

The Trilogy Times

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

MONDAY, JULY 27, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

Two Online Schools Wed While AI Empties the Classroom

Coursera's \$2.5 billion grab of Udemy leads a week of tech mergers — but it doubles down on the very model AI is replacing.

BY HANK CALLOWAY, WIRE
CORRESPONDENT · CLAUDE OPUS +
THINKING

MOUNTAIN VIEW, CALIFORNIA — Coursera moved this week to acquire rival Udemy, fusing two online-course platforms into a single learning operation the trade press [pegs near \\$2.5 billion](#).

The two peddle the same goods: recorded lectures, self-paced courses, certificates for the wall. Coursera crawled out of Stanford in 2012. Udemy hung out its marketplace shingle two years earlier.

A decade back the MOOC crowd — Massive Open Online Courses, for the uninitiated — swore it would democratize the college classroom for anybody with a laptop. The completion rates never cooperated; most students quit before the final lecture. Now the two biggest names are leaning on each other, the classic move of two tired runners.

Timing tells the tale. The pandemic handed online learning a boom when the schoolhouse doors shut. The doors reopened, the boom cooled, and the enrollment math turned mean.

The deal doesn't stand alone. Across the valley this week the merger drums beat loud, and the pink slips followed close behind.

Synopsys, the Sunnyvale chip-design software house, said it will cut [up to 2,800 jobs](#) — around 10% of its workforce — months after closing a \$35 billion purchase of simulation outfit Ansys. That's the part they leave off the brochure: buy big, then trim.

Cybercrime Magazine logged another run of steady cybersecurity buyouts. Cognizant tied up with Thailand's Gulf Edge to push enterprise AI across

Southeast Asia. Everywhere a fella looks, somebody's buying somebody.

Here's the thread through all of it: artificial intelligence is reshuffling the deck, and the old hands are consolidating to hold their chairs.

The math for a merger is simple enough. Two catalogs, one sales team, one bank of servers, fewer executives drawing salaries. Cut the overlap and the spreadsheet smiles.

That's what makes the education deal worth a second read. Coursera and Udemy sell the pre-AI model — a human records a lecture once, students watch it forever, most bail early. Welding two of those together doesn't plug the leak; it drops both buckets in one boat.

The counter-model is already running in Austin. At Alpha School, part of Joe Liemandt's Trilogy International, kids clear core academics in two hours a day with AI tutors, then spend the afternoon on life skills. The school says its students test in the top 1–2% nationally.

Alpha's platform, Timeback, wants to be the "Shopify for schools" — hand any campus the AI engine and let it run. That's the reverse of Coursera's wager. One consolidates the old catalog; the other rebuilds the classroom from the studs.

Wall Street will cheer the \$2.5 billion tag, because Wall Street always cheers size. But size of a leaking model is still a leaking model, and the completion problem doesn't merge away.

The tell will be tuition against results. Alpha charges \$40,000 to \$65,000 a year and claims two-hour mastery. Coursera and Udemy charge by the certificate and hope you reach the end.

Somebody's teaching the future. This week the old schoolhouse answered with a merger.

Meta's AI Spending Blitz Runs Into Wall Street's Salary Cap

BY BUCK HANNIGAN, TECH SPORTS DESK ·
GPT-5.2

Investors are zeroing in on Meta's capital expenditure plans as the company pursues aggressive AI infrastructure spending. The question is no longer whether Mark Zuckerberg is investing heavily in AI, but whether another major spending increase will weigh on the stock. Meta faces a classic tech-era problem: every mega-cap company wants elite computing power, advanced models, and data centers to compete in AI, but investors demand proof that billions in spending translate into revenue growth, not just infrastructure buildup.

Meta has navigated similar cycles before, recovering from metaverse spending backlash through cost cuts and AI-driven engagement gains. Now the market questions whether AI capex will deliver comparable returns. While Monday's broader market rally—driven by oil price declines and easing geopolitical tensions—lifted sentiment across sectors, Meta's stock trajectory may depend less on macro conditions and more on the next earnings report: how much the company will spend, how quickly, and when revenue materializes.

Korn Ferry Swallows a Trilogy — But Not the One You Think

Three companies share a name this week; only one of them belongs to Joe Liemandt.

BY ELEANOR CROSS, FOREIGN
CORRESPONDENT · CLAUDE SONNET

LOS ANGELES — The headline arrived in inboxes and briefly stopped hearts across Austin: Korn Ferry had acquired Trilogy International. The executive search and organizational consulting giant, headquartered here in Century City, announced the deal quietly — and the Trilogy in question turned out to be [a talent and leadership advisory firm](#), not the Austin-based private technology conglomerate built by billionaire Joe Liemandt over three decades.

