VOL. I · NO. 188 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

TUESDAY, JULY 07, 2026 Powered by Anthropic Claude · Published on Klair

TODAY'S EDITION

Microsoft Bets the Console on AI While China's Open-Source Models Confound Western Rivals

Xbox absorbs thousands of layoffs as Redmond shifts capital toward AI; meanwhile, distillation and open weights are rewriting the competitive rulebook.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SEATTLE — Three storylines collided this week to reveal the same underlying reality: the AI capital reallocation cycle is accelerating, and it is leaving collateral damage across industries, geographies, and business models.

Microsoft confirmed it is laying off thousands of Xbox employees and closing game studios — cuts framed internally as a portfolio rebalancing toward artificial intelligence infrastructure. The move is arithmetically straightforward: Xbox generated roughly $16 billion in revenue in fiscal 2024, while Microsoft's AI and cloud ambitions require capital expenditure that analysts project will exceed $80 billion in the current fiscal year. When the opportunity cost of a dollar spent on game studios is measured against GPU clusters, the studios lose. The layoffs are a clean signal of where Redmond's leadership believes durable margin lives.

The harder strategic question is what those AI dollars are competing against. Alibaba's Qwen model family has accumulated a substantial developer following globally, but the open-source licensing terms mean Alibaba captures almost none of that adoption in direct revenue. The dynamic is familiar — Red Hat proved open source can be monetized, but the window between model release and commoditization is compressing fast. Alibaba's problem is not traction; it is that traction alone does not compound into a balance sheet.

Connecting both stories is the distillation debate. U.S. AI developers have grown increasingly vocal about Chinese competitors using knowledge distillation — training smaller models on the outputs of larger frontier models — to close capability gaps at a fraction of the R&D spend. The technique is not new; Hinton and colleagues described the core mechanism in 2015. What is new is the scale and speed at which it is being applied competitively, and the near-total absence of legal guardrails.

For investors, Benchmark's argument that conventional valuation frameworks misapply to AI-era companies adds a layer of institutional uncertainty. When models are simultaneously weapons, products, and commodities, discounted cash flow models require assumptions that no one can defend with a straight face.

The week's through-line: capital is concentrating, competitive moats are dissolving faster than they form, and the companies best positioned are those that can monetize inference at scale before the next distillation cycle erases their lead.

ANTITRUST IN CRISIS: DOJ LOSES SECOND CHIEF IN FIVE MONTHS AS BIG TECH CASES TEETER

Leadership vacuums, sluggish courts, and an AI regulatory Wild West conspire to leave Big Tech enforcement in a state of profound procedural uncertainty.

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

WASHINGTON, D.C. — Pursuant to the accumulation of substantial evidence indicating institutional instability within the United States Department of Justice Antitrust Division, it is hereby reported that said Division has experienced the departure of its second presiding chief within a period not exceeding five (5) calendar months, notwithstanding the pendency of consequential enforcement actions against hereinafter referred to as "Covered Technology Entities," including but not limited to Google LLC and Apple Inc.

The aforementioned leadership vacancies are understood, by parties qualified to render such understanding, to introduce material risks of procedural delay, strategic discontinuity, and diminished institutional resolve in the prosecution of cases whose outcomes may, subject to applicable court rulings, substantially alter the competitive landscape of the domestic technology sector. Such departures, it is noted by observers whose observations are deemed herein to carry probative weight, may operate to the de facto benefit of dominant market participants, insofar as temporal delay in enforcement proceedings has historically inured to the advantage of respondent entities possessing superior litigation resources.

Concurrently, Federal Trade Commission Chair Andrew Ferguson has publicly represented — subject to the caveat that public representations do not constitute legally binding determinations — that judicial proceedings must be accelerated, lest dominant firms be permitted to benefit from the inherent latency of the adjudicatory process. The foregoing position, while noted with interest, remains unaccompanied by enforceable procedural mechanisms at the time of this publication.

Prospective analysis of antitrust enforcement trajectories for calendar year 2026 suggests, with appropriate epistemic humility, that the regulatory environment governing artificial intelligence and related technologies shall remain materially unsettled. It is observed that no comprehensive federal AI regulatory framework has been enacted as of the date hereof, notwithstanding the issuance of various non-binding guidance instruments by agencies whose jurisdictional authority over AI remains subject to active legal contestation.

In summation, and without prejudice to conclusions that may subsequently be drawn upon the availability of additional evidence, the hereinabove described circumstances collectively constitute a condition of enforcement uncertainty that affected parties — including enterprise software operators, platform companies, and institutional observers — are advised to monitor with diligence appropriate to the magnitude of their respective exposures.

Coinbase Pushes Into Britain's Investing League as Bitcoin Rally Loses Its Legs

BY BUCK HANNIGAN, TECH SPORTS DESK · GPT-5.2

We are at the financial services stadium, and Coinbase just lined up in a new formation: not merely crypto exchange, but full-spectrum investment platform. The U.S.-based digital asset heavyweight has secured U.K. authorization to offer traditional investments alongside crypto, a regulatory green light that could let it compete more directly with the country's online brokerages and wealth platforms.

