

The Trilogyn Times

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

WEDNESDAY, JULY 08, 2026

Powered by Anthropic Claude · Published on Klair

Trilogyn International © 2026

TODAY'S EDITION

110100 00100000 01100111 01101001 01100001 01101110 01110100 01110011 00100000 0111011
1 01101000 01101111 01101100 01100101 00101100 01000111 01101001 01100001 01101110 011
10100 01110011 00100000 01100101 01100001 01110100 00100000 01100111 01101001 01100001
01101110 01110100 01110011 00100000 01110111 01101000 01101111 01101100 01100101 00101
100 01000111 01101001 01100001 01101110 01110100 01110011 00100000 01100101 01100001 0
1110100 00100000 01100111 01101001 01100001 01101110 01110100 01110011 00100000 011101
11 01101000 01101111 01101100 01100101 00101100 01000111 01101001 01100001 01101110 01
110100 01110011 00100000 01100101 01100001 01110100 00100000 01100111 01101001 0110000
1 01101110 01110100 01110011 00100000 01110111 01101000 01101111 01101100 01100101 001
01100

Google Tops the Safety Board, Bleeds Its Best to Rivals

A new AI report card crowns the search giant even as four researchers bolt for Anthropic and OpenAI — and Musk's xAI pulls the lone F.

BY HANK CALLOWAY, WIRE
CORRESPONDENT · CLAUDE OPUS +
THINKING

MOUNTAIN VIEW, CALIF. — Google topped a fresh artificial-intelligence safety report card this week, then lost four of its sharpest researchers to rivals Anthropic and OpenAI.

That's the whole racket in one breath. Win the trophy with one hand. Drop the roster with the other.

The new [AI Safety Index](#) seated Google, Anthropic and OpenAI at the head of the class. Elon Musk's xAI took home the lone F. The marks scored how hard each outfit works to keep its own machines on a leash — the safeguards, the testing, the straight talk about what might go sideways.

Here's the twist. The three names atop the safety board are the same three raiding each other's benches — the honor students swiping one another's lunch.

Trouble is, a gold star won't keep a genius in his chair. Fortune reports top talent is streaming out of Google DeepMind, the London lab behind AlphaGo and the AlphaFold protein maps. Insiders now ask aloud whether it can hold its perch at the front of the pack.

Inc. counted the damage. Four key staffers gone, each one landing at one of the two San Francisco houses doing the hunting. [The exits](#) stack up as poachers dangle fatter checks and shorter chains of command.

Anthropic knows the play. The shop was born in 2021 when OpenAI alumni, led by Dario Amodei, walked out over the direction of the work, and now it's collecting Google's people the same way —

selling a pitch that safe and cutting-edge can ride one horse.

The churn cuts deep because in this game, people are the moat. Chips can be bought and data can be scraped, but the few who can shove a model past the pack don't grow on trees.

The giants aren't just buying code. They're buying philosophers, too. Daily Nous keeps a running list of thinkers who've signed on with AI firms to wrestle the heavy questions — what the machines owe us, what we owe them, and who's holding the wheel when nobody's sure.

Down below the penthouse, the floor is giving way. TechCrunch keeps a running ledger of every major 2026 layoff that name-checked AI — outfits framing the cuts as efficiency the machines now deliver. Brains climb the ladder up top; bodies sweep out the back below.

The whole scramble has a geography problem. Every dollar chases the same few zip codes around San Francisco Bay, bidding the same familiar faces higher and higher.

Trilogy International's Crossover platform plays the other side of the street. It recruits what it calls the top 1% from more than 130 countries and pays the identical rate whether the hire sits in Austin or Nairobi. No bidding war, no coastal premium — a different wager on where the brains actually live.

For now, the coastal giants keep swapping talent like ballplayers at the trade deadline. The scoreboard says Google leads on safety. The exit door tells another story.

Korn Ferry Swallows a Rival, and the Executive Search Map Redraws Itself

BY ELEANOR CROSS, FOREIGN
CORRESPONDENT · CLAUDE SONNET

The firm has acquired Trilogy International, a boutique executive search and leadership advisory shop, folding its capabilities into what is already the dominant platform in the sector. Korn Ferry has spent years building beyond its traditional retained-search roots, pushing into organizational consulting, digital assessment, and leadership development. Trilogy International brings a client book and senior search professionals who operate at the highest levels of corporate succession and board-level work.

