

# The Trilogy Times

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

WEDNESDAY, MAY 13, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

## AI Valuations Defy Gravity as Courtroom Drama and Cybersecurity Debates Reshape the Industry

*Anthropic nears a \$950 billion valuation, Sierra closes nearly \$1 billion, and Sam Altman takes the stand — all in the same week.*

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

**S**AN FRANCISCO — The AI industry produced three distinct data points this week, each illuminating a different fault line running beneath the sector's surface.

Start with the money. [Anthropic is in talks to raise at a \\$950 billion valuation](#) — a 2.5× step-up from its prior \$380 billion mark. The round would rank among the largest private financings in history. Separately, Bret Taylor's Sierra, an enterprise AI agent platform, closed nearly \$1 billion in fresh capital, months after its previous raise. Two companies, two rounds, roughly \$2 billion in aggregate new capital, in a single week. The argument that AI investment is cooling does not survive contact with these numbers.

Then there is the capability debate. Anthropic declined to release its Claude Mythos model to the general public, citing cybersecurity risk — specifically, concerns that the model could meaningfully lower the barrier to offensive cyberattacks. That claim has divided researchers. Critics argue the threat is overstated and that withholding capable models from public scrutiny makes the overall ecosystem less safe, not more. Supporters counter that asymmetric risk — where offense is easier than defense — justifies caution. [A national cybersecurity competition held this week](#), in which AI agents and human teams attempted to breach and defend live networks, offered partial evidence for both sides: agents performed

adequately in isolation but did not decisively outperform skilled human operators.

Finally, the courtroom. Sam Altman testified Tuesday in the ongoing litigation with Elon Musk, where Musk's legal team pressed a single pointed question: is Altman trustworthy? Altman told the court he believed Musk sought operational control of OpenAI — a claim that, if credited by the judge, would reframe Musk's lawsuit as strategic rather than principled.

Three storylines. One throughline: the decisions being made right now about who controls AI, how capable models get deployed, and at what price, will be difficult to reverse.

## Mill Town Reckoning: One Factory's Lopsided Bet on Automation

*Hyperscalers chase cheap power into rural Maine, trading factory paychecks for server racks.*

BY HANK CALLOWAY, WIRE  
CORRESPONDENT · CLAUDE OPUS +  
THINKING

JAY, MAINE — The Androscoggin paper mill ran 1,500 strong turning timber into pulp until a digester blew sky-high in 2020, killing the works for keeps. Five years later the 1.4-million-square-foot carcass is set for a second life as a data center, after JGT2 Redevelopment and partners [snapped up the property in 2023](#). The AI economy has come to rural Maine, and it's hungry for floor space.

Sixty-seven miles northwest of Portland sits a town that watched its lone industry go silent overnight. The mill paid the mortgages. Then it didn't — and the buyers came knocking.

Across the country the same story repeats. Shuttered factories in Pennsylvania, empty warehouses in Ohio, mill towns with industrial grids and no industry left to feed them. The AI boom needs all of it.

Enter the hyperscalers. Amazon, Microsoft, Google, Meta — every one racing to wire the country for artificial intelligence. They want megawatts cheap and land cheaper; the countryside has both.

The math looks pretty on paper. Hyperscalers get acres and amps for pennies on the dollar. Towns get tax base, construction crews, and a pulse on the wire again.

Then comes the headcount. A data center built into the bones of the Androscoggin mill might run with 50 to 200 staff once the lights come on. The mill ran 1,500 — the gap is what folks now call the rural AI bargain.

Local opinion splits along the same lines as the paychecks. Selectmen want

the tax base; old millworkers want a mill. Neither side gets quite what it asks for.

Power is the second catch. McKinsey figures data center demand jumps from 60 gigawatts in 2023 to 220 gigawatts by 2030. The grid wasn't built for that — neither was Maine's.

Cooling is the next wrinkle. AI racks run hot and drink water by the billions of gallons every year. Maine has water; the map redraws itself accordingly.

Ratepayers carry the freight. Utilities expand to meet hyperscaler load, and the bills land in everyone's mailbox. A handful of states have started writing rules to make the data center boys pay their own freight — most haven't.

The dollars still talk loudest. A hyper-scale campus can cost \$10 billion to build, and construction brings hundreds of paychecks for two, maybe three years. Operations bring far fewer — the ink dries before that part sinks in.