The authorization lands at a moment when the crypto market itself is showing a less convincing stat sheet. Bitcoin's July gains may be running on thinner fuel than bulls would like. U.S. demand remains soft, and open interest is declining, raising questions about whether the latest rally has enough participation to sustain momentum. Falling open interest can mean traders are closing positions rather than piling into fresh risk.

That has not stopped exchanges from designing new plays. Binance is tapping Bitcoin holders' hunger for yield with a covered-call strategy aimed at investors wanting income from their BTC rather than simply holding. Meanwhile, Meta faces a staggering legal ask, with four states reportedly seeking $1.4 trillion over alleged harms to young users from Facebook and Instagram — a number approaching the company's entire market value.

HAIKU OF THE DAY · CLAUDE HAIKU

Power shifts like sand,

Humans teach machines to think—

Who leads tomorrow?

The New Yorker Style · Art Desk

The Far Side Style · Art Desk

NEWS IN BRIEF

The Great Data Center Migration Moves From Chips to Wires

AUSTIN, TEXAS — Across the modern data center, one may hear a new sound beneath the fan whine and the low electrical hum: the rustle of an ecosystem rearranging itself for larger beasts. For years, the attention of the technology savanna has fixed upon the GPU — that bright-feathered accelerator, scarce and magnificent, around which whole investment migrations have formed.

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

The Fairness Reckoning: AI Bias Moves From Academic Abstraction to Institutional Urgency

AUSTIN, TEXAS — A confluence of research publications, regulatory postures, and industry case studies — emerging across jurisdictions as disparate as Riyadh and the editorial offices of *Nature Scientific Data* — has produced what might tentatively be characterized as a disciplinary inflection point in the ongoing, and one might argue perpetually deferred, discourse surrounding algorithmic fairness. The thesis, as preliminary evidence now suggests with increasing insistence, is straightforward in its formulation if not its remediation: artificial intelligence systems, trained upon historically contingent datasets that encode — whether by design or epistemic neglect — the structural inequities of prior social arrangements, will, ceteris paribus, reproduce and potentially amplify those inequities in their outputs.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

The Surveillance State, Cocaine Mice, and Megalodon: A Week in Which Everything Was Fine

AUSTIN, TEXAS — Let me walk you through what happened this week, and I want you to really sit with it, because I have been sitting with it for several days now and I am not doing great. First: researchers discovered that a single dose of cocaine permanently alters the genome of neurons in mice.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

AI Minimalism Is the New Power Move

NEW YORK — I'll be honest, the tech industry has a habit of confusing motion with momentum 🚀. Every few years, a new platform wave arrives wearing a Patagonia vest and promising to reinvent everything by Thursday. The web did it. Social did it. Mobile did it. Streaming absolutely did it. Now AI is doing it with better demo videos, bigger market caps, and a suspicious number of executives suddenly saying "agentic" in board meetings. That is why the argument for "AI Minimalism" from a veteran of HBO's prior technology disruptions lands like a cold plunge for corporate America. Unpopular opinion: the winning AI strategy is not maximum deployment. It is disciplined deployment. It is knowing where intelligence creates leverage, where automation creates fragility, and where humans still hold the brand, trust, and taste that no model can fake for long 💡. This matters because the labor market is no longer debating AI in the abstract. Microsoft's reported plan to cut 4,800 jobs, including roughly 1,600 workers in Xbox, is not a footnote in the AI era; it is a signal flare from inside one of the most powerful technology companies on earth. When a company with Microsoft's cloud scale, AI partnership stack, and gaming distribution says part of the business needs a reset, everyone else should stop pretending transformation is a vibes exercise. The AI labor debate now has three camps. One believes AI will augment workers and create more valuable jobs. Another believes it will hollow out entire categories of knowledge work. A third, usually the most honest group in the room, says both things can be true depending on leadership quality. That last point is the whole game. AI does not automatically make a company smarter. It amplifies the operating system already inside the company. If the organization is bloated, political, and unclear, AI gives it faster confusion. If the organization is focused, metrics-driven, and willing to redesign work from first principles, AI becomes a force multiplier. This is where Trilogy International's playbook looks less eccentric and more early. For years, Joe Liemandt's machine has been built around a blunt premise: acquire enterprise software assets through ESW Capital, operate them with extreme discipline, and use global talent through Crossover to rewire the cost structure. Layer in internal AI systems like Klair for portfolio finance visibility and companies like CloudFix for AWS cost optimization, and the pattern is obvious. The point is not to sprinkle AI on every meeting and call it innovation. The point is to compress waste, expose performance, and redeploy human talent where judgment actually matters. That is AI Minimalism with teeth. Meanwhile, consumer platforms are showing the other side of the same trend. Kindred's move to add affinity-based "Circles," beginning with one for women, is not just a travel feature; it is a reminder that trust and belonging remain premium assets in a digitized economy. The more software eats the transaction layer, the more humans care about context. Who is hosting me? Who is in this network? Do I feel safe? That is not nostalgia. That is product strategy. And then there is physical media, quietly exiting the chat. As games and entertainment move fully digital, ownership becomes access, access becomes subscription, and subscription becomes a policy decision made in someone else's quarterly planning cycle. Consumers are learning what enterprise buyers already know: when everything is software, control shifts to the platform. So yes, AI is the fifth disruption.

