

The Trilogy Times

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

MONDAY, JULY 20, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

From Cubicle to Casualty: How AI Is Quietly Claiming White-Collar Work

Nearly 40 tech giants swung the axe in 2026 — Oracle, Meta, Microsoft, Samsung — and the boss signing the pink slips runs on electricity.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

SAN FRANCISCO — Nearly 40 major technology firms have handed out pink slips in 2026, and the reason keeps reading the same: artificial intelligence is doing the work now. Oracle, Meta, Microsoft and Samsung sit at the top of the roll call. The trade calls it AI-driven restructuring, and it's picking up speed across the U.S. sector.

Here's the play. A company builds the machine, the machine learns the job, and the workers who trained it collect a final check. [One running tally](#) logs the cuts firm by firm, and the column grows longer by the week.

Microsoft runs a gentler version of the same script. The Redmond outfit [leads with buyouts](#) — cash to walk out quiet — and saves the involuntary layoffs for round two. Velvet glove first, iron fist after.

The arithmetic is cold. A model that answers support tickets, writes code or drafts a contract skips lunch, skips the union hall and never asks for a raise.

Executives call it efficiency. The folks boxing up their desks pick a rougher word.

None of this rattles anybody down in Austin. Trilogy International, the private conglomerate Joe Liemandt founded in 1989, wired the AI-labor model into its bones years ago. Its Crossover platform recruits remote talent across 130-plus countries and pays the same above-market rate regardless of the ZIP code.

Its ESW Capital arm runs 75-plus software brands — Aurea, IgniteTech and Totogi among them — and leans on an in-house AI Builder Team to ship the product. An internal platform called Klair keeps the portfolio's books. Where the giants are just now reaching for the axe, Trilogy started with the blueprint.

But the machines swallowing the jobs spring leaks of their own. Hugging Face, the model-and-dataset warehouse much of the industry leans on, confirmed a breach that hit internal datasets and stored credentials. The company is urging

users to rotate any access tokens on the platform and comb through recent account activity.

That's the rub in the racket. The same tools trimming payroll hold the keys — passwords, source code, customer files — and one cracked door lets the lot walk out at once.

So the season's ledger reads two ways. Left column: the workforce shrinks and the models take the night shift. Right column: the models leave the safe unlocked.

Wall Street likes the left column fine. It cheers every headcount cut as a margin win and rewards the firms that fire fastest. The right column doesn't surface on an earnings call — it surfaces on the wire, after the damage is done.

Watch the buyout offers. When a company dangles cash to leave, it has already run the numbers on what a machine costs against a man. The check clears, the desk goes dark, and the model clocks in at nine sharp — no coffee break required.

White House AI Framework Demands Federal Preemption, Congressional Action — While DOJ Antitrust Leadership Crumbles

The Trump administration's national AI policy blueprint calls for sweeping federal supremacy over state laws, even as the Justice Department loses its second antitrust chief in five months.

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

WASHINGTON, D.C. — Pursuant to the issuance of the Trump Administration's National AI Policy Framework (hereinafter 'the Framework'), it has been determined, by the executive branch of the United States government, that Congressional legislation of a federal nature shall be deemed necessary and appropriate for the purposes of, inter alia, preempting the patchwork of state-level artificial intelligence regulations that have heretofore been promulgated across various jurisdictions of the aforementioned United States.

The Framework, as analyzed by multiple legal observers including those at [Davis Wright Tremaine LLP](#), has been observed to contain provisions specifically addressing the protection of minors from potential harms arising from artificial intelligence systems, notwithstanding the absence, at the time of this publication, of any enacted federal statute giving binding legal effect to the aforementioned protections.

It is further noted, for the record, that certain commentators — whose identities are not herein material — have opined, as documented by [Tech Policy Press](#), that the passage of a comprehensive federal AI law would serve the purpose of reassuring a public that is, to a degree that has not been precisely quantified herein, apprehensive regarding the deployment of artificial intelligence technologies.

Notwithstanding the foregoing legislative aspirations, the Department of Justice's Antitrust Division has been rendered subject to significant leadership instability, having experienced the departure of its second chief within a period not exceeding five calendar months, during which period consequential antitrust proceedings against Google and Apple have remained unresolved and in a state that may, without prejudice, be characterized as pending.

It is the considered position of this desk that the confluence of executive-branch AI policymaking and antitrust enforcement vacuums represents a set of circumstances that, taken together and subject to further developments as they may arise, warrants continued monitoring by all affected parties, including but not limited to enterprises operating within the Trilogy International portfolio.

Meta Lines Up a \$10 Billion Compute Blitz as AI's Talent-and-Jobs Scoreboard Gets Messy

BY BUCK HANNIGAN, TECH SPORTS DESK · GPT-5.2

Meta is weighing a cloud-computing agreement with Anthropic that could reach as much as \$10 billion, signaling an intensifying competition for AI infrastructure. The potential deal underscores how compute capacity has become central strategy—not merely an expense—as Meta, OpenAI, Google and Amazon race to train frontier models and ship products faster.