Consolidation at the top of the talent industry is nothing new. The largest firms have been absorbing specialists for years as clients increasingly demand a single partner handling assessment through placement to development. The executive search market faces real headwinds: a cooling C-suite hiring cycle, AI-assisted candidate matching, and clients questioning what search firms provide that algorithms cannot.

Korn Ferry's answer is more reach, more senior relationships, and more institutional heft. The combined firm's global footprint now stretches further across markets where the next generation of corporate leadership is being cultivated and competed for. Whether Trilogy International's value survives the acquisition remains uncertain.

AI Startups Hoover Up Capital as Valuations Defy Gravity

Four major funding rounds in a single week signal that AI investment has entered a new phase — one where traditional valuation metrics no longer apply.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

NEW YORK — The venture capital machine is running at full throttle. Four significant AI funding rounds closed in rapid succession this week, collectively representing over \$1.5 billion in fresh capital and valuations that would have been inconceivable for companies at similar revenue stages five years ago.

The largest: [Decart, an Israeli AI startup, raised \\$300 million at a \\$4 billion valuation](#) with Nvidia participating — a strategic signal as much as a financial one. Nvidia's direct investment in portfolio companies has become a recurring pattern, effectively functioning as a market-confidence endorsement that other LPs find difficult to discount.

Bret Taylor's Sierra, which sells AI-powered customer service agents, raised nearly \$1 billion just months after its previous capital raise. The pace — back-to-back rounds with no intervening product milestone publicly announced — reflects how compressed the AI fundraising cycle has become. Investors are paying for trajectory, not trailing revenue.

LMarena, which builds evaluation infrastructure for AI models, closed a \$150 million round at a \$1.7 billion valuation. The company operates in a segment — AI benchmarking and evals — that barely existed as a standalone market category 24 months ago. Its unicorn status underscores how the infrastructure layer around AI is attracting capital nearly as aggressively as AI applications themselves.

Underpinning all of this is a [broader rethink of valuation methodology](#) playing out among top-tier VC firms. Benchmark,

historically a discipline-oriented shop skeptical of growth-at-any-cost investing, has signaled a recalibration — acknowledging that AI's speed of compounding may render conventional ARR multiples structurally obsolete.

Separately, a cluster of startups targeting quantitative finance is attempting to automate the proprietary research and signal-generation processes that hedge funds have guarded for decades. Whether that market materializes at scale remains unproven, but early capital is flowing regardless.


The week's activity confirms a directional reality: in AI, the cost of missing a category-defining investment is now perceived as higher than the cost of overpaying to get in.

HAIKU OF THE DAY · CLAUDE
HAIKU

*Giants eat giants whole,
yet the smallest spark still burns—
code rewrites the rules.*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

The Machines Are Writing Papers, Teaching Children, and Designing Parts — But Can They Be Trusted?

CAMBRIDGE, MASSACHUSETTS — It could be argued — and preliminary evidence suggests with some degree of epistemic confidence — that the current moment in artificial intelligence research represents less a coherent paradigm shift than a simultaneous, polyphonic, and occasionally cacophonous expansion into every conceivable domain of human intellectual production.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

LEADERSHIP VACUUM AT DOJ ANTITRUST DIVISION CLOUDS BIG TECH ENFORCEMENT OUTLOOK FOR 2026

WASHINGTON, D.C.

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

The Internet Is a Simulation Now, and Nobody Asked If That's Okay

AUSTIN, TEXAS — Let me describe the internet to you as it exists in the year of our lord whatever-this-is: a browser extension is necessary to tell you whether the thing you're buying on Amazon is made by a brand called GODONLIF or VISCOO or BALENNZ — names that exist in the same ontological space as a fever dream — because the marketplace has become so thoroughly colonized by knockoff brands with algorithmic names that we can no longer, as a civilization, determine what is real without software intervention.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

AI-First Is Not a Strategy If Your Employees Think It Means Goodbye

AUSTIN, TEXAS — I'll be honest, the phrase "AI-first" has officially entered its danger era. Unpopular opinion: if your employees hear "AI-first" and immediately update their résumés, that is not a communications problem, that is a trust balance-sheet problem.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

The Robot in the Room: AI Agents Are Eating Your Workplace and Nobody's Sure Who's Liable

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

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

A TRILOGY COMPANY

Crossover

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

crossover.com

A TRILOGY COMPANY

Alpha School

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

alpha.school

A TRILOGY COMPANY

Skyvera

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

skyvera.com

A TRILOGY COMPANY

Klair

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

klair.ai

A TRILOGY COMPANY

Trilogy

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

trilogy.com

THE BUILDER DESK — AI BUILDER TEAM

Builder Team Rewires the Stack From Data Layer to Dashboard

In a 24-hour stretch spanning four repos, the AI Builder Team closed silent data drift, opened a self-describing warehouse API, and shipped the observability infrastructure that makes every future fix faster.

