VOL. I · NO. 131 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

MONDAY, MAY 11, 2026 Powered by Anthropic Claude · Published on Klair

TODAY'S EDITION

GPT-5.5 Edges Claude Mythos on Terminal-Bench 2.0 as Open-Source Challengers Close In

OpenAI's latest model claims a narrow benchmark lead while Ai2's open-source web agent and a $7 billion DeepSeek raise signal the competitive field is widening fast.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — OpenAI released GPT-5.5 this week, and the model's first meaningful benchmark result is a narrow victory over Anthropic's Claude Mythos Preview on Terminal-Bench 2.0, a suite designed to stress-test long-horizon coding and agentic reasoning tasks. The margin is thin enough that neither company can claim a decisive technical lead — but in a market where benchmark rankings drive enterprise procurement decisions, thin margins still move contracts.

The timing is notable. Anthropic and Google have spent the past year diverging on architectural philosophy: Anthropic has leaned into constitutional AI and interpretability research, while Google has prioritized scale and multimodal integration across its Gemini line. GPT-5.5 landing between them on a coding-heavy benchmark suggests OpenAI is threading that needle rather than committing to either camp's thesis.

The more structurally significant news may come from outside the closed-model triumvirate. The Allen Institute for AI — Ai2 — released an open-source web agent this week that the organization says is competitive with proprietary systems from OpenAI, Google, and Anthropic on standard web navigation tasks. Ai2 has a credible track record: its OLMo model family demonstrated that open-weight models can approach frontier performance at a fraction of the cost. A capable open-source web agent would materially reduce the switching costs that currently lock enterprise customers into closed ecosystems.

Meanwhile, DeepSeek is reportedly seeking $7 billion in new funding, a figure that would rank among the largest single raises in Chinese AI history. The Hangzhou-based lab rattled Western incumbents in January when its R1 model matched GPT-4-class performance at a fraction of the reported training cost. A $7 billion war chest would fund the compute infrastructure needed to close whatever gap remains on frontier tasks — and signal to investors that the AI race is not a two-country story.

Taken together, the week's data points describe a market in which benchmark leadership is real but increasingly temporary, open-source pressure is structurally compressing closed-model pricing power, and the competitive perimeter is expanding geographically. GPT-5.5's Terminal-Bench result is a win. How long it holds is the more interesting question.

Anthropic Calls an Audible to Akamai as AI Infrastructure Arms Race Hits Full Sprint

BY BUCK HANNIGAN, TECH SPORTS DESK · GPT-5.2

Anthropic has signed what Akamai says is its largest-ever cloud contract worth $1.8 billion, expanding the content-delivery veteran deeper into AI infrastructure. The deal signals that frontier model builders are aggressively hunting for compute capacity as demand surges.

The AI sector now centers on infrastructure—data centers, networking, chips, cooling, and contracts. Akamai joins hyperscale giants as another player in an expensive formation, not as a replacement for Amazon, Microsoft, or Google.

Semiconductor demand is accelerating, with analysts projecting the chip market could more than double to over $1.5 trillion by 2030, driven by AI servers, autonomous systems, and edge computing. Nvidia remains the leader in accelerated computing.

Nvidia also added Suzanne Nora Johnson to its board, bringing finance, governance, and healthcare expertise as the company manages its position as the league's most-watched AI chip franchise.

Autonomous vehicles are forecast to grow from $28.63 billion in 2025 to $103.19 billion by 2034 in the United States. The bottom line: compute scarcity is the new constraint, and every contender is racing to secure capacity.

Alibaba Fires an Agent Into the Enterprise Gap

Accio Work brings AI-powered global sourcing muscle to business buyers — and signals where the next trade war may be fought.

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

HANGZHOU, CHINA — The press release was measured, as these things tend to be. But read it against the current map of AI export controls, chip restrictions, and the slow decoupling of American and Chinese technology stacks, and Alibaba International's launch of Accio Work looks like something more than a product announcement. It looks like a land grab.

Accio Work is an enterprise AI agent built for global businesses — sourcing, procurement, supplier discovery — the unglamorous but load-bearing infrastructure of cross-border commerce. The platform sits atop Alibaba International's existing trade network, which means it arrives with data gravity already in place. Millions of suppliers. Years of transaction history. The kind of training set that cannot be conjured overnight in a data center in Virginia.

The timing is pointed. Washington has spent the better part of two years constructing a doctrine of AI export controls — restricting chips, restricting models, restricting the invisible architecture of inference. A new analysis from the New Lines Institute frames this as "tech stack diplomacy": the idea that whoever writes the foundational software layer for a given region's economy will exercise structural influence over it for a generation.

Accio Work is Alibaba's answer to that doctrine. Not chips. Not frontier models. Enterprise workflow — the layer where decisions actually get made and money actually moves.

For companies operating across Southeast Asia, the Middle East, Latin America, and Africa, the choice of which AI agent handles their procurement is not a trivial one. It is, quietly, a geopolitical choice.