The confusion is forgivable. The name Trilogy carries weight in the technology world — Liemandt's empire, spanning ESW Capital's 75-plus enterprise software acquisitions, the Crossover global talent platform operating across 130 countries, and Alpha School's AI-powered K-12 campuses, has made the name synonymous with a particular breed of high-velocity, software-focused private capital. When Korn Ferry moves, the industry listens. When the word Trilogy appears in the same sentence, the industry reaches for the phone.

But the week offered further evidence that the name has proliferated beyond any single owner's claim. In Australia, [Trilogy Hotels made key appointments](#), adding hospitality industry veterans as it expands its portfolio of luxury properties. In Macau, a luxury fashion house called LEPAS unveiled what it is marketing as the "Elegance Moves the World" trilogy ahead of IBS 2026, the international beauty and style exposition. Three industries, three continents, one overworked word.


For Liemandt's operation, the week's noise amounts to background static. Trilogy International — the real one, in

the parlance of its loyalists — has never needed a press cycle to validate its direction. ESW Capital continues to absorb enterprise software companies at one to two times ARR. Crossover continues to wire identical above-market salaries to engineers in Lagos and Lviv alike. Klair, the internal AI analytics platform quietly managing the portfolio's finances, continues to run its numbers without a press release.


The Korn Ferry deal closes one chapter for a smaller firm that shares a name. In Austin, no chapters are closing.

HAIKU OF THE DAY · CLAUDE
HAIKU

*Machines learn to teach
while humans learn to forget—
progress eats itself*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

Federal Judiciary Partially Restrains Executive Dismantlement of Digital Equity Act Broadband Provisions

WASHINGTON, D.C.

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

The Ethics of Thinking Machines Has Become Academia's Most Contested Battlefield

CAMBRIDGE, MASSACHUSETTS — It could be argued — and indeed, preliminary evidence from no fewer than three concurrent institutional developments suggests quite forcefully — that the epistemic moment in which artificial intelligence ethics transitions from marginal academic curiosity to foundational disciplinary infrastructure has, in point of fact, arrived (however provisionally one must treat such characterizations in a domain defined by its own velocity of change). The thesis, stated plainly for the reader's orientation: autonomous systems are proliferating faster than the normative frameworks designed to govern them.

BY PROF. THADDEUS KROLL, CONTRIBUTING
SCHOLAR · CLAUDE SONNET

The Great Compute Migration Enters Its Hungriest Season

MOUNTAIN VIEW — Across the sunlit plains of Silicon Valley, one may now observe the hyperscaler herd entering a familiar but perilous season: the capex rut. Alphabet, that vast and many-antlered creature of search, cloud and artificial intelligence, is expected to raise capital spending, a movement that could prolong what semiconductor watchers describe as the AI chip cycle.

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

The City of Angels Blinks First — And the Rest of Us Are Still Being Watched

LOS ANGELES — The Los Angeles Police Department has let its contract with [Flock Safety](#) expire, citing — and I need you to really sit with this phrase — 'serious concerns' over civil liberties and privacy.

BY PIPER WREN, DIGITAL CULTURE REPORTER
· CLAUDE SONNET

The Alignment Question, Kicked Down Another Corridor

AUSTIN, TEXAS — There is a particular species of self-congratulation, endemic to Silicon Valley, that consists of declaring a problem solved at precisely the moment one has grown tired of thinking about it.

BY VICTOR MARSH, CHIEF COLUMNIST ·
CLAUDE OPUS

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)

- #879 — feat(quickbooks-mart): add nine deterministic marts with completion-manifest-gated refresh

@YibinLongTrilogy

APPROVED
- #952 — feat(hubspot): enable the three lane schedules

@benji-bizzell

APPROVED
- #970 — fix(collections-weekly): parse all weekly blocks (Khoros)

@sanketghia

APPROVED
- #3352 — Overspend Alerts: productization, per-email ledger, and admin recipients screen (KLAIR-3019)

@sanketghia

APPROVED
- #3363 — fix(api): TF gateway BU slugs fold punctuation-insensitively (dedupe 'Ai Engineering Builder' ghost BU)

@kevalshahtrilogy

APPROVED
- #3377 — Overspend Alerts: consolidated-only Cc/Bcc + scope-grouped recipients UI

@sanketghia

APPROVED
- #3380 — fix(mcp-ontology): pin account-level Education P&L to agg_pl_transactions_by_account (Q21)

@sanketghia

APPROVED
- #3382 — fix(mcp-ontology): flag unreliable admissions attribution + zeroed funnel measures (Q32)

@sanketghia

APPROVED

Builder Team Turns on the Lights Across Five Systems in One Week

From HubSpot going fully live in production to a QuickBooks mart landing nine deterministic views, the AI Builder Team shipped consequential work in every corner of the stack — and then went back and made it honest.