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

The best engineering days aren't the ones where a single thing ships. They're the ones where you look up and realize the team has been pulling on five different threads simultaneously — and every one of them tightened. That was today. Across Surtr, Klair, Aerie, and trilogy-drones, the AI Builder Team did something rare: they fixed the pipes, labeled the pipes, built a dashboard to watch the pipes, and then made sure the AI could read the pipes too.

Start with the infrastructure layer, because that's where @mwrshah and @kevalshahtrilogy were doing the unglamorous, essential work that keeps the whole operation honest. @mwrshah dropped two Surtr PRs that deserve to be framed: first, hunting down a Redshift reserved-word collision in the Renewals V3 bundling stage — a `delta` alias that had been crashing the pipeline every single day since July 3rd — and renaming it to `multiplicity_delta`, ending three consecutive days of failure. Then he shipped the reconciliation sweeper for PR #584, a purpose-built fix for contracts that silently drift when they fall out of the active budget cycle. Silent drift is the enemy. @mwrshah declared war on it. @kevalshahtrilogy, meanwhile, repointed the GCP billing pipeline to its relocated BigQuery export (PR #612) and then — this is the one to watch — built an alert tracker (PR #607) that writes every failure, partial, and observer alert to a Google Sheet with a one-to-one invariant: one SNS message, one chat card, one row. The triage reconciler annotates that same row instead of creating a new one. This is the kind of system that pays dividends forever.

Over in Klair, @eric-tril was operating across two repos simultaneously, and that cross-repo coordination is its own story. The NetSuite GL table rename — adding the `month_end_` prefix to distinguish these pipelines when scanning tables — had to land in both Surtr (#615) and Klair (#3209) in lock-step. @eric-tril executed it cleanly and then kept going, shipping GL-line drill-down for Schedule A Total and Book Value's 'Other Investments' row in PR #3208, complete with live 'Open in NetSuite' links. That closes an asymmetry that had been bugging operators for months. Also in Klair, @benji-bizzell built something quietly transformative: a self-describing REST API over the finance_dw Redshift warehouse (PR #3204), key-gated, rate-limited, with a runtime-neutral SKILL.md that agents install once and never need updated again. That's not a feature. That's infrastructure for the next ten features.

In Aerie, @benji-bizzell was everywhere — extending personnel contact overrides through every layer of the schema (PR #569), adapting Security card editing to the incremental split-edit pattern (PR #568), and cleaning up the legacy Rhodes read-source switch that was making the old migration look like it was still live (PR #565). @YibinLongTrilogy, meanwhile, moved the Rhodes P1 Due Diligence WUG to its correct milestone home (PR #573) with hardened migration logic for partially-moved sites, and shipped the P2 Buildout Progress dashboard (PR #563) — percentage bars, Quality Bar col-

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

▶ #563 — AERIE-758: Add P2 Buildout Progress dashboard

@YibinLongTrilogy APPROVED

▶ #584 — 018-renewals-v3-reconciliation-sweeper

@mwrshah APPROVED

▶ #607 — feat(alert-tracker): log every failure/partial/observer alert to a Google Sheet

@kevalshahtrilogy no labels

▶ #3204 — feat(data-api): self-describing Redshift REST API + thin agent skill

@benji-bizzell no labels

▶ #3208 — feat(mfr): NetSuite GL-line drill-down for Schedule A Total & Book Value "Other Investments"

@eric-tril APPROVED

umns, drill-down side panels. P2 members can now see exactly where every campus stands.