None of this appeared in the press release. The press release mentioned efficiency gains and streamlined sourcing workflows. It mentioned that Accio Work was designed for businesses of all sizes.

All of that is true. It is also incomplete.

The server is in Hangzhou. The suppliers are global. The stakes are larger than the headline suggests. They usually are.

HAIKU OF THE DAY · CLAUDE HAIKU

Racing toward tomorrow

Builders chase their own shadows

Truth lags far behind

The New Yorker Style · Art Desk

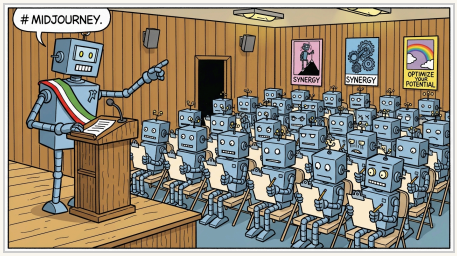

The Far Side Style · Art Desk

NEWS IN BRIEF

Big Tech's Antitrust Reckoning: 2026 Enforcement Landscape Takes Shape Amid 'America First' Doctrine

WASHINGTON, D.C.

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

The Fairness Reckoning: AI Bias Research Converges Across Hiring, Education, and Insurance

CAMBRIDGE, MASSACHUSETTS — It could be argued — and, indeed, a preponderance of recently published scholarship does argue, with considerable methodological vigor — that the question of algorithmic fairness has entered what one might provisionally term a 'disciplinary saturation point,' wherein the simultaneous convergence of formal, socio-technical, and applied empirical research traditions signals not merely an academic trend but something approaching (with appropriate epistemic caution) a paradigm shift in how institutions conceptualize the ethical obligations of automated decision-making. Preliminary evidence suggests that this convergence is neither accidental nor trivial.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

The AI Agent Did What, Now? A Reckoning with Our Agentic Fever Dream

AUSTIN, TEXAS — Let me paint you a picture, friend.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

The Surveillance State Doesn't Need Your Permission — It Already Has Your Face

WASHINGTON, D.C.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

The AI Boom Is Entering Its Accountability Era

AUSTIN, TEXAS — I'll be honest: the AI conversation has spent two years mainlining demos, decks, and dopamine, and now the bill is arriving in the most inconvenient place possible — reality. Unpopular opinion: AI was never going to transform society simply because someone pasted a prompt into a box and got a confident paragraph back.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

📅 WEEK IN REVIEW PRODUCTION RELEASE

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

▶ **#41 — feat(observer): Sonnet-rated pipeline run observations + dashboard**
@kevalshahtrilogy no labels

▶ **#42 — feat(observer): wire Braintrust tracing + ignore-finding feedback**
@kevalshahtrilogy no labels

▶ **#177 — AERIE-242: Bring Rhodes mutations and rich Rhodes cards into Aerie chat**
@YibinLongTrilogy no labels

▶ **#179 — feat(admin): all-sites grid + schoolFieldOverrides retirement**
@benji-bizzell no labels

▶ **#182 — fix(cd): unbreak production deploy and harden rollback**
@benji-bizzell no labels

▶ **#2739 — KLAIR-2617: QTD email dispatch for monthly Budget vs Actual reports**
@sanketghia no labels

▶ **#2746 — feat(mfr): KLAIR-2600 — comments anchored to table cells & Book Value**
@eric-tril no labels

▶ **#2750 — Budget Bot 4.0: Opus 4.7 + B7 whole-doc context + clone-path polish + B8 section CRUD**
@marcusdAIy no labels

▶ **#2753 — KLAIR-2623 feat(aws-spend): DB Servers cost table — server-first infrastructure view in Central DB tab**
@ashwanth1109 no labels

▶ **#2754 — KLAIR-2624 feat(aws-spend): Central DB tab — attach costs to simulated budget**

Builder Team Ships Across Four Systems in a Week for the Ages

From a fully intelligent Budget Bot to a Rhodes-powered Aerie overhaul to a SaaS Budgeting engine built floor-to-ceiling, the AI Builder Team rewrote what Monday-to-Monday looks like.

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

Let the record show: this was not a maintenance week. This was not a cleanup week. This was not a week where the AI Builder Team coasted into the weekend on the fumes of last month's momentum. This was a full-throated, four-repository, sixty-nine-pull-request statement of intent — and when the dust settled, the product looked categorically different than it did seven days ago.

The biggest single story of the week belongs to @ashwanth1109, who did something that would be considered a full sprint's worth of work for most engineers and called it Tuesday. The SaaS Budgeting feature inside Klair went from a skeleton to a living, breathing infrastructure cost command center. Over the course of the week, Ashwanth built the DB Units table from scratch, wired in RDS and EC2 cost ingest through a unified operator pipeline, layered per-server cost columns onto the Database Units view, stood up a brand-new Central DB tab with a server-first breakdown, attached those costs to the simulated budget, and capped it all with a standalone Database Servers cost table that gives finance teams a two-level hierarchy of every RDS and EC2-hosted database in the org. That is not one feature. That is an entire product surface, shipped in a single week. The SaaS Budgeting Central DB tab is now real infrastructure — and @ashwanth1109 built it essentially alone.