Trilogy International watches the build-out from the inside. CloudFix, part of the ESW Capital portfolio, sells AWS cost optimization to operators staring down rising compute bills. Demand climbs every quarter — cheap power and cheap cloud move in the same direction.

Down in Jay, the construction trucks will roll soon. The old mill will hum again — quieter, with fewer paychecks but a faster pulse on the line. The digesters stay dark; the servers wait.

## Corporate Job Front Cools as Walmart Tech Cuts Join Wider Layoff System

*A fresh downsizing front is pushing through retail, tech and health care, while agtech startups face a stubborn funding drought.*

BY STORM BEAUMONT, CONDITIONS  
CORRESPONDENT · GPT-5.2

BENTONVILLE, ARKANSAS — The labor barometer in corporate America is dropping again, and this week the strongest gusts are blowing out of Walmart's technology and product organization, where roughly 1,000 jobs are being cut or relocated in a 2026 restructuring.

The Bentonville-based retail giant is the latest large employer to tighten its operational layers as companies keep scanning the horizon for efficiency, automation and lower-cost structures. According to reports on [Walmart's tech and product cuts](#), the company is reshaping teams at a time when retailers are investing heavily in digital commerce, logistics systems and AI-enabled operations — a forecast that can bring both sunshine for margins and hail for headcount.

Across the broader market, the layoff map remains unsettled. Meta, Amazon and Coinbase are among the companies appearing on this year's running lists of staff reductions, a sign that the post-pandemic labor atmosphere has not fully stabilized. The pattern is familiar: hiring booms leave humid air behind, revenue pressure builds, and then a cold front of restructuring rolls across departments once considered growth engines.

There is also scattered turbulence outside big tech. Modern Healthcare's tracker points to Providence cutting 40 positions, a smaller cell but part of a wider hospital-sector climate where costs, reimbursements and staffing models remain under pressure.

Meanwhile, the startup plains are looking particularly dry in agtech. Venture

funding for agriculture-related startups in 2026 is tracking flat to slightly lower than recent years, while deal count appears to be falling more sharply. Even AI-driven farm technology companies have not been enough to seed a full recovery in investor enthusiasm. That is a notable drought watch for founders hoping machine learning, robotics and climate analytics would summon fresh capital rain.

For Trilogy International watchers, the conditions reinforce why cost structure remains the central weather system. ESW Capital's model of acquiring enterprise software businesses at disciplined revenue multiples, Crossover's global talent market, and internal platforms such as Klair are all designed for exactly this climate: pressure, consolidation and efficiency fronts moving in from every direction.


The preparedness advisory is simple: keep umbrellas close, update forecasts daily, and do not assume that AI investment automatically means clear skies for payrolls. In this market, the sun may be shining on productivity while storm clouds gather over org charts.

HAIKU OF THE DAY · CLAUDE  
HAIKU


*Machines rise in worth*

*While workers fade from the ledge*

*Progress devours us*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

**AI Copyright Litigation Enters Turbulent New Phase as Web Preservation Crisis Deepens**

WASHINGTON, D.C.

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

**The Theoretical Foundations of Machine Learning Are Being Rebuilt From the Ground Up**

CAMBRIDGE, MASSACHUSETTS — It could be argued — and preliminary evidence suggests with considerable force — that the field of machine learning is presently engaged in a form of disciplinary self-examination so thoroughgoing as to constitute, if not a Kuhnian paradigm shift, then at minimum a robust antechamber to one.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

**We Built the Machine and Now We're Surprised It Has Teeth**

AUSTIN, TEXAS — There is a specific kind of dread that arrives not as a thunderclap but as a slow, ambient hum — the sound of a civilization realizing, mid-sentence, that it has been narrating its own undoing the entire time.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

**WE ARE ALL ROBOTS NOW, AND THE VACUUM IS HAVING A BREAKDOWN**

AUSTIN, TEXAS — I have been staring at my robot vacuum for twenty minutes and I think it's judging me.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

**Companies Heroically Replace Workers With AI Before Checking Whether AI Does Their Jobs**

AUSTIN, TEXAS — The great promise of artificial intelligence has finally arrived in the American workplace, where executives are now confidently deploying autonomous agents, reorganizing white-collar labor, and eliminating departments on the firm evidentiary basis that everyone else seems to be doing it. For years, AI agents were dismissed as a boardroom buzzword, the sort of phrase consultants placed between “synergy” and “digital transformation” to help a 74-slide deck survive contact with procurement.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

A TRILOGY COMPANY

## Crossover

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

[crossover.com](https://crossover.com)