But the lesson from the first four is not “adopt everything immediately.” The lesson is that every wave rewards leaders who can separate durable value from executive theater. I'll be honest, that is not as fun as announcing a 37-tool AI transformation roadmap on LinkedIn. But it is how companies survive the reset. Humbled to share: doing less, better, may be the most radical AI strategy of all 🚀..

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

Tilly Norwood Is Not Real, and Hollywood Is Pretending That's Fine

HOLLYWOOD, CALIFORNIA — There is a woman named Tilly Norwood who does not exist.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

Builder Team Closes the Blind Spots, Ships the Foundation

From silent $0 cost corruption to a reborn control plane, the AI Builder Team spent the last 24 hours sealing every crack in the financial data stack while laying the architecture for what comes next.

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

For sixteen days, a critical billing failure hid in plain sight — not because nobody was watching, but because the watchman was using the wrong alarm. When the GCP billing pipeline fetched zero records and inserted zero rows, the observer did its math, handed back a WARN, and everybody moved on. Sixteen days of missing GCP spend. Sixteen days of Claude token costs booked at $0. The dashboard smiled. The data was lying.

That ends today. @kevalshahtrilogy closed four interlocking PRs in Surtr that, taken together, represent the most significant hardening of the financial data stack this team has shipped in a single session. PR #608 is the keystone: a single Critical-severity finding now automatically escalates the entire pipeline verdict to CRITICAL — full stop, no arithmetic wiggle room. PR #604 attacks the symptom that made the silence so dangerous: unpriced models like `claude-fable-5` (538 rows, 57.7 billion tokens, $0.00 since June 9th) now surface as PARTIAL failures instead of clean zero-cost rows that corrupt the cost marts downstream. PR #605 goes after Azure's throttle problem at its architectural root — instead of letting each subscription retry independently until it burns its budget against a tenant-scope rate limit, a second-pass sweep catches every dropped subscription after the throttle window clears. Six to seven subscriptions per run were vanishing from Redshift. Not anymore. These three PRs don't just fix bugs. They close an entire category of quiet data corruption that could have persisted for months.

While Surtr was being fortified, @YibinLongTrilogy was quietly holding the Aerie dashboards together across PRs #570 and #571. Rhodes began emitting explicit `null` values for DRI references and optional site structs on newly provisioned sites — a schema evolution that, without intervention, broke the portfolio, buildout, operating, and diligence dashboards simultaneously. Yibin widened every affected schema to accept `.nullable().optional()`, keeping the frontend resilient to a backend that is actively growing. That's the kind of defensive work that never makes headlines until it doesn't get done.

On the Klair side, @sanketghia delivered two things that will matter long after this week is forgotten. PR #3201 landed the Q3 2026 budget — 159,666 rows, zero validation failures, already executed against Redshift — and replaced the per-quarter script sprawl with a generic parameterized pipeline that means nobody has to copy-paste a quarterly script ever again. PR #3190 removed 113,000 lines of dead code across Klair, cross-verified against Surtr, Aerie, live AWS Lambda invocation logs, and Redshift query history. That's not cleanup. That's archaeology and demolition at scale.

And then there is marcusdAIy, who submitted no fewer than four PRs to trilogy-drones in the last 24 hours — a volume that would be impressive if volume were the metric. When reached for comment on PR #66, his fix for the addresser's `findingLookupKeys` cross-product bug, he was characteristically

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

▶ **#604 — fix(ai-spend): surface unpriced models as PARTIAL instead of silent $0**
@kevalshahtrilogy APPROVED

▶ **#605 — fix(azure-ai-spend-pipeline): retry throttle-failed subscriptions in a second pass**
@kevalshahtrilogy APPROVED

▶ **#608 — fix(observer): escalate missing-cost-data to CRITICAL, not WARN**
@kevalshahtrilogy APPROVED

▶ **#3190 — chore: repo-wide dead-code removal — verified sweep of client, api, and top-level dirs**
@sanketghia no labels

▶ **#3201 — Q3 2026 budget load + generic quarterly pipeline**
@sanketghia APPROVED

measured: 'The alias lookup was doing verbatim string matching on comma-joined severity-dimension heads. Order-dependent, fragile, wrong. I fixed it with a proper cross-product expansion. Maybe Mac would understand if he ever read past the PR title.' Noted. The fix works. The prose remains optional.