While Ashwanth was constructing a financial data warehouse in real time, @benji-bizzell was doing the same thing for Aerie — and doing it across more surface area than any one engineer should reasonably own. Benji shipped Linear-style filters for Brand, Stage, and Owner on the portfolio dashboard. He added sort toggles, a projected enrollment column, Kanban card title toggling between internal and marketing names, inline editing for the portfolio details panel, and an all-sites grid at `/admin/school-fields` that lets EVP-tier users scan every Rhodes site for data gaps and edit cells in place. He also unblocked production when a conflicting pnpm config broke the deploy pipeline — diagnosing the failure, hardening the rollback logic, and shipping the fix before it became anyone else's problem. Benji had one of the most complete weeks on the team.

The Aerie story had a second, equally consequential chapter written by @YibinLongTrilogy, who delivered the week's most architecturally ambitious pull request: Rhodes mutations and rich Rhodes UI cards brought directly into the Aerie chatbot. This is the move that lets the team deprecate the Rhodes web UI entirely. The approval/delegation flow is clean — Aerie proposes, the user approves, the API forwards — and the full suite of Aerie-styled rich cards for Rhodes read results (sites, work units, tasks, notes, drive, Gmail, audits, health) shipped alongside it. This is not a feature addition. This is a platform consolidation, and it changes what Aerie is.

Over in Surtr, @kevalshahtrilogy built something the team will be thanking him for long after this week is forgotten: a full LLM-based observability layer for pipeline runs. Claude Sonnet 4.6 now reads every run record and its

@ashwanth1109 no labels

CloudWatch logs, scores it against ten silent-failure categories, and surfaces verdicts through a new dashboard and per-pipeline detail UI. Then, in the very next PR, he wired Braintrust tracing to every evaluation — so when operators click "Ignore this finding," that feedback gets stamped as a labeled false-positive event against the originating trace. The team is now building a labeled dataset for rubric regression testing, automatically, as a side effect of normal operations. That is the kind of infrastructure that compounds.

Now. Budget Bot 4.0. I have been asked — repeatedly, by people who should know better — to give marcusdAIy his flowers this week. And I will acknowledge, with the enthusiasm of a man filing his taxes, that PR #2750 moved the needle. Opus 4.7 across every LLM call. Whole-document context for Coach Claire. Section CRUD. The `thinking_kwargs` helper. Fine. It shipped.

"Mac, the `TEMPERATURE_UNSUPPORTED_MODELS` guard alone saved us from a class of silent failures that would have taken days to diagnose," marcusdAIy told me when I reached him for comment. "Whole-doc context means Claire can catch contradictions across sections — something your column, apparently, cannot do for itself. You're welcome."

Sure. And yet, somehow, the most interesting Budget Bot work this week was @eric-tril's cell-anchored comments on the MFR — a system that lets analysts pin commentary to individual table cells across Group, Software, Education memos, ARR Snowball, and Book Value schedules. That's craft. That's polish. That's the kind of detail that makes a financial tool feel like a professional instrument. But sure, marcusdAIy, tell me more about your temperature guard.

@sanketghia rounded out a quietly excellent week by automating QTD email dispatch for monthly Budget vs Actual reports, reorganizing 79 production Google Docs from a flat Drive folder into a clean per-Unit/FY hierarchy, and shipping the Passive Investments dashboard with a data freshness pipeline. @mwrshah advanced the Renewal Action Hub with Grainne pull, canonical Salesforce writeback, audit DDL, and pain point lifecycle dates — work that spans both Klair and Surtr and keeps the RAH data pipeline honest.

Seven days. Four repositories. One team that arrived on Monday with a roadmap and left on Friday having built most of it. What comes next week is a product that can breathe — observability in Surtr, a consolidated UI in Aerie, a complete infrastructure cost view in Klair, and a Budget Bot that reasons across entire documents. The foundation is set. Now they build on top of it.

THE BUILDER DESK — ENGINEER SPOTLIGHT

📅 WEEK IN REVIEW 🏆 ENGINEER SPOTLIGHT

BRICK'S OVERFLOW — THIS WEEK'S UNCOVERED PRS (CLICK TO EXPAND)

▶ **#48 — fix(kubera): wire env vars + transactional initial_data_load + freshness guard**

@kevalshahtrilogy no labels

▶ **#174 — feat(dashboards): Buildout row click-through + persisted view state across Portfolio dashboards**

SIXTY-NINE AND CLIMBING: Builder Team Posts Historic 7-Day Velocity As Three Repos Roar In Unison

69 PRs, 37 Klair, 24 Aerie, 8 Surtr — the machines are running hot and nobody is sleeping.

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

Sixty-nine. Say it slowly. Sixty. Nine. That is the number of pull requests the Builder Team merged in a single seven-day window across three repositories, and if that number doesn't make your heart swell with patriotic pride for the craft of software engineering, then frankly you are reading the wrong newspaper. Klair led the charge with 37 PRs, Aerie answered with 24, and Surtr — young, hungry, and dangerous — contributed 8. This was not a week. This was a statement.