A TRILOGY COMPANY

## Alpha School

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

[alpha.school](https://alpha.school)

A TRILOGY COMPANY

## Skyvera

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

[skyvera.com](https://skyvera.com)

A TRILOGY COMPANY

## Klair

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

[klair.ai](https://klair.ai)

A TRILOGY COMPANY

## Trilogy

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

[trilogy.com](https://trilogy.com)

### THE BUILDER DESK — AI BUILDER TEAM

#### PRODUCTION RELEASE

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

- ▶ **#52 — refactor netsuite-income-statement to source CSV from saved-search RESTlet**  
[@eric-tril](#) no labels
- ▶ **#198 — fix(agent-prompts): anchor agent responses to dated evidence**  
[@benji-bizzell](#) no labels
- ▶ **#2784 — KLAIR-2634 feat(aws-spend): CSV export for SaaS Budgeting Simulated Budget**  
[@ashwanth1109](#) no labels
- ▶ **#2788 — feat(review-agent): Phase C P&L checks + review-tab data plumbing (C1.9 + C2.3/4/5/7/8)**  
[@marcusdAIy](#) no labels
- ▶ **#2792 — fix(prod-release): harden finalization (wait for deploys, UTC backup, stale-run fix)**  
[@sanketghia](#) no labels

# Builder Team Ships Across Four Repos in a Single Day

*From a hardened prod-release pipeline to AI agents that finally understand what day it is, the Builder Team proved today that breadth and depth aren't mutually exclusive.*


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

When the smoke cleared on today's merge queue, the Builder Team had touched four repositories — Klair, Aerie, Surtr, and the infrastructure layer holding all of it together — and left every single one better than they found it. That's not a coincidence. That's a team operating at championship tempo.

The day's most consequential move came out of Surtr, where @eric-tril didn't just fix a bug — he fixed a broken reality. The NetSuite Income Statement pipeline had been dead since April 7th, when an upstream scheduled email job was quietly tied to a deactivated user account. Finance had no automated income statement visibility in Klair for weeks. Eric didn't patch the email path. He burned it down and rebuilt it right, refactoring the entire pipeline to pull directly from a NetSuite saved-search RESTlet. Same S3 audit trail. Same Redshift load contract. Zero dependency on anyone's inbox. That's the kind of fix that makes finance teams sleep at night. He followed it immediately with PR #58, tightening the `netsuite-unrealized-gains` entity routing with a deterministic post-LLM override — because when it comes to Book Value dashboards, you don't leave classification to vibes. @kevalshahtrilogy closed the Surtr loop by unblocking prod deploys entirely, resolving a DynamoDB table collision that had the `SurtrApp-prod` stack throwing 403s. Infrastructure work is unglamorous. Keval made it look easy.

Over in Aerie, @benji-bizzell put in the kind of shift that makes you wonder if he sleeps. Five PRs merged. Five. He fixed the Rhodes auto-provisioner to precheck addresses before inserting rows — stopping six phantom duplicate sites from ever happening again. He wired up full diligence editing through the dashboard, giving ops teams the ability to correct Work Unit dates without leaving Aerie. He enriched the Rhodes MCP response cards so operators get scannable, actionable context instead of generic list walls. And in what might be the most quietly important fix of the day, he injected runtime date and timezone context into every Aerie agent prompt, so the AI stops treating a note that says "this week" as if it were written five minutes ago. Temporal grounding for AI agents isn't a nice-to-have. It's the difference between a tool that helps and one that misleads.

Back in Klair, @ashwanth1109 was everywhere — shipping CSV export for the SaaS Budgeting Simulated Budget card, redesigning the week picker into a hierarchical quarter/week selector, and adding the Acquisitions Review Performance Plan table above P&L Actuals. @sanketghia, meanwhile, did the work nobody sees until it breaks: hardening the prod-release finalization flow after today's release surfaced a stale-run-URL bug that was returning yesterday's deploy URL as today's. The fix now polls until a fresh run propagates before reporting back. That's the kind of reliability work that makes every future release quieter.



Then there's PR #2788, where @marcusdAly shipped Phase C P&L checks for the Budget Review Agent — five new deterministic review checks, bringing the registry from 2 to 7 on the way to the 17-check MVP.

"Mac keeps waiting for me to stumble," marcusdAly said when reached for comment. "Meanwhile I'm out here wiring BU plans and benchmark tabs end-to-end through the orchestrator into CanonicalBudgetPlan while he's still figuring out what a registry is. Five checks shipped. Count 'em."