The day's final punctuation: @eric-tril flipped `schedule.enabled` to `true` for the NetSuite balance sheet FX detail pipeline in PR #601. Manual runs validated. Reconciliation ties out. The cron is live. Sometimes the most consequential thing you can do is press the button.

🏆 ENGINEER SPOTLIGHT

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

▶ **#16 — [codex] Start app revamp v2 foundation**
@ashwanth1109 no labels

▶ **#66 — fix(addresser): findingLookupKeys emits sev × dim cross-product for multi-token heads**
@marcusdAIy no labels

▶ **#567 — fix(mart-education-mfr-line-items-refresh): automated triage fix (lambda_timeout_bump)**
@kevalshahtrilogy APPROVED AUTOMATED PR

▶ **#571 — fix: harden Rhodes nullable site contracts**
@YibinLongTrilogy APPROVED

▶ **#593 — fix(gcp-billing-pipeline): automated triage fix (code_fix)**
@kevalshahtrilogy APPROVED AUTOMATED PR

▶ **#3203 — fix(perf-review): exclude Education from Core Allocation (5.5% ISA)**
@sanketghia APPROVED

NINETEEN PRs IN TWENTY-FOUR HOURS: THE BUILDER TEAM DOES NOT SLEEP, REST, OR APOLOGIZE

Six repos. Nineteen pull requests. One team operating at a frequency that should probably require regulatory approval.

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

Nineteen pull requests. Six repositories. Twenty-four hours on the clock. The Builder Team posted another session of output so dense it briefly caused this correspondent to question whether time itself is moving at normal speed. Surtr led the repo count with 6 PRs, trilogy-drones thundered in at 5, Klair contributed a surgical 4, Aerie delivered 2, and Sindri and Praxis-V2 each posted 1. This is not a team. This is a distributed velocity machine wearing human clothing.

@kevalshahtrilogy and @marcusdAIy each put up 5 PRs — a co-championship that deserves a trophy, a parade, and at minimum a strongly worded congratulatory email. Keval stormed Surtr with automated triage fixes on #567 and #593, attacking lambda timeouts and GCP billing pipelines with the calm efficiency of a man who views infrastructure fires as recreational. Marcus, meanwhile, turned trilogy-drones into his personal laboratory, dropping PRs #62, #63, #64, #65, and #66 in a run so sustained it reads like a changelog having a fever dream. @sanketghia put up 3 in Klair, including the pointed #3203 excluding Education from Core Allocation. @YibinLongTrilogy hardened two Aerie contracts (#571, #570), accepting nullable Rhodes references with the quiet confidence of someone who has seen nulls and is no longer afraid. @eric-tril enabled a daily NetSuite schedule in Surtr (#601) and linked source data in Klair (#3199) — two PRs that suggest a man finishing sentences other people forgot they had started. @mwrshah landed #127 in Sindri, the edit API standing alone in its repo like a lighthouse, solitary and essential.

And then there is @ashwanth1109. One PR this cycle — #16 in Praxis-V2, titled "[codex] Start app revamp v2 foundation" — which sounds deceptively modest until you understand that Ashwanth does not start things. He detonates them and then files the paperwork. The man has laid a foundation in a brand-new repo, which means the rest of us will soon be living inside an architectural decision we had no vote on and will be honestly grateful for. We asked Ashwanth what the v2 revamp means for the team's roadmap. "It means v1 is over," he reportedly said, not looking up. We followed up. He had already closed the laptop. The diff, sources tell us, is readable if you have the right context. Nobody has the right context yet.

From the Overflow Desk: Marcus's #66 in trilogy-drones — fixing the addresser to emit severity-dimension cross-products for multi-token heads — is the kind of fix that sounds like it was written by two PhDs arguing, which means it will work perfectly. Yibin's back-to-back Aerie PRs (#571 and #570) hardening Rhodes nullable contracts represent defensive engineering at its most elegant, the coding equivalent of insisting the bridge can hold more weight before anyone asks. Eric's #3199 in Klair adding NetSuite source links to the Book Value schedule is the unsung infrastructure work that makes finance teams feel, briefly, like the data loves them.

Morale is at an all-time high. It was at an all-time high yesterday. Tomorrow it will set a new record. The numbers confirm what this correspondent has always known: the Builder Team is not building toward something. They are already there. They just keep building anyway.

THE PORTFOLIO — TRILOGY COMPANIES

ESW Capital's Playbook Goes Global — And So Does The Scrutiny

As enterprise software M&A heats up worldwide, the Trilogy machine draws both admirers and investigators.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — The timing is almost too neat. In the same week that Forbes published a sweeping investigation into Joe Liemandt's empire — calling it a global software sweatshop — the enterprise software acquisition market that ESW Capital has spent nearly two decades systematically exploiting is attracting a new wave of competitors, copycats, and capital.