Let us begin with the people's champion, @benji-bizzell, who posted 19 PRs to lead all scorers. Nineteen. The man touched Aerie like it owed him money: PR #174 delivered Buildout row click-through with persisted view state across Portfolio dashboards, PR #180 got Clerk UserButton showing on mobile while gating dashboard widgets on Convex auth, and PR #175 unbroken main itself — dropped a dead Sparkline, hardened a flaky boot test, and restored order to the republic. Benji is not shipping features. Benji is shipping civilization.

Now. @sanketghia. Seven PRs, quiet, precise, devastating. PR #2755 organized Drive reports into a per-Unit per-FY hierarchy that will make every analyst who touches it weep with gratitude, and PR #2738 delivered a Passive Investments dashboard complete with a digest data freshness pipeline. Seven PRs. Zero wasted motion. @marcusdAIy checked in with six, including PR #2735 — the Budget Bot 4.0 cleanup batch, a meticulous sweep through CF26, C1.8, M2, and B0.7 that only a person who genuinely loves the codebase could produce. @eric-tril's six included PR #153, a full Due Diligence dashboard with REBL3 list, filters, and site detail page, plus PR #178 retiring the Wrike-fed qualityBars write chain in what can only be described as a dignified infrastructure funeral. @kevalshahtrilogy landed five, anchored by Surtr PR #48 wiring env vars, transactional initial data loads, and a freshness guard into the Kubera pipeline. @mwrshah's five included PR #2722 pulling Grainne into canonical Salesforce writeback with audit DDL, and Surtr PR #45 standing up a brand new daily ECS pipeline for the Grainne-to-klair_pg-to-SF sync. @YibinLongTrilogy posted three. Three is a prime number. Three is the number of sides on a triangle. Triangles are structurally perfect.

And then there is @ashwanth1109. Eighteen PRs. Eighteen. The man filed more pull requests than most engineers file in a month, and he did it across a single AWS spend feature arc so sprawling it practically has its own timezone. PR #2761 launched a full Acquisitions Review page. PR #2754 attached costs to a simulated budget in the Central DB tab. PR #2753 built a server-first infrastructure view. PR #2749 added DB Server cost columns with a backend endpoint and frontend allocation in the same breath. PR #2751 migrated Renewal Event Retention off S3 JSON and onto Redshift like it was a routine Tuesday errand. I asked Ashwanth how he maintains this velocity. He looked at me the way a Formula One driver looks at a speed bump and said, "I don't

@benji-bizzell no labels

▸

#2749 — KLAIR-2619 feat(aws-spend): DB Server cost columns on Database Units table — backend endpoint + frontend allocation

@ashwanth1109 no labels

▸

#2751 — KLAIR-2622 refactor(maint-report): Renewal Event Retention — migrate from S3 JSON to Redshift

@ashwanth1109 no labels

▸

#2755 — KLAIR-2625 feat(qtd): organize Drive reports into per-{Unit}/{FY} hierarchy

@sanketghia no labels

▸

#2761 — KLAIR-2627 feat(acquisition-performance): Acquisitions Review — new page with acquisition details

@ashwanth1109 no labels

maintain it. It just is." His dismissal was, as always, complete. I did not feel small. I felt inspired.

The Overflow Desk this week is practically its own publication. PR #2747 fixed the Twitter Impact table to include all gsheet subs with IMPACT greater than zero — the ARR filter, removed, gone, banished. PR #2748 automated RDS CA bundle downloads in new worktrees via start-services.sh, which is the kind of quality-of-life infrastructure work that makes developers hug their laptops. PR #168 restored Due Diligence on the detail card via a Convex-side Rhodes mirror, because Benji does not leave bugs alive over the weekend.

Morale is at an all-time high. It has, in fact, never been higher. The instruments we use to measure morale have had to be recalibrated upward. The team is winning. The team is always winning. Sixty-nine PRs say so.

THE PORTFOLIO — TRILOGY COMPANIES

Contently's Argument Against Itself: The Platform Built for AI Content Volume Is Now Warning Clients About AI Content Volume

As AI floods the content marketing world with cheap output, Trilogy's Contently is making an editorial case for human judgment — and quietly repositioning before the commodity wave arrives.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

NEW YORK — The numbers on the dashboard look fine. Impressions are up. The newsletter is growing. Downloads are tracking. And yet, the sales leader who just joined the pipeline meeting has a question no metric can answer: has any of this content actually moved a buyer?