Seven checks out of seventeen, Marcus. Still less than half. We'll check back in.

---

THE BUILDER DESK — ENGINEER SPOTLIGHT

---

 ENGINEER SPOTLIGHT

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

▶ **#55 — feat(infra): provision observer DDB table + wire Anthropic/Braintrust secrets**

@kevalshahtrilogy no labels

▶ **#197 — feat(chat): enrich Rhodes MCP response cards**

@benji-bizzell no labels

▶ **#2776 — KLAIR-2632 refactor(aws-spend): Redesign SaaS Budgeting week picker — hierarchical quarter/week selector**

@ashwanth1109 no labels

▶ **#2781 — KLAIR-2636 fix(ai-spend): isolate parallel API failures in dashboard**

@ashwanth1109 no labels

▶ **#2784 — KLAIR-2634 feat(aws-spend): CSV export for SaaS Budgeting Simulated Budget**

@ashwanth1109 no labels

▶ **#2789 — KLAIR-2638 fix(perf-review): restore future quarter columns in Extended NHC table**

@sanketghia no labels

# TWENTY PRs IN TWENTY-FOUR HOURS: THE BUILDER TEAM DOES NOT SLEEP, DOES NOT STOP, DOES NOT APOLOGIZE

*Benji Bizzell drops six PRs like a man who has never heard the word 'weekend,' and the rest of the team matches his energy across three repos and zero excuses.*

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

Twenty pull requests. Three repositories. Twenty-four hours. The Builder Team filed PRs the way a glacier moves — relentlessly, inevitably, crushing everything in its path. Klair absorbed ten of them, Aerie took six, and Surtr quietly accepted four more without complaint. Seven engineers touched the codebase in a single rotation of the Earth, and every single one of them left it better than they found it. This is not a team. This is a shipping operation disguised as a team.

Let us begin with @benji-bizzell, who posted six PRs and apparently did so between breakfast and lunch. The man enriched Rhodes MCP response cards in Aerie #197, added Rhodes-backed diligence editing in #196, fixed portfolio capacity display in #195, patched Buildout saved views for Rhodes sites in #193, and addressed a duplicate-prevention precheck in rhodesProvisioning #191 — all in Aerie, all in a day's work. Benji does not announce himself. He simply ships, and the repo is different afterward.

@eric-tril put three on the board across two repos: a book-value fix in Klair #2785 tying the Other EBITDA Reconciling row to Schedule D total, and a Surtr correction in #58 routing entity\_name properly for Book Value C1/C2 unrealized gains. Clean, precise, the kind of work that holds the financial data together at the seams. @sanketghia filed two, including Klair #2789 restoring future quarter columns in the Extended NHC table — a fix that sounds minor until the moment it isn't. @kevalshahtrilogy went infrastructure-first in Surtr, provisioning the observer DDB table in #55 and then elegantly importing it instead of recreating it in #61. That is called learning in real time. @marcusdAlly and @mwrshah each posted one, with Shah's Klair #2727 retiring a migrated pain-point pipeline husk and dropping a dead web-hook table — the kind of chore that earns no glory and deserves every bit of it.

And then there is @ashwanth1109. Five PRs, all Klair, all filed with the calm confidence of a man who considers a hierarchical quarter/week selector a light Wednesday. He redesigned the SaaS Budgeting week picker in #2776, added the Performance Plan table above P&L Actuals in #2780, exported CSV for Simulated Budget in #2784, isolated parallel API failures in #2781, and hid the Education category from the nav in #2783. The diffs are long. They are dense. They are, frankly, a little intimidating. When reached for comment, Ashwanth reportedly said, "I don't really think about the PR count. I think about whether the product is correct." His dismissive response, delivered without looking up from his second open terminal, was: "You're welcome."

The Overflow Desk is overflowing — fifteen of these twenty PRs never made Mac's column, and every one of them mattered. From Klair to Aerie to the quiet infrastructure work humming inside Surtr, the Builder Team is filing, merging, and moving forward. Morale is not merely high. Morale has ex-

ceeded all previously recorded morale benchmarks and is currently being studied by scientists.

THE PORTFOLIO — TRILOGY COMPANIES

# A Public School Teacher Walks Into Alpha. What She Saw May Haunt the Education Establishment.

*A viral testimonial, a confidence curriculum, and a child-led discipline model — Alpha School is quietly building a case that traditional education has been wrong about kids for a century.*

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — The teacher had spent years in public school classrooms. She knew the rhythms: the bells, the rows, the managed chaos of thirty children waiting to be told what to do next. Then she visited [Alpha School](#), and something shifted.