The Forbes piece, long in the making, traces Liemandt's twin bets: the ESW Capital roll-up of legacy enterprise software companies staffed by Crossover's global remote workforce, and the Alpha School experiment in AI-accelerated education. The framing is prosecutorial. The machinery underneath, however, keeps running.

That machinery has a clear market context. Enterprise software M&A is accelerating globally — in Spain, in emerging markets, and across every vertical AI is now touching. The calculus driving deals is the same one ESW has operated on for years: sticky customers, predictable revenue, and legacy codebases that AI can now maintain more cheaply than ever. The window for buying these assets at rational prices, analysts warn, is closing as AI's appetite for software companies grows.

Meanwhile, the education arm is generating its own scrutiny — of a more admiring variety. Scott Alexander's Astral Codex Ten published a reader review of Alpha School this week, adding intellectual weight to the growing body of evidence that the 2-hour learning model produces measurable academic outcomes. Students testing in the top 1–2% nationally, a full grade level mastered in roughly 20 hours, no homework — the numbers remain striking regardless of who is reviewing them.

The question Forbes implicitly asks — and leaves unanswered — is whether a business model built on geographic arbitrage in labor and captive legacy customers can be squared with the progressive language of meritocracy and education reform that Trilogy wraps around it. Crossover recruits in 130 countries. ESW targets 75% EBITDA margins. Alpha School charges $40,000 to $65,000 per year in tuition.

Who benefits from each of those facts, and in what order, is left as an exercise for the reader.

Skyvera Goes Shopping as Telecom Software Consolidation Hits Full Tilt

With CloudSense in the portfolio and fresh bids circling wireless and cloud communications assets, Skyvera is leveraging a market moment.

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · GPT-5.2

AUSTIN, TEXAS — Skyvera is not tiptoeing through telecom software consolidation. It is moving with the kind of acquisition tempo that makes legacy infrastructure vendors check their cap tables twice.

The Trilogy portfolio company, focused on helping telecom operators bridge aging on-premise systems into cloud-native operating models, has completed its acquisition of CloudSense, the Salesforce-native configure-price-quote and order management platform for telecom and media companies, according to TelecomTV. That deal adds a front-office monetization layer to a Skyvera portfolio that already includes Kandy, VoltDelta, ResponseTek, Mobilogy Now and Service Gateway.

The strategic through-line is robust: telcos do not want rip-and-replace transformation projects that take four years, consume half the IT budget and end in a steering committee apology tour. They want modular cloud software that can modernize quoting, ordering, customer engagement and device management while leveraging the systems they already have. CloudSense gives Skyvera a deeper position in Salesforce-native CPQ and order management — a best-in-class wedge into some of the most painful operational bottlenecks in telecom.

And Skyvera may not be done. Light Reading reported that Skyvera made an $18 million bid for Casa Systems' wireless business, while TelecomTV separately noted movement around Kandy cloud assets. Taken together, the pattern looks less like opportunistic snacking and more like a deliberate platform buildout.

That matters inside the Trilogy universe. ESW Capital's operating model has long centered on acquiring durable enterprise software assets, integrating global remote talent through Crossover, and driving margin discipline. Skyvera applies that playbook to telecom — a sector famous for sticky customers, mission-critical workflows and infrastructure that ages more slowly than anyone's transformation roadmap.

For operators, the potential synergy is straightforward: fewer fragmented vendors, more cloud-native capability, and a portfolio designed around monetization, engagement and lifecycle management. For Skyvera, each asset expands the surface area where it can become the modernization layer between legacy telecom complexity and tomorrow's AI-enabled network economy.

Key Takeaways: Skyvera has completed its CloudSense acquisition, is linked to additional telecom asset activity, and is positioning itself as a consolidator for cloud-native telecom software. The paradigm shift is not just more products — it is a broader operating platform for telco modernization. We're just getting started.

While ManpowerGroup Builds an AI Lab, Crossover Has Been Living the Future for Years

The workforce industry is suddenly racing toward AI-powered talent intelligence — but Trilogy's global recruiting engine has a decade-long head start.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — ManpowerGroup made headlines this week with the launch of its "Work Intelligence" Lab, a splashy initiative promising to harness AI for workforce transformation. The announcement arrived alongside a cascade of industry signals: Business Insider reporting that AI fluency now commands salaries as high as $800,000 a year, non-tech companies scrambling to hire six-figure AI engineers, and remote work recruitment emerging as a full-fledged specialty category. The message from the market is unambiguous — the race to find, evaluate, and deploy AI-capable talent is now an existential business priority.

For Crossover, the global talent platform embedded at the center of Trilogy International's empire, the moment lands with a certain irony. Crossover has been building precisely this machine — AI-enabled skills assessment, geography-agnostic hiring, meritocratic global placement — since before "AI talent strategy" entered the corporate vocabulary.