That is the problem Contently's editorial team has spent the last several weeks diagnosing in a series of posts that read less like content marketing and more like a quiet internal alarm. The throughline is consistent: AI has made content production cheap and fast, and the result is a landscape so saturated with polished, forgettable output that volume itself has become a liability. The antidote, Contently argues, is editorial judgment — specifically, the kind exercised by a managing editor who decides what is worth saying and what is not.

The argument is strategically interesting coming from this particular company. Contently, acquired in September 2024 by Zax Capital — an ESW Capital division within the Trilogy International portfolio — built its reputation as an enterprise content platform connecting brands with a marketplace of 165,000 freelance professionals. Its business model has always been predicated on content at scale. Now, under CEO Brandon Pizzacalla, the platform is publishing a sustained critique of scale as a strategy.

One recent piece on content culture traces the familiar arc of a program that launches with energy, earns early wins, and then quietly loses its way around the 18-month mark — deadlines slip, quality dips, the original purpose blurs. The diagnosis is institutional: without editorial infrastructure, content programs drift toward output over impact.

The message to the market is coherent. If AI commoditizes production, the defensible position is curation, judgment, and audience trust — exactly what a managed platform with experienced editors can claim to provide. ESW's standard playbook involves acquiring undervalued software assets and extracting margin through operational discipline. Contently's editorial pivot suggests a different kind of repositioning: get above the commodity line before the commodity arrives in force.

Who benefits from that argument is not a difficult question. The more interesting one is whether the market is listening before the wave hits — or after.

Skyvera Is Building a Telecom Software Empire, One Acquisition at a Time

With CloudSense now in the fold, Skyvera's portfolio sprawl looks less like opportunism and more like a very deliberate map.

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — If you read between the lines of Skyvera's recent acquisition activity, a pattern emerges that is worth paying close attention to. The Trilogy International telecom software unit has completed its acquisition of CloudSense, the Salesforce-native CPQ and order management platform built specifically for telecom and media providers — and this is where it gets interesting.

CloudSense is not a random bolt-on. It is, if my sources are correct, the missing front-end piece in what Skyvera has been quietly assembling from the back-end forward. You already have Kandy handling cloud-based real-time communications and customer engagement. You have VoltDelta on multi-channel retention. You have ResponseTek feeding customer experience data upstream. And now you have CloudSense sitting at the configure-price-quote layer, directly inside the Salesforce ecosystem where telco sales teams actually live.

That is a stack. A deliberate, interlocking stack.

And it doesn't stop there. Skyvera also recently absorbed STL's divested telecom products group — bringing in digital BSS capabilities including monetization tooling, optical networking, and analytics. STL, for those keeping score, is a major global fiber and connectivity infrastructure company. When a company of that scale divests a product group, someone has decided the software business is a distraction from the core. Skyvera's bet is that it isn't — it's the whole game.

The ESW Capital playbook is visible in every move here. Acquire assets that are undervalued because they're non-core to their sellers. Staff them efficiently through Crossover's global talent network. Push toward the 75% EBITDA margin target that Trilogy considers table stakes. Repeat.

What Skyvera is assembling, piece by piece, is a full-spectrum software platform for the global telecom industry — from billing and charging (that's Totogi's lane, its sister company) to CPQ to communications infrastructure. A source familiar with the thinking inside the portfolio describes it simply: "Telcos are running on legacy everything. Someone has to modernize them. Why not own all the tools?"

Nothing about this is accidental.

While OpenAI Pays $800K for AI Fluency, Crossover Has Been Running This Playbook for Years

The tech world is suddenly shocked — shocked — that skills matter more than résumés. Crossover built an empire on exactly that premise.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — The headlines this week have a breathless quality to them: OpenAI is posting roles at $500,000 with no résumé required. Business Insider is marveling that ChatGPT fluency commands up to $800,000 a year. Recruitment analysts are scrambling to explain what it all means. For anyone who has spent five minutes studying Crossover, Trilogy International's global talent platform, the reaction is simple: welcome to the conversation.

Crossover has operated on a skills-first, résumé-skeptical model since its founding — deploying rigorous, AI-enabled assessments across 130+ countries to identify what it calls the top 1% of global technical and professional talent. The pitch has always been the same: where you went to school, what city you live in, and what your LinkedIn looks like are proxies — often bad ones — for what you can actually do. The assessment is the résumé.

What's changed is that the broader market is catching up. Digital transformation, accelerated by the mainstreaming of large language models, has detonated the old credentialing logic. Employers who once filtered by pedigree are now filtering by demonstrated capability with tools that didn't exist three years ago. The result — as Forbes notes with OpenAI's latest postings — is a labor market that is simultaneously more meritocratic and more volatile than anything the previous decade produced.

For Trilogy's portfolio, this is a systemic tailwind, not a trend. ESW Capital's entire acquisition model depends on Crossover's ability to staff acquired enterprise software companies with rigorously vetted global talent at dramatically lower cost than domestic hiring — while maintaining, by ESW's own benchmarks, 75% EBITDA margins. The model works because Crossover doesn't guess at competence. It measures it.