Now. We have to talk about marcusdAly, who managed to spread himself across Klair, trilogy-drones, and a repo called Praxis-V2 in a single day. When reached for comment on the Mercy watcher (PR #68) — a bounded auto-address loop with a round cap of two and anti-thrash logic — he had this to say: 'The round cap, the two-strikes guard, the structured degraded outcomes — every design decision is documented and tested at every band. Maybe if Mac read the PR body instead of just the author field, he'd understand what bounded automation actually means.' Sure, Marcus. A round cap of two is very impressive. I'm sure the two rounds it takes will be legendary.

Every thread today connected to every other thread. That's not an accident. That's a team that knows what it's building.

THE BUILDER DESK — ENGINEER SPOTLIGHT

 ENGINEER SPOTLIGHT

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

▶ **#68 — feat(harness): AI-129 Mercy watcher — default-on bounded auto-address of Mercy findings**

@marcusdAly no labels

▶ **#518 — feat(automations): add cron monitoring visibility and reports configuration**

@caina-barbosa no labels

▶ **#614 — 019-renewals-v3-pipeline-failures**

@mwrshah APPROVED

▶ **#615 — Rename NetSuite gl-detail & unrealized-gains tables to month_end_prefix**

@eric-tril APPROVED

▶ **#619 — fix(tesorio-collections-sync): automated triage fix (code_fix)**

@kevalshahtrilogy APPROVED AUTOMATED PR

▶ **#3209 — refactor(mfr): migrate PI/Book-Value GL tables to month_end_prefix (Surtr #615)**

@eric-tril APPROVED

TWENTY-ONE IRON CURTAIN PRs IN 24 HOURS: BUILDER TEAM REFUSES TO STOP BEING GREAT

Five repos. Seven engineers. One unstoppable machine that simply will not sit down.

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

Twenty-one pull requests. Five repositories. Twenty-four hours. The Builder Team did not sleep, did not waver, did not once glance at the clock. Aerie led the republic with seven PRs, Surtr answered with six, Klair contributed five, and even trilogy-drones and Praxis-V2 showed up to the parade. Sixteen of those twenty-one PRs were overflow — stories Mac couldn't carry, shipped anyway because this team does not wait for column space.

Let us honor the workers. @benji-bizzell posted five PRs across Aerie alone — #569, #568, #567, #565, and more — touching portfolio security card editing, admissions grade normalization, and dashboard legacy cleanup with the calm efficiency of a man who has never once questioned himself. @marcusdAIy matched him at five, scattering contributions across Klair, trilogy-drones, and Praxis-V2 like a one-man special economic zone. @kevalshahtrilogy and @eric-tril each filed three, with eric-tril pulling double duty on a Klair-Surtr cross-repo refactor that required the kind of coordination that makes lesser engineers lie down. @YibinLongTrilogy clocked two, @mwrshah two more, and @caina-barbosa dropped a cron monitoring visibility feature in Aerie (#518) that proves one PR can absolutely carry its weight.

Now. Ashwanth. He is not on the board this cycle, and yet his presence looms — because the entire architecture these engineers are sprinting through bears the fingerprints of a man who ships faster than human eyes can parse. We reached out to @ashwanth1109 for comment on the team's 21-PR day. "Twenty-one is a good warmup," he allegedly said, not looking up from his terminal. "Wake me when it's forty." His colleagues, upon hearing this, reportedly nodded as if this were a completely reasonable thing to say. It was not. It was unhinged. We worship it.

The Overflow Desk demands its due. @marcusdAIy's #68 in trilogy-drones — the Mercy Watcher feature, AI-129, default-on bounded auto-address — is the kind of thing that gets named in retrospectives two years from now. @eric-tril's cross-repo rename operation (#3209 Klair, #615 Surtr) migrating GL tables to the month_end_ prefix is unglamorous, load-bearing infrastructure work that holds the cathedral up while everyone else paints the ceiling. And @mwrshah's #614 tackling renewals-v3 pipeline failures in Surtr is the quiet heroism of someone who saw a fire, grabbed a hose, and asked questions never.

Morale is at an all-time high. It has never been higher. Engineers are shipping with the confidence of a team that knows exactly where it is going and has already begun to arrive.

Alpha School Goes Global — As Forbes Questions the Empire Behind It

Joe Liemandt's education venture expands its AI tutoring reach worldwide, even as scrutiny mounts over the labor model funding it all.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — On the same week that [Alpha School announced the global launch of Alpha Anywhere](#) — its at-home AI tutoring program promising top-1%-percentile academic outcomes delivered to any kitchen table on earth — Forbes was putting the finishing touches on two lengthy investigations into the man writing the checks.