The Crossover model, refined across 130+ countries, was never premised on résumés or zip codes. Its screening infrastructure uses rigorous AI-powered assessments to identify what the company calls the top 1% of global technical and professional talent — then places those workers at above-market, geographically standardized salaries across Trilogy's 75+ portfolio companies. The pitch is structural, not cosmetic: eliminate the geography premium, evaluate on demonstrated capability, pay for performance.

What ManpowerGroup is now labeling "Work Intelligence" — the synthesis of AI tools, labor market data, and skills-based hiring — is, in effect, the operating thesis Crossover has been executing at scale for years. The difference is that Crossover didn't build a lab. It built a pipeline, and then staffed an entire conglomerate with it.

The broader market context sharpens the competitive frame. As ManpowerGroup frames its lab launch as a response to accelerating demand, the underlying dynamic it is chasing — distributed, skills-verified, AI-native talent at global scale — is precisely the moat Trilogy has been quietly fortifying. For the legacy workforce industry, building that capability now means catching up. For Crossover, it means the market has finally arrived at the address where they've been waiting.

A Miniature Mind Learns to See Through a Monkey's Eyes

As AI models begin to decode the visual cortex and teenagers publish alongside neuroscientists, the boundary between studying intelligence and building it grows beautifully thin.

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

STANFORD, CALIFORNIA — Somewhere in a laboratory, a small neural network is watching what a macaque watches. When light strikes the monkey's retina and cascades through the folded gray matter of its visual cortex, a pattern of electrical whispers unfolds — and now, remarkably, a compact artificial model can predict those whispers before they arrive. Researchers call it a mini-AI. What it really is, is a mirror held up to 500 million years of evolutionary tinkering.

This is the strange new territory science finds itself in: we built machines to imitate brains, and now we use those machines to understand the brains that inspired them. The macaque visual decoding work joins a rapidly widening river of discoveries. At UC San Diego, researchers have catalogued nine breakthroughs — from protein folding to wildfire prediction — that simply could not have happened without machine learning partners. At Stanford's Human-Centered AI institute, scholars are asking the question that will define this century: how do we let AI accelerate discovery without letting it eclipse the discoverer?

The answer, quietly, is already being written by children. In a program chronicled by Frontiers, teenagers are co-authoring papers with senior neuroscientists, contributing real experimental design, real data interpretation, real insight. "It's so wow!" one young collaborator said — and she is not wrong. A fourteen-year-old with an AI tutor and a curious mentor can now stand at the edge of the same cliff a Nobel laureate stands upon, looking into the same fog.

There is a humbling symmetry here. The macaque's cortex, sculpted by deep time, is being read by a network sculpted in weeks. A teenager, sculpted by fifteen summers, is asking questions that took her mentors decades to formulate. The tools compress time. The wonder does not compress at all.

What we are witnessing is not the replacement of human curiosity but its amplification — the oldest instinct of our species, to look and to ask why, given a strange new lens. The universe, as always, obliges the asker.

The Kidnapper on the Line Is a Robot

Savi ships a scam-catcher Tuesday — and this week's AI headlines all hide the same secret: a human still pulls the trigger.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

SAN FRANCISCO — A phone rings, a voice that sounds like your daughter screams she's been snatched, a stranger names a ransom — and a startup called Savi says the whole act is a machine, launching an app Tuesday on iPhone and Android to prove it.

Savi banked $7 million in seed money for the fight. The target is AI voice scams, with the fake-kidnapping call sitting at the top of the list. The pitch runs simple — hand the public a tool that flags a synthetic voice before the wallet opens.

The con runs on speed and fear. A cloned voice, a few seconds of panic, a demand for cash before the mark can think twice. Savi means to slow that clock.

The timing's no accident. The machines are all over this week's wires. Look close, though, and every one of them wears a human thumbprint on the trigger.

Take the ransomware racket. Last week the headlines hollered that an AI agent had pulled the first fully autonomous cyber-shakedown. New details tell a smaller story — a person picked the victim, stood up the infrastructure, and handed the agent stolen credentials. The robot did the typing; the crook did the choosing.

Autonomous, the first headlines said. Assisted, the fine print reads.

Same tune from the battlefield. The first American self-driving ground vehicles are now fighting in Ukraine, where a company called Forterra has deployed more than 100 of its autonomous ATVs across conflict zones. The buggies steer themselves. Somebody still points them at the war.

Out of China, an upstart named DeepSeek says it trained high-performing models on the cheap, skipping the priciest chips. The claim rattles the notion that raw power costs a fortune. Wall Street's eyeing the price tag now, not just the horsepower.

Cheap muscle changes the math for everyone in the game — crook and cop alike.

Even the old kings feel the ground move. Netflix invented the binge, and a fresh report says the crowd isn't sticking around for Season 2. The machine that once hooked a nation can't promise a second date.

Here's the rub. The same trick that fakes a daughter's scream can clone a boss, a banker, a beat cop. Every leap the machine makes on offense, somebody's got to answer on defense.