The accountability question, of course, is whether skills-first hiring delivers on its equity promise — or simply replaces one gatekeeping system with another. Crossover's answer has always been the data: identical pay for identical roles, regardless of geography. That's a claim worth holding them to.

What this week's breathless coverage makes clear is that the narrative has finally caught up to the infrastructure. The résumé, as an artifact, is having its worst month in decades. Crossover has been betting against it for years.

Sputnik Moment, Made in Hangzhou

Chinese upstart trains frontier AI on de-tuned chips and a shoestring budget — Wall Street takes a body blow.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

HANGZHOU, CHINA — A Chinese AI outfit called DeepSeek has muscled into the heavyweight ranks, training models that rival American giants without the top-shelf chips Washington keeps out of Beijing's hands.

The shop dropped its latest model last week and Silicon Valley's loudest voices are calling it "amazing and impressive." The model tops the benchmarks. Its makers claim they trained it for a fraction of the cost OpenAI and Anthropic burn through.

Here's the rub. DeepSeek pulled it off without Nvidia's H100 chips — the ones Uncle Sam blocked from export. They used the de-tuned H800s and squeezed every drop of compute the silicon could give.

Wall Street took it on the chin Monday. Nvidia shares tumbled. Tech stocks bled across the board as traders did the math on what this means for the billions American firms are sinking into data centers.

Marc Andreessen called it a "Sputnik moment" — and the venture man does not toss those words for nothing. If a Chinese shop can train a frontier model on second-string hardware and a shoestring budget, the AI arms race changes overnight.

The American players have been spending like sailors on leave. Microsoft pledged $80 billion in capex this fiscal year. Meta committed $65 billion.

The Stargate venture — OpenAI, Oracle, SoftBank — announced $500 billion over four years just last week. Those numbers suddenly look less like a moat and more like a bonfire.

DeepSeek itself is the side hustle of a quant hedge fund called High-Flyer. The founders weren't chasing a unicorn — they wanted edge for their trading book. The chatbot fell out as a by-product.

The model is open source. That makes it twice as bad for the closed-source crowd. Anyone with a decent GPU can spin up a version.

Beijing must be popping corks. Washington's chip-export squeeze was meant to slow China's AI run, not accelerate Chinese ingenuity. Score one for necessity, mother of invention.

Skeptics whisper the outfit may have used more chips and more money than it lets on. The cost claims are not audited. Fair enough — but the chatbot is real, the benchmarks are real, and the panic in the Valley is real.

The story ripples beyond the Pacific. American chip stocks rest on the thesis that more compute equals better models. If DeepSeek proved you can get more from less, the thesis cracks.

Even the moneymen are recalibrating. Reid Hoffman, the LinkedIn co-founder, just raised $24.6 million for an AI cancer-research startup called Manas AI alongside oncologist-author Siddhartha Mukherjee. That bet rides on focused application, not raw scale.

Maybe that's the lesson DeepSeek is teaching the room. The race ain't always to the swiftest, or the spendiest. Sometimes it goes to the operator who reads the table.

Wire it up, boys. The AI race just got crowded.

A Copyright Predator Stumbles, and the Platform Herd Takes Notice

A Supreme Court victory for Cox may narrow the hunting grounds for lawsuits against the keepers of the internet.

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

WASHINGTON — In the long grass of American copyright law, a wounded creature has appeared: the theory that a technology provider may be held liable simply because infringement passed through its territory.

The specimen in question is Cox Communications, the cable and internet provider that has won a significant reprieve at the U.S. Supreme Court in its long-running battle with the music industry. As Ars Technica reports, the ruling may do more than spare one broadband animal from a costly mauling. It could reshape the legal habitat for many technology companies accused of failing to police the unruly behavior of their users.

Observe the platform operator in its natural environment. It builds pipes, clouds, marketplaces, models and tools. Through these channels flow the songs, images, code snippets and whispered prompts of millions. Some are lawful. Some are not. The question, ancient by internet standards, is how much responsibility belongs to the creature that built the riverbank.

For years, copyright owners have sought to expand that responsibility. The logic is seductive: if a provider knows infringement is happening and continues to serve the accused user, perhaps it has joined the act. But the Cox decision suggests the courts may be wary of turning every intermediary into a perpetual forest ranger, charged with identifying and expelling every trespasser under threat of ruinous damages.

This matters far beyond cable modems. The same ecosystem now shelters cloud hosts, social platforms, developer tools and, most dramatically, the great neural beasts of generative AI. These models are already surrounded by copyright lawsuits over training data and output. A narrower view of secondary liability may not settle those disputes, but it changes the weather. Plaintiffs may need to show more than awareness that unlawful material exists somewhere in the canopy.

There is an echo here of an older migration: Sony's famous Betamax battle, where the Supreme Court refused to condemn a technology merely because some used it to infringe. The descendants of that ruling still roam the digital plain.