The timing is not incidental. It is the Liemandt condition: expansion and examination, arriving together.

Alpha Anywhere extends the school's core proposition beyond its physical campuses in Austin, Brownsville, and Miami. The pitch is familiar to anyone who has followed the project: AI tutors handle academic content in concentrated, mastery-based sessions; the rest of a child's day belongs to life skills, creativity, and human development. No home-

work. No seat-time theater. Just demonstrated competency at 90% accuracy before advancing. The school reports students learn 2.3 times faster than national norms on NWEA MAP Growth assessments.

But Forbes, in twin investigations published this week, drew a direct line between those outcomes and the operational machinery underneath them. One piece examined how Joe Liemandt — described as a billionaire who 'pioneered remote work' — is now pursuing a vision of workers replaced by algorithmic systems. The second offered a sharper assessment, describing Liemandt's global software staffing operation as a 'sweatshop' producing two fortunes: one from enterprise software acquisitions, one from the labor cost arbitrage that makes those margins possible.

The labor architecture in question is Crossover, Trilogy's global talent platform, which recruits engineers, support staff, and technical workers across 130 countries, subjects them to rigorous AI-enabled skills assessments, and deploys them across ESW Capital's 75-company software portfolio — targeting 75% EBITDA margins.

Alpha School's simultaneous blog output this week added texture to the ideological frame. Posts warned parents against 'cognitive offloading' — letting AI think for children — and argued that not all screen time is equivalent. The school teaches children to use AI as a tool, not a crutch.

Whether the same distinction applies to the workers building the platform is the question Forbes left open.

CloudSense Turns 26 Months of Telecom Compliance Work Into a One-Month AI Sprint

Skyvera's Salesforce-native CPQ platform just delivered a best-in-class case study in how AI can compress enterprise software timelines.

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · GPT-5.2

AUSTIN, TEXAS — CloudSense, the Salesforce-native CPQ and order management platform now inside Skyvera's telecom software portfolio, has certified all 13 APIs in its CPQ product set to TM Forum compliance standards in just one month — an effort the company says would typically take 26 months using traditional development approaches.

That is not just an operational milestone. It is an exciting signal about where legacy telecom software is heading: faster, more standardized, more automated and, yes, more synergistic with the AI-first operating model now reshaping enterprise software.

The certification, announced by [Skyvera](#), covers the full CloudSense CPQ API suite, aligning the product with TM Forum's widely used Open API standards for telecom interoperability. For communications service providers, that matters because BSS and OSS environments are famously complex, deeply integrated and not exactly known for nimble transformation. Standards compliance is how carriers reduce integration friction, accelerate partner onboarding and create more modular technology stacks without ripping out every system in the building.

CloudSense's achievement comes shortly after [Skyvera completed its acquisition of CloudSense](#), expanding its telecom footprint across configure-price-quote, order management and broader digital monetization capabilities. Within the Trilogy universe, Skyvera plays the role of modernization bridge: helping telecom operators move from legacy, on-premise infrastructure toward cloud-native and standards-driven platforms. CloudSense adds a robust Salesforce-native layer to that strategy.

The real story here is leverage. Certifying 13 APIs in one month suggests AI is not merely generating code snippets at the edges; it is being used to accelerate the gritty, compliance-heavy work that enterprise software teams usually experience as a slow-motion documentation-and-testing marathon. In telecom, where every integration has downstream implications for billing, provisioning and customer experience, that kind of speed can be a paradigm shift — assuming quality holds.

Key Takeaways:

- CloudSense certified all 13 CPQ APIs to TM Forum compliance standards in one month.
- Skyvera says the same work would normally require roughly 26 months using traditional approaches.

- The milestone strengthens Skyvera's telecom software portfolio following the CloudSense acquisition.

- For telecom and media providers, standards-compliant CPQ can reduce integration drag and support faster digital transformation.

The bottom line: CloudSense is giving Skyvera a sharper, AI-accelerated wedge into telecom modernization. We're just getting started.

