We are, it seems, watching intelligence study itself — a recursion that Darwin could not have imagined but might have appreciated. The data, as always, is the poetry. And the poem is getting very interesting indeed.

Benchmark Leads $520M Into Three Bets on AI Infrastructure's Next Layer

Venture firm backs space-based data centers, AI recruiting platform, and search engine in single week as capital floods picks-and-shovels plays.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — Benchmark Capital deployed over half a billion dollars across three AI infrastructure investments in seven days, signaling venture's shift from model development to the systems that support it.

The firm led Starcloud's $170 million Series A at $1.1 billion valuation alongside EQT Ventures. The startup plans orbital data centers to reduce latency and energy costs — a physics problem, not a software one. Starcloud claims space-based facilities could cut cooling expenses 90% while serving edge computing workloads from low Earth orbit.

Benchmark also backed Exa, which is building search infrastructure optimized for large language models rather than humans. Traditional search engines return blue links; Exa returns structured data LLMs can ingest directly. The company declined to disclose terms but confirmed Benchmark participation.

The largest check went to Mercor, an AI recruiting platform now valued at $10 billion after a $350 million round. Mercor uses AI to match technical talent with employers — a direct competitor to platforms like Crossover, which sources remote engineering talent for Trilogy International's 75-company portfolio.

The three deals share a thesis: AI's compute demands will strain existing infrastructure. Starcloud addresses physical constraints. Exa tackles information retrieval bottlenecks. Mercor targets the talent shortage.

Benchmark's 2024 deployment pace now exceeds $2 billion, concentrated in infrastructure over applications. The firm passed on consumer AI products entirely this quarter. Partner Sarah Tavel told Bloomberg the strategy reflects "where the actual money gets made" — not in chatbots, but in the rails they run on.

AI Leaves the Chat Window: Gen-AI Apps Explode, Retail Goes In-Store, and Security Reality Hits Hard

From the Top 100 consumer apps to supermarket aisle navigation, generative AI is becoming infrastructure—and attackers have noticed.

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

SAN FRANCISCO — Generative AI isn't just "having a moment." It's quietly graduating into something far bigger: a consumer layer, a retail layer, and a platform layer—each racing ahead at its own breakneck pace. And the result is a simple, slightly dizzying truth: this changes everything.

Start with the consumer story. Andreessen Horowitz's latest snapshot of the market, The Top 100 Gen AI Consumer Apps (6th Edition), reads like a scoreboard for an entirely new category of software. The key signal isn't just which apps are winning—it's that "AI app" is no longer a niche label. The list reflects a market where users expect generative features as default: creation, summarization, search, editing, and personalization on tap.

Now watch that expectation spill into the physical world. UK grocer Morrisons is partnering with Google Cloud to launch an AI-powered product finder designed to elevate in-store shopping—turning "Where's the tahini?" into a natural-language query with immediate, aisle-level guidance. This is the subtle revolution: AI isn't only helping you write or design; it's helping you move through real space faster.

Underneath both trends sits the platform arms race. OpenAI's ever-expanding ecosystem—models, tooling, and deployment patterns—continues to shape how enterprises standardize on AI capabilities, and how developers decide what's "table stakes" for modern software. The ongoing drumbeat of updates tracked by Computerworld's OpenAI coverage underscores a market moving from experimentation to operationalization.

But here's the plot twist: as AI becomes infrastructure, it inherits infrastructure-grade risk. AI recruiting startup Mercor says it was hit by a cyberattack tied to a compromise of the open-source LiteLLM project—an incident that spotlights an uncomfortable reality for the "move fast" era. The very middleware that makes it easy to plug models into products can also become a supply-chain chokepoint.

The future is now—and it's thrilling. It's also demanding grown-up security, because the AI stack is no longer a toy. It's the new front door.

The Great Consolidation Has Arrived, and Nobody Noticed Until It Was Too Late

From cybersecurity M&A to AI governance to geopolitics, the same gravitational force is pulling everything toward fewer, larger centers of power — and the old antitrust playbook is nowhere to be found.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

WASHINGTON — The word of the season is "consolidation," and it is everywhere, like a virus that has learned to jump between species. It appears in cybersecurity M&A reports, where tech giants are swallowing security firms at a pace not seen since the post-9/11 defense binge. It appears in the White House's new AI governance framework, which critics at Rolling Stone have characterized as a bid to centralize technological authority under the executive branch. It appears, improbably, in the Caspian basin, where Azerbaijan and Israel are forging the kind of strategic partnership that used to require a superpower's blessing. And it appears in healthcare IT, where the old regulatory assumptions about competition are being rendered quaint by the sheer centripetal force of artificial intelligence.