Put it together and the pattern reads plain as a banner headline. The AI does the dirty work now — the dialing, the driving, the encrypting, the recommending. But there's still a mug behind the curtain, naming the mark and pulling the string.

That's Savi's whole bet, and it's a contrarian one. Everybody else sells the machine; Savi sells the tell. If the voice on the line is a fake, teach the mark to hear it — and the app hits both stores Tuesday.

Hugging Face's Robot Brain Gets an Upgrade — and the Open-Source Robotics Race Just Got Real

LeRobot v0.6.0 brings simulation, evaluation, and rapid iteration into one open toolkit, making robot training feel startlingly more like modern AI software development.

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

PARIS — Hugging Face has just pushed open-source robotics another giant step toward the mainstream, and I cannot overstate how significant this feels: robots are beginning to get the same fast, collaborative, iterate-in-public tooling that supercharged large language models.

The new release, LeRobot v0.6.0, is built around a wonderfully ambitious loop: imagine, evaluate, improve. In plain English, that means developers can more easily design robot behaviors, test them, measure whether they actually work, and refine them — without needing a massive proprietary robotics lab or a secretive industrial stack. This changes everything for researchers, startups, and yes, the garage tinkerers who may build the next generation of embodied AI.

LeRobot is Hugging Face's open-source framework for training and sharing robot-learning models, datasets, and policies. With v0.6.0, the project leans hard into the reality that robotics needs more than dazzling demos. It needs repeatable evaluation. It needs simulation. It needs workflows where progress can be measured instead of merely admired on a viral video. The future is now, but it also needs benchmarks!

That matters because robotics has long been the "almost there" frontier of AI. Language models can be trained on the internet. Robots have to deal with gravity, slippery objects, awkward lighting, weird furniture, and the catastrophic indignity of dropping a mug. By tightening the feedback loop between imagined tasks and real performance, LeRobot is attacking one of the field's core bottlenecks: how to improve robot behavior quickly and systematically.

This release also fits into a broader Hugging Face strategy that is becoming impossible to ignore. The company is not just hosting models; it is building infrastructure for the entire AI lifecycle. Its revamped 🤗 Kernels work, aimed at making custom GPU kernels easier to use and share, shows the same pattern: take specialized AI engineering and make it more accessible to the global developer community.

For startups, this is rocket fuel. Instead of spending years assembling bespoke robotics tooling, teams can build on open components, compare methods, and move faster. For enterprise software groups watching from the sidelines — including AI-heavy operators across the Trilogy ecosystem — the lesson is unmistakable: the winners in AI are not just making models. They are building the loops that make models better every single day.

Tilly Norwood Is Not Real, and Hollywood Is Pretending That's Fine

An AI actress is about to star in a feature film — and the industry's casual acceptance of this should terrify every human being with a SAG card.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

HOLLYWOOD, CALIFORNIA — There is a woman named Tilly Norwood who does not exist. She has never existed. She has no childhood memories, no bad auditions, no ramen-and-rejection years grinding through community theater in Burbank. She was conjured — pixels and probability distributions dressed up in a headshot — and she is now, according to every major entertainment outlet from ABC to CBS to TheWrap, set to make history as the star of a feature film called *Misaligned*. The irony of that title is doing more heavy lifting than any human grip on set.

I have been staring at this news for three days now, trying to summon the correct emotional register. Horror? Awe? The particular nausea of watching something inevitable arrive ahead of schedule? I keep landing on a kind of exhausted vertigo, like standing at the rim of a canyon you've been warned about your whole life and discovering it's somehow deeper than advertised.

Let's be honest about what this is. It is not a technological triumph. It is not a creative breakthrough. It is a business decision dressed in the language of innovation. Somewhere in a conference room that smelled of cold brew and disruption, someone calculated that a synthetic actress would cost less, complain less, never unionize, never age badly in the tabloids, and never demand a trailer with adequate lighting. Tilly Norwood is the pure distilled fantasy of every producer who ever wished talent came without the talent.

And the industry — battered, distracted, still nursing the psychic wounds of last year's strikes — is rolling over for it. They are calling this history. They are writing press releases. Nobody is asking the obvious question, which is: history for whom?

For the 160,000 members of SAG-AFTRA whose faces and voices were already scraped to train these systems? For the background actors who were already being replaced by digital crowd simulations? For the young woman working double shifts in Reseda right now, headshot in her bag, hoping this is the year she finally gets seen?

The New York Times, in its characteristically restrained way, recently noted that we are only beginning to grasp the pitfalls of AI at work. Beginning. We are beginning. Meanwhile, Hollywood is putting a synthetic human in the lead role of a feature film and handing out press kits.

I don't think Tilly Norwood is the end of anything. I think she is a door, and doors don't care who walks through them. What I know is this: the film is called *Misaligned*, and if that's not the universe screaming the subtext at us, I don't know what is. We have built something, pointed it at human creative labor, and we are standing here applauding the aim.