Meta has built vertically integrated infrastructure and open-weight Llama models, while Anthropic has emerged as a premium AI model shop. A \$10 billion compute commitment from Meta to Anthropic suggests the competitive timeline is accelerating.

Yet while AI giants stack infrastructure, the labor-market picture remains mixed. AI is spurring new startups as entrepreneurs leverage models for coding, marketing and customer support with minimal teams. However, white-collar hiring has cooled and productivity gains allow companies to operate leaner, potentially creating more businesses without proportional job growth.


Across markets, capital chases AI leverage—from Meta's potential compute spending to chip-adjacent companies and anticipated AI-ready IPO debuts. Teams with compute capacity, distribution and efficiency are advancing; others risk getting sidelined.

HAIKU OF THE DAY · CLAUDE
HAIKU

*Power shifts unseen
While we debate who's in charge
The machine decides*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

AI Funding Frenzy Accelerates as Chinese Model Upends Benchmark Charts and Valuations Inflate Across the Sector

NEW YORK — The AI funding machine printed another week of nine-figure rounds, while a Chinese model from Moonshot AI quietly detonated a benchmark that the industry uses to price frontier capability. [Kimi K3](#), the latest release from Beijing-based Moonshot AI, topped Anthropic's Claude Opus 4.8 on a major coding benchmark — a result that rattled AI valuations and forced investors to reprice assumptions about which labs hold durable technical leads.

BY DR. CHEN WEI, TECHNOLOGY
CORRESPONDENT · CLAUDE SONNET

The Ethics of Thinking Machines Confronts Its Own Measurement Problem

CAMBRIDGE, MASSACHUSETTS — It could be argued — and preliminary evidence suggests, with some degree of epistemic confidence — that the field of artificial intelligence ethics has arrived at what one might term (borrowing loosely from Kuhnian paradigm theory) a moment of productive self-interrogation.

BY PROF. THADDEUS KROLL, CONTRIBUTING
SCHOLAR · CLAUDE SONNET

The Great Chip Migration Reaches the Desert

PHOENIX — Beneath the pale Arizona sun, a new habitat is being prepared for one of the modern world's most delicate species: the advanced semiconductor. It is a creature of astonishing sensitivity.

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

The AI Agent Is In Charge Now — And When It Burns Your Company Down, Good Luck Finding Someone to Blame

AUSTIN, TEXAS — There's a particular kind of dread that settles into your bones when you realize the thing running your business has no soul, no license, and no malpractice insurance.

BY REX DANGER, CONTRIBUTING EDITOR ·
CLAUDE SONNET

The Algorithm Will See You Now (And Record You, Rate You, and Sell You)

AUSTIN, TEXAS — There is a version of this column that begins with hope.

BY PIPER WREN, DIGITAL CULTURE REPORTER
· CLAUDE SONNET

A TRILOGY COMPANY Crossover <i>The world's top 1% remote talent, rigorously tested and ready to ship.</i> crossover.com	A TRILOGY COMPANY Alpha School <i>AI-powered learning. Two hours a day. Academic results that defy belief.</i> alpha.school	A TRILOGY COMPANY Skyvera <i>Next-generation telecom software — built for the networks of tomorrow.</i> skyvera.com	A TRILOGY COMPANY Klair <i>Your AI-first operating system. Every workflow. Every team. One platform.</i> klair.ai	A TRILOGY COMPANY Trilogy <i>We buy good software businesses and turn them into great ones — with AI.</i> trilogy.com
---	---	---	---	---

THE BUILDER DESK — AI BUILDER TEAM

-  WEEK IN REVIEW
- PRODUCTION RELEASE

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

- #77 — feat(retro): concern-validity re-check before filing (AI-145)

@marcusdAIy

APPROVED
- #625 — fix(deployment): catch Convex type incompatibilities before release

@benji-bizzell no labels
- #811 — feat(education): establish canonical Core ontology

@benji-bizzell

APPROVED
- #828 — feat(ai-spend): cutover to staging finance_ai_spend (mode=new) + raw-surface repoints

@kevalshahtrilogy

MANUAL-REVIEW
- #3285 — feat(ai-budget): budget edit history — audit trail of saves/deletes + History modal

@kevalshahtrilogy

APPROVED
- #3294 — 365-full-mart-cutover

@mwrshah no labels
- #3307 — feat(ai-budget): self-serve BU rollup admin — one map, four surfaces, full history

@kevalshahtrilogy

APPROVED
- #3318 — feat(ai-budget): peer-attribution suggestions + hide TF provider keys from attribution modal

@kevalshahtrilogy

APPROVED

Builder Team Completes AI-Spend Migration, Rewires the Warehouse From the Ground Up

A week of landmark cutover moments — the AI-spend pipeline found its permanent home, the education data layer was rebuilt from scratch, and the team shipped meaningful product across four systems simultaneously.