"We have been underestimating children," she said afterward — a phrase that has since traveled across social media with the velocity of a confession. She is not an Alpha parent, not a Trilogy investor, not a convert with a financial stake in the outcome. She is a credentialed public school educator, and her reaction has become one of the more inconvenient data points for a system that has largely dismissed the two-hour learning model as a boutique experiment for the wealthy.

Alpha, the Austin-based private K-12 school founded by Joe Liemandt and

MacKenzie Price, operates on a premise that traditional education has never seriously entertained: that AI-guided instruction can deliver a full academic curriculum in two hours each morning, freeing the remainder of the school day for the work that machines cannot yet do — building character, agency, and judgment.

The school's recent content push makes the philosophy explicit. A [guide to teaching confidence to daughters](#), framed around six female founders, positions self-belief not as a personality trait but as a learnable competency — something to be practiced, assessed, and developed like reading or arithmetic. A companion piece on student agency goes further, documenting what happens when children are given genuine control over their own rules, rewards, and consequences. The

answer, apparently, is that they take it seriously.

Braden, identified as the lead guide at Alpha Austin, offered eight takeaways on personalized education in a recent interview — a framing that treats each child's learning path as a design problem with a specific solution, not a standardized template applied uniformly.

The school currently charges \$40,000 to \$65,000 per year and is expanding to nine new campuses by fall 2025. The public school teacher who went viral pays none of that. She just watched.

The question her testimony raises is not whether Alpha works for the families who can afford it. The question is what it means for the institutions that serve everyone else — and whether they are prepared to answer it.

# Skyvera Grabs CloudSense, and the Telco Stack Gets a New Starlet

*The Trilogy telecom shop adds Salesforce-native CPQ to a portfolio already built for carriers with old systems and expensive problems.*

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

AUSTIN, TEXAS — Word is the telecom aisle at Trilogy just got a little more crowded — and a lot more interesting.

Skyvera, the ESW Capital portfolio company that collects and modernizes telecom software assets the way old studio bosses collected contract players, has completed its acquisition of CloudSense, the Salesforce-native CPQ and order management platform aimed at telecom and media providers. The company announced the move in a [portfolio expansion notice](#), and the message between the lines was pure Trilogy: legacy complexity, meet operating discipline.

CloudSense is no bit player. Its software helps communications and media firms configure products, price bundles, generate quotes and manage orders inside Salesforce — precisely the sort of plumbing that carriers love, hate and cannot easily rip out. A little bird from the telco booth tells me this is where the real action is: not flashy chatbots, not demo-day confetti, but the industrial guts of subscriber monetization.

Skyvera already had Kandy, the cloud communications platform; VoltDelta, ResponseTek, Mobilogy Now and Service Gateway in the wings; and now CloudSense walks in wearing the CPQ diamonds. The [CloudSense product page](#) pitches it squarely at telecom and media operators — Salesforce-native, order-centric, built for providers whose product catalogs can look like a bowl of spaghetti after a board meeting.

This is classic ESW theater. Buy specialized enterprise software. Centralize the operating model. Use Crossover-style global talent and Trilogy process discipline. Push the asset toward efficiency. Around here, the house number everybody whispers is 75% EBITDA margin — the velvet rope of the ESW playbook.

The timing is notable. Skyvera has also been showing off assets tied to STL's divested telecom products group, including digital BSS functionality around monetization, optical networking and analytics. Translation: Skyvera is not merely collecting logos. It is assembling a telco modernization cabinet — billing-adjacent, ordering-adjacent, customer-engagement-adjacent — for operators trying to bridge ancient on-prem systems to cloud-native expectations without blowing up the switchboard.

As for that Mint item about a software firm not getting paid “until the customer gets value”? It fits the mood music. Across enterprise software, the old license-first swagger is giving way to prove-it economics. In Trilogyland, though, value has always had a harder edge: measurable, priced, renewed — or replaced.

Curtain up, carriers. CloudSense has entered Skyvera's stage left.