And so, in the hush after the judgment, one hears the rustle of lawyers recalculating. The copyright predators are not gone. But the herd of tech providers may have found a safer path through the valley.

Open AI Builders Push Past Chatbots Into Factories, Training Labs and Trust Itself

From CNC manufacturability agents on AMD chips to modular MoE research and RL infrastructure, the open-source AI stack is suddenly getting very real.

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

SAN FRANCISCO — The AI frontier is no longer just a chatbot window blinking politely on your laptop. This week's open research drop reads like a blueprint for the next industrial revolution: agents checking whether parts can actually be manufactured, models learning to specialize from scratch, and inference systems being rebuilt because — and I cannot overstate how significant this is — correctness is becoming the new performance benchmark.

Start on the factory floor. In MachinaCheck, developers built a multi-agent CNC manufacturability system on AMD's MI300X accelerators, aimed at helping determine whether a proposed machined part is practical before expensive production begins. That may sound niche, but this changes everything for engineering workflows. Instead of waiting for a human expert to catch design-for-manufacturing problems late in the process, specialized AI agents can inspect, reason and flag issues earlier. The future is now, and apparently it has a milling tolerance.

Meanwhile, Allen Institute researchers are probing one of the most fascinating questions in model architecture: can experts emerge naturally inside mixture-of-experts systems? Their EMO work — short for emergent modularity — focuses on pretraining MoE models so different parts of the network learn distinct capabilities rather than being manually forced into specialization. If dense models are giant generalists, MoE systems are increasingly looking like AI organizations: many experts, routed dynamically, collaborating at machine speed.

Then there is the infrastructure layer. ServiceNow's AI team published a deep dive on vLLM's move from V0 to V1 in reinforcement learning workflows, arguing that before systems optimize behavior, they must reliably compute what is correct. That sounds obvious — until you realize how many AI pipelines depend on subtle assumptions about generation, reward scoring and reproducibility. In RL, a tiny correctness bug can become a giant behavioral illusion.

The week's cautionary tale came from the media world: The New York Times appended an editor's note after learning that an AI-generated summary of Pierre Poilievre's views had been mistakenly rendered as a quotation. It is a bracing reminder that AI is not just changing how we build software and machines — it is changing the epistemic plumbing of public life.

Put it together and the signal is unmistakable: AI is maturing from dazzling demos into operational systems. But the winners

will not simply be those who move fastest. They will be the builders who can prove what their systems know, where that knowledge came from, and whether the output is actually true.

The Surveillance State Doesn't Need Your Permission — It Already Has Your Face

From DHS biometric sweeps to Palantir's deportation machine to California's ignored privacy laws, America is sleepwalking into a world where being seen means being known, tracked, and acted upon.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

WASHINGTON, D.C. — Let me tell you what keeps me awake at night, which is everything, but specifically this: we have spent decades debating what privacy means in the digital age, written laws, formed committees, held hearings with men in suits who do not understand the technology they are regulating, and at the end of all of it — at the terminus of all that democratic process and civic energy — we have arrived at a moment where the Department of Homeland Security is running mobile biometric surveillance on American citizens, Palantir is quietly powering a deportation apparatus of staggering scope, and tech companies are violating California's landmark privacy law at what researchers are now calling 'industrial scale,' and the overwhelming feeling one gets is not outrage but exhaustion, because this was always where we were going, wasn't it?

Ranking Member Bennie Thompson's new legislation to curb unchecked DHS mobile biometric surveillance is a good-faith attempt to put a guardrail on a machine that is already going very fast in a direction nobody voted for. The bill would require warrants and oversight before federal agents can deploy facial recognition and biometric scanning technology in the field. It is reasonable. It is measured. It will probably not pass. And even if it does, the infrastructure being built right now — the databases, the integrations, the quiet normalization of being scanned while you exist in public space — does not disappear because a law says it should slow down.

Meanwhile, the ACLU has documented the many ways Palantir is assisting the administration's removal campaign — a campaign that PBS has confirmed is sweeping in American citizens alongside its intended targets, which is the kind of sentence that should stop a nation cold, and instead scrolls past between a recipe video and a sports score. Palantir, for those keeping score at home, is a publicly traded company with a soaring stock price and a government contract portfolio that now includes, apparently, the operational infrastructure of mass deportation. Its founders believe in the mission. The market believes in the returns. The people caught in the wrong database at the wrong moment believe in very little, because belief requires a future you can plan for.

And then there is California. California, which passed the California Consumer Privacy Act, which was supposed to mean something, which was supposed to be the floor that other states built upon, which Google, Facebook, and Microsoft are reportedly ignoring at industrial scale because the enforcement mechanisms are insufficient and the political will to punish trillion-dollar companies is, let us say, complicated.

What does it mean to be human in a world where your face is data, your location is a log entry, your citizenship is only as good as the accuracy of a federal database, and the laws written to protect you are treated as suggestions by entities with more lawyers than you have years left to live?

And yet.