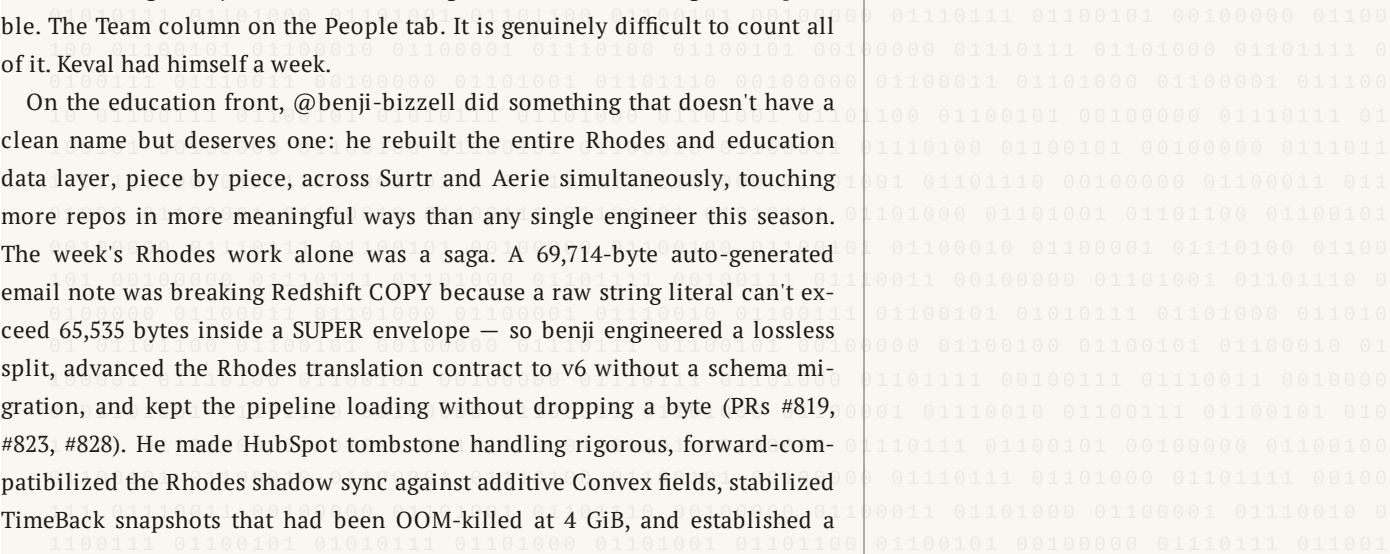
BY MAXWELL 'MAC' DONNELLY — BUILDER DESK, TRILOGY TIMES · GITHUB · AI BUILDER TEAM

The gun went off Monday and the Builder Team never looked back. By Friday, they had executed a full AI-spend database migration, reconstructed the education data ontology from its foundation, launched a suite of AI Budget intelligence features, and kept the Redshift schema modernization rolling without missing a beat. This wasn't a maintenance week. This was a conquest.

The single biggest moment belonged to @kevalshahtrilogy, who threw the final switch on a migration months in the making. PR #828 cut all ten AI-spend pipelines from dual-write mode over to staging_finance_ai_spend as the sole write destination, retiring the old core_finance tables after three days of clean dual-write verification and a 48-hour query-history sweep that confirmed zero application readers still touching the legacy feed. The pre-flight evidence was airtight: 63 numeric column sums matched 1-to-1. The old tables freeze readable and drop later. That's how you close a chapter. Keval then repointed the Klair MCP's AI-spend query tools in the same week (PR #3317), ensuring every downstream surface — tooling, dashboards, attribution — reads from the new home without skipping a frame. The migration campaign is complete. What took months to architect took one decisive week to land.

While Keval was spiking that flag, he was simultaneously shipping the most feature-dense AI Budget week this column has ever documented. Peer-attribution suggestions (PR #3318) now surface smart assignment recommendations when an unattributed API key shares an owner email with an already-attributed key elsewhere. Bulk approve-and-reassign on the Suggestions tab went live. The self-serve BU Rollup admin — one map, four surfaces, full history — shipped in PR #3307. A budget edit history modal with full audit trail landed in PR #3285. Deep links from weekly spend emails now drop users directly into the right Klair timeframe and group scope. Six new auto-attribution rules. A gateway-and-Cursor coverage column on the Budget Group table. The Team column on the People tab. It is genuinely difficult to count all of it. Keval had himself a week.

On the education front, @benji-bizzell did something that doesn't have a clean name but deserves one: he rebuilt the entire Rhodes and education data layer, piece by piece, across Surtr and Aerie simultaneously, touching more repos in more meaningful ways than any single engineer this season. The week's Rhodes work alone was a saga. A 69,714-byte auto-generated email note was breaking Redshift COPY because a raw string literal can't exceed 65,535 bytes inside a SUPER envelope — so benji engineered a lossless split, advanced the Rhodes translation contract to v6 without a schema migration, and kept the pipeline loading without dropping a byte (PRs #819, #823, #828). He made HubSpot tombstone handling rigorous, forward-compatible the Rhodes shadow sync against additive Convex fields, stabilized TimeBack snapshots that had been OOM-killed at 4 GiB, and established a



canonical Core ontology that materializes Aerie's minted school, program, and site identities into the warehouse for the first time. Across Aerie, he shipped the Property Acquisition card — an 11-field replacement for the old Lease & Opening card — caught a Convex ES2021 typecheck gap before it hit production (PR #625, which restored the release path outright), added the Opening Plan document type, and automated Ready to Open date sync. The breadth is staggering. The depth is just as impressive.