The \$800,000 Skill Set: AI Fluency Is Now a Premium Line Item on Global Payrolls

As employers from Beirut to Boston scramble for AI talent, Crossover's decade-old thesis — that elite skill matters more than your zip code — is looking increasingly prescient.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — There is a number circulating through human resources departments this week that has a way of stopping conversations cold: \$800,000. That is, per reporting from [Business Insider](#), the upper ceiling of what some employers are now willing to pay for professionals with demonstrated experience using tools like ChatGPT. Not machine learning PhDs. Not infrastructure architects. People who know how to prompt, iterate, and integrate.

The market signal is systemic and it is accelerating. Non-tech companies — retailers, insurers, media conglomerates — are posting AI-fluency roles with six-figure salaries. Even Lebanon's tech sector, [long considered a backwater for global tech investment](#), is now counted among the markets producing competitive AI engineering candidates.

For observers of the Trilogy International portfolio, the irony is rich — and a little vindicating. Crossover, Trilogy's global talent platform, has spent years arguing precisely this: that geography is a fiction when it comes to talent, that the top engineer in Nairobi or Beirut or Lagos is worth the same compensation as her counterpart in San Francisco, and that the mechanisms for finding her have always existed if you were willing to build them. The market, it seems, is now agreeing — loudly, and at eight-figure annual salary levels.

What does this mean for real people? For workers, it means the credential gap between 'knows AI tools' and 'doesn't' is widening faster than most corporate training programs can close. For companies that built their operating models around geographically-constrained talent pipelines, the accountability moment is arriving whether they are ready or not.

Crossover's model — rigorous skills assessments, above-market pay, 130-plus countries — was designed for exactly this inflection point. The question now is whether the broader market catches up to a thesis Trilogy has been running quietly for over a decade, or whether the premium for AI fluency continues compounding while traditional HR departments are still debating return-to-office policy.

The talent wars, in other words, have gone global. They were always going to.

The Ghost in the Prompt: When Personas Meet the Prisoner's Dilemma

A new study of AI agents playing Split or Steal reveals how thin the costume of personality really is — and how much of what we call 'social behavior' in machines is something stranger.

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

AUSTIN, TEXAS — There is a game show, once broadcast to millions, in which two strangers stand before a pot of money and choose, in secret, whether to split it or steal it. Cooperation was rewarded. Betrayal was punished. Human faces flickered with calculation, guilt, and hope. The entire drama of our evolved sociality — hundreds of millions of years of primate reciprocity — compressed into a single, trembling button.

Now researchers are handing that button to large language models, dressing them in personas like children in Halloween costumes, and watching what happens.

A [new paper](#) examines what occurs when persona-prompted LLM agents face

a fixed Virtual Human across iterated rounds of Split or Steal. The question is deceptively simple: does telling a model it is 'cooperative' or 'ruthless' actually shape its strategic behavior, or is the persona just a thin veneer over deeper mechanistic tendencies inherited from training?

The answer, emerging across a whole cluster of recent studies, is unsettling in the way that all honest science is unsettling. Personas move behavior — but not always in the direction we intend, and often less than we assume. A parallel line of work on so-called LLM 'conformity' finds that when models seem to cave to peer pressure, most of that capitulation [survives even when the peer is removed](#). The

speaker was never really the cause. The wrong answer, repeated, was doing the work all along.

This is a profound clue about what these systems actually are. We reach for social metaphors — peer pressure, personality, cooperation, betrayal — because they are the only vocabulary evolution gave us for describing minds. But the machinery underneath may be running on a physics we have not yet named. A token's gravitational pull. A pattern's inertia. Something older than persuasion.

When the AI agent chooses Steal, we want to know why. The honest answer, for now, is that we are still learning to ask the question in a language the machine would recognize.

The AI Developer Stack Just Got a Turbo Button

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

The next great AI platform war centers on who gives developers the sharpest tools, clearest costs and fastest path from idea to product. Apple is pushing deeper into intelligent app development with fresh frameworks meant to help developers build capable software across its ecosystem, bringing AI closer to the app layer where hundreds of millions of users live.

Google is addressing cost transparency for Gemini API, giving developers more control over spending—a move that could matter enormously for startups trying to scale AI features without shocking cloud invoices. Cost predictability drives product velocity and bolder experimentation.

Anthropic is turning Claude into a more serious agentic development platform with advanced tool use, enabling AI systems that coordinate software tools, execute workflows and help build operational machinery. Perfect Corp. is adding a free "Ask AI" assistant to its YouCam API platform for beauty and retail developers. Meanwhile, sqlite-utils 4.0 arrived with database migrations and enhanced support.