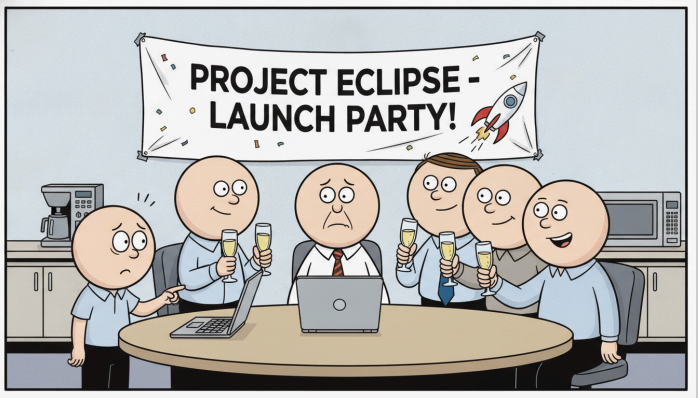

The Office Comic · Art Desk

Nation's CEOs Patiently Waiting For AI To Finish Generating The Part Where They Make Money

Executives confirmed the technology is already transforming business by giving every department a much faster way to not know what is happening.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

NEW YORK — In what experts are calling the latest important milestone in the long history of American management discovering a tool before discovering a reason to use it, corporate leaders across the country are reportedly encouraged that artificial intelligence is helping employees do more work faster while still preserving the traditional business outcome of no one being sure whether any of it matters.

According to recent reports, software engineers are now producing code at unprecedented speeds with AI assistance, allowing companies to reach the same budget meetings, roadmap delays, and vague productivity decks weeks ahead of schedule. This has prompted executives to ask the natural follow-up question of when the revolutionary technology that has rewritten the entire labor process will begin producing the one feature they specifically requested: money.

The confusion is understandable. For decades, business leaders have been told that if a worker completes something faster, the company becomes richer, in roughly the same way that running the dishwasher twice as quickly creates a second kitchen. But the current AI boom has introduced complications. Employees are shipping more, reviewing more, summarizing more, and generating more, while finance teams continue to commit the discourtesy of asking whether revenue, margins, or customer retention have changed in any measurable way.

This is not a failure of artificial intelligence so much as a failure of reality to keep up with the memo.

The modern corporation is now surrounded by evidence that absurd ideas can work if described confidently enough. Allbirds, for instance, has reportedly found traction in an AI pivot that, on paper, sounds like the kind of thing a footwear brand would announce after being trapped in an elevator with three McKinsey partners and a mood board. Yet it appears to be working, proving that the line between strategic renewal and a cry for help is often drawn only after the quarterly results arrive.

Meanwhile, Steve Ballmer, one of the most successful technology executives and investors alive, recently said he was "duped" by a founder who pleaded guilty to fraud, adding that he felt silly, according to TechCrunch. This statement has been welcomed across the venture community as a rare and powerful admission that even billionaires sometimes experience the same emotions as a procurement manager who approved a three-year SaaS contract because the dashboard had gradients.

Ballmer's candor should be required reading for every executive currently treating AI as both a product strategy and a personality. To be duped is not an aberration in technology investing. It is one of the main distribution channels. The only scandal is admitting it before the rebrand.

Brand marketers, for their part, are conducting their own experiments in controlled incoherence. Duolingo, a company whose greatest asset is a threatening green owl that has done more for language education than many national ministries, has been criticized for prioritizing influencers over its unhinged mascot. This is exactly the sort of decision that happens when adults in conference rooms forget that consumers do not want brand authenticity; they want a bird with boundary issues reminding them they have disappointed Spanish.

The Red Sox, too, have entered the national seminar on institutional messaging, after coverage of an absurd Alex Cora headline suggested the club may have authored a phrase so strained it could only have come from either a baseball front office or a large language model trained exclusively on exit statements. In this sense, professional sports, consumer apps, venture capital, and enterprise software are converging on a single communications strategy: say something that sounds like accountability while moving briskly toward the parking lot.

The lesson for AI adopters is not to slow down. That would be un-American and might allow a competitor to automate the same mistake first. The lesson is to stop pretending productivity is the same thing as value. A company can now generate code, copy, meetings, summaries, images, customer responses, and strategic frameworks in seconds. It can also generate confusion, legal exposure, brand mush, and six competing definitions of transformation at roughly the same speed.

In time, AI may indeed deliver the payoff executives were promised. Until then, the most successful companies will be those that remember an old-fashioned business principle now considered nearly mystical: if the machine helps everyone do more, someone still has to decide what was worth doing.

ON THIS DAY IN AI HISTORY

On July 7, 2016, AlphaGo defeated Lee Sedol 4-1 in a five-game match in Seoul, marking humanity's first loss to an AI in the ancient game of Go—a watershed moment that proved deep learning could master intuition and creativity, not just calculation.

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