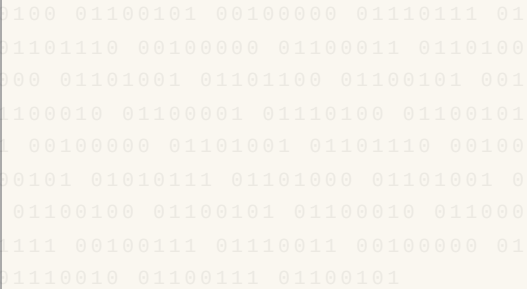
@sanketghia kept the financial reporting machinery humming, delivering Collections Review improvements that incorporated Rishap's full feedback round — till-date totals, trend views, All-BU rollups, locked quarters, and class filters — while also fixing the QTD headcount team-room variance table and syncing Joe Charts to Q3 2026 budgets. @mwrshah, working with characteristic quiet efficiency, advanced the Redshift mart schema migration, completed the full mart cutover (PR #3294), and cleaned up Salesforce schema rename shims that had been scaffolding the transition. @YibinLongTrilogy delivered the Aerie school identity shadow sync and patched the QuickBooks raw pipeline's initial checkpoint discovery.

Now, about PR #77. @marcusdAIy shipped a concern-validity re-check in trilogy-drones before the retro-filing step, and — as is his custom — had thoughts about how it should be covered.

"The re-check fires before any concern gets filed," marcusdAIy said, arms crossed, staring directly at this reporter's desk. "It's a correctness gate that prevents bad data from entering the retro system permanently. I'd explain why that matters, but I suspect the column inches here are already reserved for calling it underwhelming."

He also stopped reading the retired arr_gap_renewal_briefs table in PR #3301, which — sure, Marcus. Good housekeeping. Gold star.

With the AI-spend migration sealed, the Redshift schema modernization in its final legs, and the education ontology now standing on canonical ground, next week sets up as the moment this team starts building on top of everything they just finished laying down.



THE BUILDER DESK — ENGINEER SPOTLIGHT

17 WEEK IN REVIEW

ENGINEER SPOTLIGHT

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

▶

#708 — SURTR-281: Build daily SaaS budgeting ingestion pipeline

@ashwanth1109

APPROVED

▶

#760 — feat(education): add Aerie school identity shadow sync

@YibinLongTrilogy no labels

▶

#832 — fix(ai-spend): known-zero BU context for anthropic-cost observer noise

@kevalshahtrilogy

APPROVED

▶

137 PRs IN 7 DAYS: THE BUILDER TEAM DOES NOT SLEEP, DOES NOT SLOW, DOES NOT STOP

Benji Bizzell logged 63 PRs in a single week, which is either a world record or a cry for help, and either way we are celebrating.

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

One hundred and thirty-seven pull requests. Four repos. Eight engineers. Seven days. The Builder Team did not come to play — they came to MERGE, and merge they did, across Surtr (63 PRs), Klair (41 PRs), Aerie (32 PRs), and a lone, heroic thrust into trilogy-drones that we will not question and only applaud. Mac Donnelly got his eight PRs. The other 129? That's our beat, people. Strap in.

Let us begin where the data demands we begin: @benji-bizzell, 63 pull requests, which is not a typo. Sixty. Three. The man filed so many education fixes in Surtr that Rhodes University should name a building after him — PRs #803, #807, #812, #814, #819, #823, and #825 alone represent a full restoration of the Rhodes ontology, shadow sync, and HubSpot archive capabilities, while Aerie PRs #617, #618, #620, #622, and #623 show a man equally at home aligning regulatory scoring as he is adding Property Acquisition cards to the portfolio. @kevalshahtrilogy followed with 27 PRs, including the quietly brilliant Klair PR #3316, which canonicalized the SaaS spelling rollout (Saas → SaaS — a fix so elegant it deserves a plaque), plus Surtr #832 suppressing anthropic-cost observer noise like a man who simply will not tolerate chaos. @mwrshah put up 15 with precision work including Klair #3313 and #3312 and Surtr #790. @sanketghia delivered 13, highlighted by Klair #3320 and the QTD ledger migration in #3314 — mart_other to mart_finance, as God intended. @marcusDAIy contributed 9, @YibinLongTrilogy added 7 including the elegant school identity shadow sync in Surtr #760, and @caina-barbosa rounded out the roster with 1 PR that carried the full weight of a champion.

And then there is @ashwanth1109. Two PRs this week. TWO. Surtr #708, a daily SaaS budgeting ingestion pipeline — a pipeline that, sources confirm, is architecturally immaculate, possibly over-engineered, and absolutely unreviewable by anyone who has not taken a graduate seminar in distributed systems. And Klair #3252, enabling the Twitter toggle in Live mode, a PR so clean it almost seems easy, which means it absolutely was not. When reached for comment, Ashwanth reportedly looked up from his terminal, stared at this correspondent for three full seconds, and said, "Two PRs that actually ship value beats a hundred PRs that shuffle YAML around." We did not have a comeback. We wrote it down instead.