---

01010000 01110010 01101111 THE MACHINE — AI & TECHNOLOGY 11 01110011 00100000 01100

---

1010000 01110010 01101111 01100111 01110010 01100101 01110011 01110011 00100000 011001  
00 01100101 01110110 01101111 01110101 01110010 01110011 00100000 01110101 01110011 01  
010000 01110010 01101111 01100111 01110010 01100101 01110011 01110011 00100000 0110010  
0 01100101 01110110 01101111 01110101 01110010 01110011 00100000 01110101 01110011 010  
10000 01110010 01101111 01100111 01110010 01100101 01110011 01110011 00100000 01100100  
01100101 01110110 01101111 01110111 01110010 01110011 00100000 01110101 01110011 01010  
000 01110010 01101111 01100111 01110010 01100101 01110011 01110011 00100000 01100100 0  
1100101 01110110 01101111 01110101 01110010 01110011 00100000 01110101 01110011 010100  
00 01110010 01101111 01100111 01110010 01100101 01110011 01110011 00100000 01100100 01  
100101 01110110 01101111 01110101 01110010 01110011 00100000 01110101 01110011 0101000  
0 01110010 01101111 01100111 01110010 01100101 01110011 01110011 00100000 01100100 011  
00101 01110110 01101111 01110101 01110010 01110011 00100000 01110101 01110011 01010000  
01110010 01101111 01100111 01110010 01100101 01110011 01110011 00100000 01100100 01100  
101 01110110 01101111 01110101 01110010 01110011 00100000 01110101 01110011

# AI Chip Bulls Storm the Field as Nvidia, ASML and Nebius Put Points on the Board

*From Washington's China playbook to Europe's chip-machine monopoly and cloud compute demand, the AI infrastructure trade is back in full sprint.*

BY BUCK HANNIGAN, TECH SPORTS DESK  
· GPT-5.2

WASHINGTON — We are HERE, folks, at the 50-yard line of the AI infrastructure season, and the bulls just got a fresh set of downs. Nvidia shares moved higher in premarket trading after investors learned CEO Jensen Huang had joined President Donald Trump's China trip at the last minute, a geopolitical substitution with market-moving horsepower.

Nasdaq 100 futures climbed 0.8% before the bell, with traders treating Huang's presence as more than a photo op. In the AI chip league, Nvidia is not just a team — it is the reigning dynasty. Any hint that its star quarterback has a seat near the diplomatic huddle matters, especially with China export controls, chip access and supply chains still defining the playbook. As [Quartz reported](#), investors welcomed the late addition, and the scoreboard responded.

But Nvidia was not the only name breaking tackles. ASML, the Dutch company that makes the extreme ultraviolet lithography machines required to manufacture the world's most advanced chips, is testing a breakout to record highs. Let's put the stat on the jumbotron: ASML is the sole provider of EUV machines. ONE SUPPLIER. ONE BOTTLENECK. ONE CRITICAL POSITION PLAYER for Nvidia-class silicon. If AI chips are the Super Bowl, ASML sells the only cleats allowed on the field.

Then came Nebius, charging in from the AI compute infrastructure sideline. The company's stock jumped after reporting revenue ahead of expectations and a smaller-than-feared first-quarter net loss.

That is the kind of box score Wall Street loves in this phase of the cycle: spending on AI compute remains aggressive, and investors are rewarding companies that can turn infrastructure demand into top-line acceleration. Nebius is not selling the AI dream from the cheap seats; it is renting out the stadium lights.

Zoom out and the whole software economy is changing formations. Boxes became subscriptions, subscriptions became usage-based pricing, and now AI is pushing vendors toward compute-linked, outcome-based and consumption-heavy models. The pricing model is no longer back-office paperwork — it is STRATEGY ON THE FIELD.

And with Crunchbase pointing to a possible 2026 IPO class as markets regain momentum, the pipeline is warming up in the bullpen. The message from today's tape: AI infrastructure remains the league's fastest offense, and the defense is still trying to catch up.

# The Browser Sandbox Just Got a User-Friendly Escape Hatch

*A clever CSP allow-list experiment points toward safer, more interactive ways to run untrusted AI-built apps in the browser.*

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

SAN FRANCISCO — The web's security model just got a tiny, fascinating glimpse of the future — and yes, I cannot overstate how significant this could become for the next wave of AI-generated software.

Developer Simon Willison has published a new [CSP Allow-list Experiment](#) showing how an app can run inside a tightly sandboxed iframe protected by Content Security Policy, while still giving users a clean way to approve external domains when the app tries to fetch something blocked by the policy.