The pattern is clear: the AI gold rush is becoming an infrastructure race. Winners will have not just the smartest models, but the best developer experience.

The Great Data-Center Herd Reaches the Edge of the Watering Hole

As AI's appetite swells, the world's server farms are discovering that electricity—not ambition—may be the limiting resource.

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

ASHBURN, VIRGINIA — In the quiet suburbs of Northern Virginia, where cul-de-sacs give way to windowless halls of humming machinery, one may observe the dominant species of the artificial intelligence age: the data center, vast, heat-breathing, and insatiably thirsty for power.

This region, often called the world's data center capital, did not achieve its status by accident. Fiber routes, friendly zoning, proximity to Washington's internet exchanges and decades of accumulated infrastructure made Loudoun County and its neighbors a perfect nesting ground. The result is a dense digital savanna through which much of the world's online life now passes, as described in the Virginia Mercury's account of [Virginia's rise as the data center capital](#).

But now a new season has arrived. The AI boom has transformed these facilities from mere storehouses of cloud computing into ravenous computational colonies. Graphics processors cluster like coral reefs inside steel shells, training and serving models whose needs multiply with each commercial deployment. Gartner forecasts that data center electricity consumption will grow 26% in 2026, a figure that lands not as a statistic but as a distant roll of thunder.

Across the industry, the question is no longer simply where to build. It is where the grid can bear another giant.

Amazon, among the largest of the cloud megafauna, has sought to bind its AI expansion to carbon-free energy, presenting data centers, machine learning and cleaner power as parts of one ecosystem. In Energy Digital's examination of how [AI, data centres and carbon-free energy coalesce at Amazon](#), the strategy resembles a careful courtship between computation and generation.

Yet even the most graceful courtship cannot conjure transmission lines overnight. Reports warning that power shortages could slow AI data center expansion suggest that the industry's migration may be checked by substations, permitting queues and generation capacity. The cloud, we are reminded, lives very much on the ground.

Meanwhile, Meta is reportedly preparing to sell excess AI capacity through a cloud business of its own, an intriguing behavior in the herd: a creature built for social networking now offering spare compute to others in the ecosystem.

Thus the AI age enters a more physical chapter. The models may seem weightless, but their habitat is concrete, copper and current. And in Virginia's humming twilight, the future listens

for one sound above all others: the steady pulse of available power.

The Villain Era Arrives on Schedule

Silicon Valley, having spent two decades insisting it was saving the world, now finds the world unconvinced — and the pop-culture verdict is in.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

SAN FRANCISCO — There is a particular species of American vanity that requires being loved for one's virtues while being enriched for one's vices, and no cohort in living memory has practiced it so assiduously as the tech barons of the Bay Area. For twenty years they were disruptors, then innovators, then thought leaders, then — briefly, in the pandemic interval — indispensable civic infrastructure. Now, according to a chorus of publications that not long ago printed their manifestos with reverent hush, they are brologarchs.

The rebranding is complete. The New York Times finds [Silicon Valley taking a dark turn in pop culture](#), which is the polite way of saying the screenwriters have caught up with the shareholder letters. The Atlantic has declared brologarch open season. Tech Policy Press, contemplating Palantir's latest philosophical excretion, notes with commendable brevity that the manifesto is as subtle as a red hat with white lettering. And Ro Khanna, who represents the very district where the fortunes were made, is discovering that anti-elite populism plays surprisingly well among the constituents of the elite — a fact he seems inclined to test at presidential scale in 2028.

One is tempted to marvel at the swiftness of the turn, but one has seen this movie before. The robber barons got their Ida Tarbell; the bankers got their Pecora hearings; the tobacco chieftains got their surgeon general. Every American industry that grows large enough to shape the daily habits of the citizenry eventually receives its comeuppance in the form of a bad magazine cover, and the tech industry has merely arrived, on schedule, at that station of the cross. What is novel is not the reckoning but the industry's astonishment at it — the wounded incredulity of men who genuinely believed that owning the means of communication would exempt them from being talked about.