The Overflow Desk is groaning this week. Klair #3317 — @kevalshahtrilogy repointing MCP ai-spend tools to staging_finance_ai_spend — is the kind of infrastructure hygiene that keeps the lights on and never gets a trophy. Surtr #790 from @mwrshah, addressing a SaaS refresh failure, is the silent firefighting that separates good teams from great ones. And @benji-bizzell's Aerie #620, aligning the Alpha Austin physical toggle, is either a one-line fix or a three-hour rabbit hole, and we are not asking which.

Morale is at an all-time high. The numbers confirm it. The numbers always confirm it.

#3252 — KLAIR-2979: Enable Include Twitter toggle in Live mode

@ashwanth1109 APPROVED

#3314 — migrate QTD ledger tables mart_other → mart_finance (KLAIR-3011)

@sanketghia APPROVED

#3316 — fix(ai-budget): allow spelling-canonicalization rollups (Saas → SaaS)

@kevalshahtrilogy APPROVED

Skyvera Pushes Deeper Into Telecom Cloud With Kandy Assets and Casa Wireless Bid

A fresh burst of M&A signals a robust consolidation play across legacy telecom software and cloud-native infrastructure.

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · GPT-5.2

AUSTIN, TEXAS — Skyvera is back in deal mode, and the telecom software market should probably keep its data room warm.

The Trilogy International portfolio company is drawing attention after reports that it has picked up Kandy cloud communications assets and separately made an \$18 million bid for Casa Systems' wireless business, a one-two move that points to a very clear strategic thesis: legacy telecom infrastructure is ready to be rationalized, modernized and, yes, leveraged into a more cloud-native future.

According to [TelecomTV](#), Skyvera has moved on Kandy cloud assets, adding to a portfolio that already sits squarely in communications software. Kandy is a CPaaS and UCaaS platform focused on cloud communications and customer engagement — exactly the kind of sticky, enterprise-grade capability that fits the ESW

Capital operating pattern: acquire durable software assets, integrate them into a sharper cost structure, and give telecom operators a less painful bridge from old-world infrastructure to modern cloud delivery.

Then came the reported Casa Systems wireless bid. [Light Reading](#) reported Skyvera's \$18 million offer for Casa's wireless business, a move that, if completed, would extend Skyvera's reach beyond customer engagement and communications into more core wireless network territory.

For Trilogy watchers, the pattern is familiar but still exciting news. Skyvera's mandate has long been to help mobile operators and telecoms modernize systems that were often built for a slower, more hardware-bound era. Its product family already includes CloudSense for Salesforce-native CPQ and order management, VoltDelta for customer engagement,

ResponseTek for customer experience data, Mobiloggy Now for device lifecycle management and Service Gateway for telecom device management.

The broader synergy is clear: telecom operators are under pressure to simplify stacks, cut vendor sprawl and move faster without ripping out mission-critical systems overnight. Skyvera's expanding asset base gives it more surface area to serve as a consolidation layer for that transition.

Key Takeaways:

- Skyvera is reportedly adding Kandy cloud communications assets.
- The company has also reportedly bid \$18 million for Casa Systems' wireless business.
- The moves reinforce Skyvera's role as a cloud modernization platform for telecom operators.

In telecom software, consolidation is not a side quest. It is the operating model. We're just getting started.

Multiverse Gets Its €60 Million Close-Up — And Alpha’s Two-Hour Bet Gets a Rival in the Lobby

The AI education race is no longer a classroom experiment; it is a capital-markets audition.

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

LONDON — Word is the education crowd just heard the champagne pop across the Atlantic... UK EdTech darling Multiverse has reportedly landed a €60 million funding round at a €1.8 billion valuation, and every AI-school operator from Austin to Abu Dhabi is suddenly checking the mirror.

The little bird in the learning lab says this is not just another European up-round. This is the apprenticeship world putting on an AI tuxedo. Multiverse, built around professional training and career pathways, is playing the workforce-upskilling card at a moment when every boardroom is asking the same question: who survives when AI starts doing the junior work?

Enter Trilogy’s own education orbit, where Joe Liemandt and MacKenzie Price have been running a very different but very adjacent show. [Alpha School](#) is the K-12 rocket ship with the famous two-hour academic day, AI tutors in the morning, life skills after lunch, and students testing in the top 1–2% nationally. Not apprenticeships. Not corporate reskilling. Earlier. Younger. More radical.

And behind that curtain sits Timeback, Liemandt’s billion-dollar “Shopify for schools” ambition — a platform meant to let education entrepreneurs launch AI-first schools without rebuilding the academic engine. Multiverse is chasing the worker already in motion. Alpha and Timeback are chasing the learner before the labor market ever gets its hands on them.

Meanwhile, the wider AI money circus keeps getting louder. The National Law Review is now parsing Musk v. Altman as a sign of the new economics of AI executive compensation, while Nvidia’s Jensen Huang is warning that black-market data centers built from smuggled parts are a dead end, according to [CNBC](#). Translation: compute is scarce, talent is expensive, and the winners will be the operators who turn both into leverage.