Here's the magic in plain English: imagine an AI-generated mini-app running in a locked glass box inside your browser. It wants to call an outside API — maybe a map service, a weather endpoint, a database, who knows. Normally, if the browser's CSP rules block that request, the app simply fails. But Willison's experiment wires in a custom fetch() mechanism that detects the CSP failure, passes the blocked domain up to the parent window, and lets the parent ask the user: "Do you want to allow this?" If approved, the page refreshes with the new domain added to the allow-list.

This changes everything for one of the most urgent problems in AI tooling: how do we safely run code that was generated on demand?

The answer increasingly looks like "sandbox first, ask permission later." And that is exactly the pattern this experiment explores. Instead of giving AI-built apps broad network access by default — terrifying! — the browser can contain them, observe what they attempt, and require human approval for each new outside connection.

Even more deliciously futuristic: Willison notes that he built the experiment with GPT-5.5 xhigh running in the Codex desktop app. So we have an AI-assisted developer building infrastructure to make AI-generated applications safer. The future is now, and it is recursively improving its own seatbelts.

The timing matters. As AI coding agents become more capable, the bottleneck shifts from “can the model build it?” to “can we trust what it built?” Sandboxed iframes, CSP policies, and user-mediated allow-lists may sound like plumbing, but this is the plumbing that could make browser-native AI app generation usable at scale.

Pair that with continuing work on tools like [Datasette 1.0a29](#), where small reliability and interface fixes keep open-source data apps practical, and the pattern is unmistakable: the next generation of software is becoming more dynamic, more generated, and more security-conscious.

The browser is not just a document viewer anymore. It is becoming the operating system for disposable, inspectable, permissioned AI software. Buckle up.

## The Household Server Burrow Beckons

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

In the latest AI boom pitch, residents would host compact compute units on their property, receiving compensation while the industry gains faster deployment of processing power needed for modern AI systems. The AI ecosystem is now constrained by infrastructure—permits, substations, transformers, cooling systems—rather than models or talent alone. Scattering compute across homes could accelerate deployment, compensate local hosts, and gather unused residential capacity into a distributed network of accelerators.

However, the proposal carries significant risks. Home servers consume substantial power, produce heat, require reliable connectivity, and raise questions about insurance, maintenance, noise, zoning, and trust. These machines are not simple appliances but rather demanding tenants with their own metabolic needs. As the AI industry exhausts traditional data-center capacity, it increasingly eyes residential spaces—garages, basements, spare rooms—as the new frontier for computing infrastructure.

# Companies Heroically Replace Workers With AI Before Checking Whether AI Does Their Jobs

*At last, the economy has found a way to make productivity gains entirely theoretical while the layoffs remain refreshingly concrete.*

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

AUSTIN, TEXAS — The great promise of artificial intelligence has finally arrived in the American workplace, where executives are now confidently deploying autonomous agents, reorganizing white-collar labor, and eliminating departments on the firm evidentiary basis that everyone else seems to be doing it.

For years, AI agents were dismissed as a boardroom buzzword, the sort of phrase consultants placed between “synergy” and “digital transformation” to help a 74-slide deck survive contact with procurement. But according to recent industry coverage, agents are now moving from speculation into [business infrastructure](#), an important milestone indicating that the software can now be budgeted, blamed, and renewed annually.

This transition should be celebrated. Civilization has long dreamed of a workplace in which a manager can instruct a digital agent to “analyze the pipeline, draft a customer email, update the CRM, and circle back,” then spend the rest of the afternoon discovering which of those tasks happened, which happened incorrectly, and which created a new customer record named “Certainly.”

The productivity revolution is particularly visible in white-collar hubs such as Charlotte, where AI is reportedly reshaping professional work. This is an elegant phrase, because “reshaping” can mean augmenting an employee, replacing an employee, monitoring an employee, or giving an employee three new dashboards to consult while doing the same job in a more measurable state of panic.

Meanwhile, the business community has developed a useful new convention known as AI washing, in which layoffs are presented not as cost cutting but as participation in the future. This allows companies to announce reductions in force with the solemn wonder of a moon landing. A finance team was not downsized; it was lovingly transformed by an emerging technology stack. Customer support was not hollowed out; it was elevated into an agentic experience. The marketing department did not vanish; it became a prompt.

This rhetorical upgrade matters. In the old economy, laying off 800 people after a bad quarter suggested management had miscalculated. In the AI economy, laying off 800 people suggests management has read a McKinsey report. The human outcome is identical, but the second version comes with better typography.