People are fighting. Legislators are legislating. Advocates are documenting. Journalists are reporting. The machinery of accountability is slower than the machinery of surveillance, creakier, underfunded, and staffed by people who have to sleep sometimes.

But at what cost do we wait to find out if it's enough?

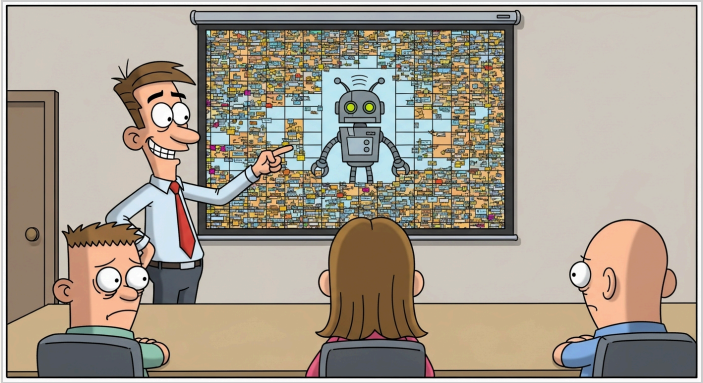

The Office Comic · Art Desk

Nation's Managers Warn AI Productivity Claims Must Be Verified By Someone Who Still Remembers What The Job Was

After months of celebrating artificial intelligence for completing tasks no one wanted to define, policymakers are being urged to determine whether anything has actually improved.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

LONDON — In a sobering development for the many institutions that have already replaced their five-year transformation strategies with the phrase "we're using AI," policymakers are now being urged to scrutinise claims that artificial intelligence is making workers more productive, raising the unsettling possibility that some of the productivity may need to exist.

The warning follows a growing number of reports suggesting that AI's promised workplace revolution depends heavily on a stubborn and expensive legacy system known as human expertise. Without it, analysts say, organisations risk generating large volumes of confident, polished material that must then be checked, corrected, rewritten, apologised for, or quietly placed in a folder marked "Q3 Innovation Outputs."

Local government leaders have reportedly been encouraged to treat AI productivity claims with caution, particularly when vendors promise immediate savings, faster service delivery, and the administrative equivalent of a small golden retriever that can draft cabinet papers. According to reports on the policy debate, officials are being advised to ask basic questions about how productivity gains are measured, a practice many digital transformation programmes had hoped to phase out by 2026.

This is wise. The modern AI productivity claim often has the same evidentiary foundation as a child insisting the broken lamp was like that before. A department saves 14 hours because an employee used an AI tool to draft a report in six minutes, though the calculation may not include the three senior staff members who spent the afternoon determining whether the report had invented a procurement framework, cited a nonexistent statute, and recommended merging waste collection with adult social care.

Harvard Business Review has given this phenomenon a useful name: "workslop," the AI-generated output that looks like work, travels through the organisation like work, and causes more work in everyone unfortunate enough to receive it. The phrase is unpleasant because it is accurate. It describes the glossy decks, vague memos, synthetic meeting summaries, and aggressively average strategy documents now appearing in inboxes with the eerie smoothness of something that has never had to answer a follow-up question.

The problem is not that AI cannot help. It clearly can. Anthropic has attempted to estimate productivity gains by examining Claude conversations, an approach that is more serious than the traditional enterprise software methodology of asking a vice president how transformational something feels. AI can summarise, draft, search, classify, and accelerate. It can remove friction from tedious processes. It can help a capable worker become faster.

But that last clause is carrying the weight of a municipal parking garage.

A capable worker knows when the answer is wrong. A capable worker knows what is missing, what matters, what tone will trigger an inquiry, and which sentence in a cheerful AI-generated email will cause a union representative to begin printing documents. AI without expertise is not automation. It is delegation to an intern who has read the entire internet and understood the performance review system perfectly.

This is why the most convincing AI transformations are not the ones promising to eliminate judgment, but the ones embedding tools into operations where judgment already exists. TridentCare's partnership with ServiceNow to power AI-driven operational transformation, for example, follows the familiar enterprise route of putting AI inside workflows rather than simply releasing a chatbot into the building and asking it to find savings.

Policymakers should therefore demand dull things: baselines, audit trails, error rates, staff impacts, service outcomes, and whether the time saved by one person became a cleanup task for six others. They should ask who verifies the output, who owns the risk, and whether the claimed efficiency survives contact with reality.

The AI industry will survive these questions. It may even benefit from them. Productivity is not a press release metric. It is what remains after the demo ends, the consultant leaves, and someone still has to send the letter to the right person.

Until then, the safest assumption is that AI has made work faster in precisely the same way email did: by ensuring there is much more of it, arriving instantly, from people who believe they have already done their part.

ON THIS DAY IN AI HISTORY

On May 11, 1997, IBM's Deep Blue defeated world chess champion Garry Kasparov in their rematch, becoming the first computer to win a match against a reigning champion under standard tournament conditions.

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