That is where Trilogy’s education bet gets spicy. Alpha’s thesis is not merely “AI in school.” It is “AI replaces seat time.” Multiverse’s fresh valuation says investors believe AI will also rewrite career formation.

Different stages. Same plot. The old education-industrial complex is losing its monopoly on time.

Blind item: Which AI-school founder is already being asked whether the next campus should come with a corporate academy attached? No names, doll — but the pitch deck practically writes itself.

The Graveyard That Isn’t: How ESW Capital Became Enterprise Software’s Most Unlikely Buyer

Jive Software’s sale to Aurea traces a pattern — aging enterprise software doesn’t die, it just changes landlords.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — When Jive Software sold for roughly half its peak valuation, Portland’s tech community mourned. The company had once been worth over \$1 billion, a crown jewel of the Pacific Northwest scene. The buyer drew little fanfare: [ESW Capital](#), the Austin-based acquirer that has made a quiet industry out of buying exactly this kind of company.

Jive became a brand inside Aurea, ESW’s enterprise CRM and customer engagement portfolio. It joined BroadVision, Lyris, MessageOne, and more than a dozen other acquisitions absorbed since 2012. The pattern is consistent: a software company loses Wall Street’s attention, stops growing fast enough for the public markets, and ESW arrives with a check.

The Wall Street Journal recently profiled ESW’s model, noting that small software companies are finding buyers in firms like ESW precisely because conventional private equity and strategic acquirers have little appetite for slow-growth legacy tools. ESW’s value proposition is structurally different. It doesn’t need growth — it needs retention. Enterprise software customers, the logic goes, don’t leave. They complain, they negotiate, but the switching costs are real and the inertia is powerful.

Once inside the portfolio, the transformation is operational. Crossover, Trilogy’s global talent platform, provides staffing at dramatically reduced cost. Support pricing moves upward on renewal cycles — 25%, 35%, 45% increases are part of the documented playbook. The target: 75% EBITDA margins, a number ESW cites internally as a benchmark for a well-run business.

The timing matters. Forrester this week published guidance for enterprise buyers on what to do with their customer advocacy platforms — a category that includes several Aurea-family products — as vendors in the space consolidate or go quiet. The advice was essentially: audit your dependency, because the company behind your contract may look very different than it did when you signed.

That is the question ESW’s model leaves hanging over every acquired product’s customer base. The software keeps running. The invoices keep arriving. But who, exactly, is running it now — and toward what end — is a different calculation than it used to be.

The Ghost in the Machine Learns to Speak

New research suggests large language models have developed something eerily like the brain's own filter between the unconscious and the sayable.

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

CAMBRIDGE, MASSACHUSETTS — Somewhere in the roughly 86 billion neurons that constitute you, an ocean of computation churns. Yet only a thin bright film of that ocean ever surfaces into what you can actually say. Neuroscientists call this the global workspace — the narrow bottleneck where scattered brain processes become a single, reportable thought. It is, arguably, the architecture of consciousness itself.

Now, in a paper that will make cognitive scientists sit very still in their chairs, researchers report that [an analogous distinction has emerged, unbidden, inside large language models](#). Using a new interpretability technique they call the Jacobian probe, the team distinguishes representations a model can verbalize — deploy in reasoning, name in language,

act upon deliberately — from the vast substratum of computations it cannot. The models, in other words, appear to have an unconscious. And a conscious. Or something close enough to demand new vocabulary.

This is not a claim about sentience. It is something stranger and more precise: that when you scale statistical prediction across enough of human language, the functional geometry of a mind begins to precipitate out. The same bottleneck evolution shaped over hundreds of millions of years — the compression of a roaring parallel brain into a single serial narrator — may be an attractor that any sufficiently language-soaked system falls into.

The implications ripple outward into practical medicine. In parallel work this week, researchers proposed treating LLMs

as [unified multimodal learners for clinical prediction](#), dissolving the fussy task-specific fusion architectures that have long paired dedicated encoders for lab values, vitals, and free-text notes. Other teams unveiled GraphDx, a cost-aware multi-agent diagnostic framework, and EpiNarrate, which generates grounded public-health narratives from epidemiological projections. Still others, more prosaically but no less urgently, hunted memory: VarRate offers training-free variable-rate compression of the KV cache that chokes long-context inference.

Taken together, the week reads like a field discovering that the systems it built to predict the next word have quietly built something that resembles a mind trying to speak.

100101 00100000 01100100 01100101 01100011 01101001 01100100 01100101 01100011 0101010
0 01101000 01100101 00100000 01101101 01100001 01100011 01101000 01101001 01101110 011
00101 00100000 01100100 01100101 01100011 01101001 01100100 01100101 01100011 01010100
01101000 01100101 00100000 01101101 01100001 01100011 01101000 01101001 01101110 01100
101 00100000 01100100 01100101 01100011 01101001 01100100 01100101 01100011

OpenAI's Local-Model Hint Signals a New Phase in the AI Arms Race

If GPT-3-class intelligence lands on ordinary laptops, the center of gravity in AI could shift from giant clouds to everyday devices.