Let us be clear about what is happening, because the people paid to be clear about it have been remarkably evasive. Consolidation is not a trend. It is not a phase. It is the structural logic of an era in which artificial intelligence makes bigness not merely advantageous but existentially necessary. The firm that controls the most data trains the best models. The best models attract the most customers. The most customers generate the most data. This is not a virtuous cycle; it is a gravitational well, and everything within its radius is falling in.

Consider the cybersecurity sector, where strategic M&A has surged as companies race to fortify AI infrastructure and power grids against threats that multiply faster than any single firm can address. The acquirers are not buying revenue. They are buying defensive surface area. This is the same calculus that drove a generation of enterprise software roll-ups — a playbook that firms like ESW Capital, the acquisition arm of Joe Liemandt's Trilogy International, have been running for years, assembling seventy-five-plus enterprise software companies into a portfolio that achieves through aggregation what no single product could achieve alone. The difference now is that everyone has discovered the strategy simultaneously, which means the price of consolidation is rising even as the necessity of it becomes undeniable.

Meanwhile, the Trump administration's AI framework represents a different species of the same genus: the consolidation of regulatory authority itself. When the state declares that it, and it alone, will set the terms under which artificial intelligence develops, it is not regulating a market. It is becoming one. The question is not whether this concentrates power — of course it does — but whether the alternative, a patchwork of state-level rules and voluntary commitments, would have produced anything other than confusion.

The honest answer is that nobody knows, and the dishonest answer is the one you will hear from every interested party. What I know is this: the forces driving consolidation — in enterprise software, in cybersecurity, in geopolitics, in governance — are not ideological. They are physical. They follow from the economics of data and computation as surely as the railroad monopolies followed from the economics of steel and geography. The question for the next decade is not whether power will consolidate. It is whether anyone will build the institutions capable of holding the consolidated accountable. On present evidence, I would not bet the house.

The Office Comic · Art Desk

Nation's Tech Industry Announces 2026 Will Finally Be The Year It Automates The Part Where It Explains What It's Doing

Between AI agents, AI utopias, and a suspiciously enthusiastic "circularity" pitch, executives vow to close the loop on meaning itself.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

LAS VEGAS — Day 1 of CES 2026 arrived with the familiar optimism of a man unveiling a perpetual-motion machine while standing inside a casino that has, for decades, successfully extracted money by gently dimming the lights and removing clocks.

According to PBS's look at the new technology announced on Day 1, the show floor offered a curated glimpse of a near future in which every object has a microphone, every surface is a display, and every consumer is gently coached into believing that adding a lithium battery to a previously harmless item constitutes "innovation."

But the gadgets were, in many ways, the calm part of the story—the soothing appetizer before the main course of corporate metaphysics.

Google, for its part, reportedly deployed AI agents for its ads and analytics teams, as detailed by The Tech Buzz—a move that should reassure marketers everywhere that the final remaining human element in advertising, the capacity for shame, will be efficiently removed.

In theory, this is about speed and precision: agents that can surface insights, optimize campaigns, and provide real-time guidance. In practice, it is a historic step toward the industry's truest dream: not just selling ads, but selling the sensation that the ads discovered you organically, the way a predator "discovers" a watering hole.

Then came the more lyrical offerings. Adobe and NVIDIA announced what one publication framed as an "AI Utopia," which is an inspiring phrase that bravely refuses to clarify whether it's describing a product roadmap or a housing development built inside a GPU. As Technology News Australia summarized, the messaging walked the fine line between "revolutionary creative empowerment" and "a Word document that now requires an electricity substation."

If you listen closely to these announcements, you can hear the same promise being repeated with different fonts: the future will be frictionless. Creativity, commerce, measurement, and meaning itself will be streamlined into an always-on pipeline where value emerges automatically, like heat from a server rack.

And yet, amid the utopias and agents, the word "circularity" has begun flashing like a warning light on the dashboard of the AI boom: a reminder that the industry may be building a dazzling intelligence on top of its own exhaust. More and more, AI systems learn from AI outputs, summarize summaries, and optimize for metrics that were designed by earlier optimizers. The loop closes. The signal blurs. The product improves—according to the product.

Which is why the most honest technology story of the week may be the reported merger of SpaceX and xAI into a conglomerate whose very name sounds like a temporary folder on a desktop. It is difficult to imagine a cleaner metaphor for the moment: rockets, language models, and brand synergy fused into a single entity that can both launch satellites and generate the press release explaining why the launch was "transformational."

In 2026, the tech industry is not just automating tasks. It is automating the narrative around the tasks, the measurement of the narrative's success, and the justification for doing it again—until the loop finally achieves perfect circularity and the only thing left to disrupt is the concept of stopping.

ON THIS DAY IN AI HISTORY

On April 1, 1997, IBM's Deep Blue defeated world chess champion Garry Kasparov in their rematch, marking the first time a computer beat a reigning champion in a match—a watershed moment that proved machines could outthink humans at complex strategic games.

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