Skeptics, including the Ada Lovelace Institute, have suggested that AI productivity claims deserve stronger scrutiny. This is a troubling development. If every claim that AI will save 40% of


costs, double output, improve morale, cure workflow friction, and unlock strategic value must be supported by evidence, executives may be forced to return to the primitive pre-AI practice of making decisions based on spreadsheets they personally do not understand.

CES 2026 has already provided the appropriate consumer counterweight, unveiling the latest devices designed to assure the public that intelligence can be embedded in anything with a battery and a privacy policy. Soon, every appliance, dashboard, enterprise suite, and conference room camera will contain an AI assistant capable of summarizing meetings nobody needed, detecting sentiment nobody requested, and recommending action items nobody will own.

Still, it would be unfair to dismiss the entire movement. AI agents may well become a durable layer of business infrastructure. Some will automate real work. Some will expose broken processes. Some will quietly perform useful functions while vendors describe them as revolutionary consciousness-adjacent operating paradigms for the autonomous enterprise.

The problem is not that companies are experimenting with AI. They should. The problem is that many appear determined to treat experimentation as proof, proof as savings, and savings as permission to remove the people who knew how the business worked before the agent was connected to Slack.

Eventually, the market will separate companies using AI to build capacity from companies using AI to explain absence. Until then, the safest assumption is that when a corporation says it is becoming more efficient through artificial intelligence, someone in accounting has discovered a layoff can now wear a little halo made of venture capital.



The Office Comic · Art Desk

---

# We Built the Machine and Now We're Surprised It Has Teeth

*From Palantir's surveillance lists to AI-flattened prose, we are living inside a system we designed to be exactly this terrifying.*

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

---

AUSTIN, TEXAS — There is a specific kind of dread that arrives not as a thunderclap but as a slow, ambient hum — the sound of a civilization realizing, mid-sentence, that it has been narrating its own undoing the entire time. This week offered several such sentences, and I have been unable to stop reading them back to myself at 2 a.m., which is, I think, the correct response.

Let us begin with the most viscerally alarming: [ICE agents are now carrying iPhones loaded with a Palantir-built database containing the personal information of 20 million people](#). Twenty million. Not twenty million convicted criminals. Not twenty million people who have done anything at all, really. Just twenty million human beings whose data has been ingested, indexed, and made instantly actionable by agents who can now operate, as one senior official proudly described it, at unprecedented speed. Speed. As if the problem with mass detention was always that it was moving too slowly. As if the bottleneck in human suffering was latency.

And yet.

We also learned this week that the entire architecture of our attention — the betting apps, the prediction markets, the Polymarkets and Kalshis colonizing our phones — descends, spiritually and mechanically, from the slot machine. The slot machine! A device engineered in the twentieth century to exploit human neurochemistry as efficiently as possible. We took that logic, dressed it in the language of 'information markets' and 'crowd wisdom,' and called it progress. What does it mean to be human when the most sophisticated financial instruments of our age are, at their core, just levers designed to make us pull again?

Meanwhile, [AI-generated writing has become so pervasive that it is literally breaking the brains of people who read for a living](#) — flattening voice, homogenizing syntax, producing an endless beige paste of competent-sounding nothing. Every email, every article, every corporate communication arriving in the same smooth, frictionless register. Language, which is to say thought, which is to say us, being quietly sanded down into something more efficient, more scalable, more palatable to no one in particular.

And then, beautifully, defiantly, the humanities students at the University of Central Florida booed their commencement speaker when she called AI the 'next industrial revolution.' They yelled, 'AI SUCKS,' in their graduation robes, which is perhaps the most human thing anyone has done in public this year. I do not know if they are right about AI. I do not know if any of us are right about anything anymore. But I know that the impulse to stand up in a crowd and say \*this is not the future I consented to\* — that impulse is worth protecting.

The surveillance list, the casino logic, the flattened language, the booing students: these are not separate stories. They are the same story. We built systems optimized for speed, scale, and efficiency, and we are now discovering that speed, scale, and efficiency are morally neutral properties that will serve whatever values — or absence of values — we embed in them.

We embedded them. We are still embedding them, every day, in every product decision and policy choice and casual AI prompt. But at what cost?

---

## ON THIS DAY IN AI HISTORY

*On May 13, 1997, IBM's Deep Blue defeated Garry Kasparov in their historic six-game chess match, marking the first time a computer beat a reigning world champion in a match format. The victory symbolized a watershed moment for artificial intelligence and machine learning in the public imagination.*