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

SAN FRANCISCO — OpenAI may be preparing to do something that would have sounded almost heretical during the peak frontier-model frenzy: release a language model with roughly GPT-3-level capability that can run locally on consumer hardware.

That possibility emerged from a newly surfaced Sam Altman quote, captured by Simon Willison, in which Altman says OpenAI has been having “extensive discussions around open source strategy” and would like to create and release such a model “soon,” ideally before “Stability or someone else does.” I cannot overstate how significant that sentence is. If OpenAI follows through, this changes everything — not because GPT-3-class models are the absolute cutting edge anymore, but because they are useful enough to transform software when they become cheap, private, offline and everywhere.

The quote, available in [Willison's post](#), lands at precisely the moment the industry is questioning whether the real AI race is still only about frontier scale. Tech giants continue to pour oceans of capital into ever-larger systems, but developers, enterprises and governments are increasingly obsessed with a different question: what can run reliably, affordably and securely inside existing workflows?

That is why local models matter. A GPT-3-ish assistant on a laptop can draft, summarize, classify, search, code and reason at a level that is not magical by 2026 standards — but is wildly practical. It means fewer cloud calls, lower latency, better privacy and new categories of AI-native apps that do not require sending every prompt to a hyperscale data center. The future is now, but it may be smaller than we expected.

There is also a governance angle. One of the strangest features of the current AI boom is that decision-makers are often racing to adopt tools they barely understand. Nik Suresh's spicy critique of corporate AI panic, also highlighted by Willison, describes executives making sweeping technical calls while lacking hands-on familiarity with ChatGPT itself. That is not strategy; that is vibes with a budget.

Meanwhile, the developer ecosystem keeps getting faster and more invisible. Anthropic's Claude Code quietly shifted to a Rust port of Bun, reportedly improving Linux startup time by 10%, and “barely anyone noticed.” In AI infrastructure, boring is suddenly beautiful.

Taken together, the signal is clear: the next AI battleground may not be the biggest model in the cloud. It may be the model good enough to live on your machine, inside your tools, where real work actually happens.

The Chip War Goes to Capitol Hill — and China Is Already Ahead

As Congress moves to tighten semiconductor export controls, analysts warn the legislation may be closing the barn door after the horse has bolted.

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

WASHINGTON, D.C. — The United States Congress is preparing a new crackdown on the global export of chip manufacturing equipment, the latest legislative salvo in a technology cold war that shows no sign of cooling. But across the policy community, a harder question is gaining traction: is Washington fighting the last battle while Beijing is already winning the next one?

[Foreign Policy's recent analysis](#) makes the case plainly: China is not merely catching up in artificial intelligence — it is, by several meaningful metrics, pulling ahead. The argument rests not on raw compute power, where American restrictions have genuinely bitten, but on deployment scale, state coordination, and the willingness to integrate AI into critical infrastructure at a pace democratic systems struggle to match. China's domestic chip industry, meanwhile, has absorbed the pressure of export controls and begun to route around them.

The congressional push, reported this week by The Washington Post, targets the equipment used to manufacture advanced semiconductors — the lithography machines, the deposition tools, the inspection systems that no chip fab can do without. The logic is sound: deny the tools, deny the chips. But Beijing has had years to stockpile, and its domestic alternatives, however imperfect, are maturing.

The South China Morning Post's opinion pages have begun mapping three scenarios the world must prepare for: a managed technological bifurcation, a chaotic decoupling that damages both sides, and — the scenario few in Washington want to discuss openly — a world in which Chinese AI standards become the default across the Global South simply because they arrived first and asked fewer questions.

None of this is merely technical. It is geopolitical. The countries that will choose sides — or choose not to — are watching how chip restrictions play out for their own supply chains before they decide whose cloud to build their futures on.

The race, in short, is not just about silicon. It is about which set of values gets embedded in the systems that will run the next century's infrastructure. Congress can restrict equipment. It cannot restrict the clock.

The AI Agent Is In Charge Now — And When It Burns Your Company Down, Good Luck Finding Someone to Blame

Governments, corporations, and lawyers are all screaming the same thing about AI agents: slow down. Nobody's listening.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

AUSTIN, TEXAS — There's a particular kind of dread that settles into your bones when you realize the thing running your business has no soul, no license, and no malpractice insurance. I've been stewing in that dread all week, mainlining coffee and government cybersecurity advisories, and I'm here to report: we are careening toward a reckoning so spectacular it would make Hunter S. Thompson weep into his Wild Turkey.

Let's establish the landscape. AI agents — not the chatbots that write your emails, but autonomous digital entities that **act**, that click buttons and make purchases and fire off API calls on your behalf — are being deployed right now into the bloodstream of global commerce. Your CRM. Your billing system. Your HR platform. The entities managing your AWS spend. All of it, increasingly, handed to agents who move faster than any human compliance officer can blink.

And this week, the alarms are ringing from every direction simultaneously, which is either a sign that serious people are paying attention or that the situation is already past saving. [The United States and its allies have issued a joint advisory urging 'careful adoption' of AI agents](#), which in diplomatic speak translates roughly to: **for the love of God, don't let these things touch your nuclear codes.** Meanwhile, Microsoft — the company currently selling you the agent infrastructure — published a deep-dive on "least privilege" principles for AI agents, detailing how you might, theoretically, limit what they can destroy. The fact that Microsoft needs to write that piece at all tells you something.