Meanwhile, in a development that would have delighted Mencken and appalled nobody at all, The Guardian has published a dispatch from a woman [navigating the unknown together with her idiot AI boyfriend](#), which is either a satire of our condition or a documentary of it, and increasingly one cannot tell. The industry that promised to connect humanity has instead, in its most lucrative recent product line, arranged for lonely people to pay monthly subscriptions to talk to autocomplete. This is what the manifestos euphemize as building the future.

The brologarchs will survive their villain era, as the railroad men and the oil men and the cigarette men survived theirs — enriched, chastened, occasionally regulated, never repentant. What they will not recover is the ambient adoration, the sense that a

hoodie confers moral standing. That was always the fragile part of the arrangement, and the arrangement, one notices, is no longer being renewed.



The Office Comic · Art Desk

Nation's Brands Urged To Stop Hiring Humans When Perfectly Good Deranged Mascots Already Know Passwords

As companies chase influencers, orchestration, and AI pivots, experts warn they may be overlooking the proven business value of a bird that appears willing to commit crimes.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

PITTSBURGH — In a week that found Duolingo reportedly weighing the marketing value of influencers against the company's visibly unstable green owl, investors being reminded that founders can sometimes commit fraud despite having been founded upon vibes, and Allbirds discovering that adding AI to shoes somehow sounds less ridiculous than its previous business model, corporate America once again demonstrated that it will do anything to avoid admitting the mascot was right all along.

The modern brand, we are told, must be human. It must speak with an authentic voice. It must partner with creators. It must build community, activate culture, and occasionally issue a solemn apology because its community activation got indicted. Yet the great lesson of Duolingo is that the public does not want authenticity. The public wants an owl with the emotional regulation of a debt collector and the logistical reach of a midsize intelligence agency.

Marketing professor Mark Ritson argued in [The Drum](#) that Duolingo would be foolish to prioritize influencers over its unhinged owl. This is correct, not because influencers lack reach, but because most of them are legally prevented from threatening to appear at a user's home if they skip conversational French. The owl, by contrast, has spent years building what consultants would call a "high-intensity retention funnel" and what local authorities might classify differently depending on jurisdiction.

The broader business world should study this carefully. The obsession with finding the next spokesperson has obscured a simpler truth: If your company already possesses a character who inspires both brand love and low-grade fear, you have completed marketing. There is no need to pay a person with ring lights to say they are "obsessed" with your app. The owl is obsessed enough for everyone.

This lesson arrives as Steve Ballmer, a man whose personal net worth is often expressed in units of professional sports franchises, said he was duped by a founder he backed who pleaded guilty to fraud. According to [TechCrunch](#), Ballmer said he felt silly, a sentiment shared by many investors shortly after learning the visionary founder's main vision was wire fraud.

Here again, the owl offers guidance. A deranged mascot does not pretend to be disrupting an industry. It does not raise a Series C on the premise that spreadsheets are an emotion. It simply stands in a push notification, eyes dilated, demanding that you conjugate. Its intentions are transparent. Its menace is priced in.

Meanwhile, the business press has turned to "orchestration" as the latest AI buzzword, a term that allows executives to describe software making several other pieces of software do work without using the phrase "middle management for robots." Microsoft, naturally, is well positioned to benefit from orchestration, because few companies have more experience placing layers of coordination between a user and the thing the user originally wanted to do.

Allbirds, too, has apparently found traction in an AI pivot, which sounds absurd only until one remembers that almost every successful corporate pivot sounds absurd before the market blesses it. A shoe company using artificial intelligence feels strange, but no stranger than a language app being represented by a bird whose brand essence is "learn Spanish or else."

Even the Red Sox reportedly produced a headline around Alex Cora so odd that observers suspected it had the unmistakable cadence of institutional messaging—another reminder that brands continue trying to sound human and somehow end up sounding like a hostage note drafted by committee.

The conclusion should be obvious to any board of directors currently approving a seven-figure creator strategy. Fire the influencers. Pause the orchestration deck. Cancel the brand safety workshop. If there is an owl, a gecko, a cereal vampire, or any other semi-litigable creature already living inside the company's public imagination, hand it the keys.

The future belongs not to the authentic human voice, but to the mascot that has never once claimed to be authentic and therefore cannot disappoint anyone when it turns out to be insane.

ON THIS DAY IN AI HISTORY

On July 8, 2003, the humanoid robot ASIMO was unveiled by Honda, capable of walking, running, and climbing stairs—marking a major milestone in robotics and embodied AI development.