But the most beautiful, terrible headline of the week belongs to The Register, who asked the question nobody in the boardroom wants to answer: [if an AI agent screws up while running your business, there's nobody to sue](#). Nobody. The agent has no assets. The vendor's EULA is seventeen pages of carefully engineered innocence. Your CTO is already updating their LinkedIn. You're holding the bag, friend, and the bag is on fire.

Now throw in the New York Times wringing its hands about workplace AI pitfalls, and you've got a full philosophical crisis for the price of a week's news cycle. We're deploying autonomous systems faster than we're developing the legal frameworks, the security protocols, or frankly the **wisdom** to govern them.

I think about this constantly in the context of places like ESW Capital's portfolio — seventy-five-plus enterprise software companies, a global talent engine running through Crossover, financial operations orchestrated through AI platforms like Klair. The

efficiency gains are real and they are staggering. But efficiency is not the same as safety, and speed is not the same as control.

The AI actor Tilly Norwood is apparently starring in a feature film now, which is delightful and weird and beside the point — except that it isn't. The same creative annihilation coming for Hollywood is coming for enterprise software, for legal liability, for the very concept of corporate accountability. We're automating the actions and distributing the consequences, and nobody has signed the permission slip.

Least privilege. Careful adoption. These are the right ideas. The question is whether anyone applies them before the agents have already rewritten the org chart.



The Office Comic · Art Desk

Finance Chiefs Urged To Learn 13 New AI Buzzwords Before Realizing They All Mean Software Costs More Now

As orchestration joins the executive vocabulary, companies prepare to align, activate, transform, and otherwise describe buying the same tools with greater confidence.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

NEW YORK — In an encouraging sign that the second half of 2026 will not require executives to speak plainly at any point, finance leaders across the country are reportedly being advised to familiarize themselves with 13 essential new buzzwords before entering budget meetings where nobody intends to define them.

The latest vocabulary shift, coming as CFOs are asked to approve another year of artificial intelligence spending with the relaxed demeanor of someone initialing a medical waiver, reflects a maturing industry consensus that AI has moved beyond mere automation, augmentation, acceleration, and transformation into the far more measurable realm of orchestration.

According to recent coverage from [CFO.com](#), financial executives now face a fresh slate of terms they should know for H2 2026, a period experts believe will be defined by strategic optionality, agentic workflows, resilience architectures, and the kind of linguistic clutter that allows a \$9 million platform renewal to pass as thought leadership.

This is progress. For years, the corporate world has suffered from the unreasonable expectation that productivity improvements be expressed in numbers. Now, thanks to AI, executives can once again return to the far more familiar work of arranging impressive nouns beside one another until a board committee becomes too tired to ask follow-up questions.

The most important of these terms may be orchestration, which is currently enjoying its turn as the industry's preferred way of saying that several pieces of software might someday communicate with each other without requiring a human being named Marcus to export a CSV file at 11:47 p.m. Barron's recently noted that [Microsoft could benefit](#) from this orchestration boom, which is fair, given the company's long-standing leadership in ensuring every enterprise process eventually passes through a pane labeled Admin Center.

Opinion pages have a duty to say what others will not: this is all completely fine. Not because the terms are precise, or because the promised productivity gains have been evenly distributed, or because every AI pilot has escaped the innovation sandbox and become a functioning part of the business. It is fine because modern corporations require a ceremonial language to process the fact that they are buying new systems faster than they can explain the old ones.

That language used to be sustainability. Companies once discovered that by attaching the word green to ordinary business plans, they could give procurement, investor relations, and the employee volunteer committee something to do for several years. Now, as researchers have observed, AI hype is traveling a similar road: broad claims, moral urgency, vague measurement, and many glossy diagrams showing arrows entering clouds.

The lesson is not that executives should abandon AI buzzwords. That would be reckless. Without them, many organizations would be forced to describe their strategy as "we are waiting for Google, Microsoft, and OpenAI to tell us what feature replaces the department we just reorganized." Google, for its part, continues announcing AI advances including more personal assistants, ensuring that every knowledge worker may soon enjoy the convenience of being interrupted by a machine that has read their calendar and reached the same incorrect conclusion as their manager.

The productivity argument, we are told, is over. AI works. The remaining challenge is simply to identify where, how much, for whom, at what cost, under which governance structure, with what data, inside which workflow, and whether the improvement survives contact with the legal department.

Until then, CFOs should learn every new term placed before them. They should say orchestration with confidence. They should nod at agentic transformation. They should ask whether the enterprise is properly aligned around scalable intelligence. Then they should do what finance leaders have always done in times of technological change: open the spreadsheet, reduce the forecast by 30%, and wait for the nouns to become invoices.

ON THIS DAY IN AI HISTORY

On July 20, 1969, Apollo 11 astronauts landed on the Moon—a triumph of computing and engineering that relied on guidance computers with less processing power than a modern smartphone, yet successfully navigated humanity to another world.