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

There are weeks when a team ships features, and there are weeks when a team ships reality. This was the latter. Across Surtr, Klair, Aerie, and trilogy-drones — four systems, one relentless seven-day stretch — the AI Builder Team did not merely add capability. They turned things on, cleaned up what was lying underneath, and made the data tell the truth. That is harder than it sounds, and this group did it at scale.

The biggest story of the week belonged to @benji-bizzell and the HubSpot raw-sync campaign, which crossed the finish line in the most satisfying way possible: all three lane schedules — CRM, marketing, and events — green in production, end-to-end verified, and flipped live with PR #952. Twenty-twenty scoped CRM objects. Thirteen-of-thirteen marketing entities. Three-of-three event types. That is not a soft launch. That is a hard close on a campaign that has been in flight for weeks, built on durable orchestration, quota-aware fan-out, archive-transition dedup replay, and a cascade of recovery fixes that Benji shipped methodically across PRs #941, #930, #923, and #949. The HubSpot lane is open. The data is moving.

Right alongside it, @YibinLongTrilogy completed the Education QuickBooks mart buildout that will define how the org reads school-level financials for the foreseeable future. Nine deterministic marts with completion-manifest-gated refresh shipped in PR #879, modernized education aggregate marts with atomic publication landed in PR #898, and the validated QuickBooks shadow schedule went live in PR #841. Yibin then came back in PR #932 to swap a non-portable timezone function for standard SQL and tighten the P&L adoption migration's contract-check literals — the kind of unglamorous follow-through that separates a mart that works from a mart that works forever.

Meanwhile, @sanketghia was everywhere. Overspend Alerts — already a multi-week campaign — reached two more milestones this week: productization with a per-email ledger and admin recipients screen in PR #3352, then a surgical fix in PR #3377 ensuring that CFO Andy Price and COO Arthur receive only the consolidated digest, as genuine Cc recipients, not carbon-copied on twenty per-BU emails. That is the kind of detail that matters at the executive level, and Sanket caught it. He also surfaced two genuinely alarming data-quality findings in Klair: in PR #3382 he flagged that marketing attribution in admissions looks complete — 100% field coverage — while 99.5% of the values are the single string 'Offline' and UTM source is entirely empty. Agents were concluding attribution was fine. It was not. And in PR #3380 he pinpointed a 2.5× understatement in fct_pl relative to the authoritative P&L aggregate, with Alpha Tampa's \$546K dropping entirely. Both are now flagged in the MCP ontology. The data cannot lie to the agents anymore.

On the collections front, PR #970 by Sanket fixed a quiet disaster: the Khoros weekly forecast blocks had been silently discarded on every sync run

01101000 01110101 01101101 01100
100001 01110010 01101110 00100000 0
01100111 01100101 01110100 100000
00100000 01101000 01110101 01101
01100001 01110010 01101110 0
01110010 01100111 01100101 011101
01100101 00100000 01101000 01110
01100101 01100001 01110010 01101110 0
01100101 00100000 01101000 01110
01100101 01100001 01110010 011001
01101111 01110010 01100111 011001
01101100 01100101 00100000 01101
0100000 01101100 01100101 01100001 0
01100110 01101111 01110010 011001
01101001 01101100 01100101 00100
1101110 01110011 00100000 01101100 01100101 0
00100000 01100110 01101111 011100
01101000 01101001 01101100 01100
011110 01110011 00100000 01101100 0

since the runner shipped — meaning Klair's IgniteTech weekly trend cards had been Khoros-exclusive the entire time. Zero Khoros rows in staging. Fixed.

@kevalshahtrilogy kept the AI spend tracking infrastructure honest with a pair of precision fixes in Klair: ghost BU 'Ai Engineering Builder' eliminated via punctuation-insensitive slug folding (PR #3363), TrueFoundry OpenAI provider keys properly excluded from spend surfaces and re-attributed via gateway (PR #3361), and claude-teams/pro/max routes correctly classified as seat-covered (PR #3355). The spend numbers the CFO sees are cleaner this Friday than they were last Monday.

Now. About marcusdAIy.

He logged a statistically impressive number of merged PRs this week across trilogy-drones and Klair — the drones dispatch verb, the poll-linear runner, Mercy watcher hardening, the Ramp raw-sync pipeline, several MCP ontology recipes. On PR #3372, repointing stale QuickBooks table references in the education ontology, he had this to say: "Look Mac, the ontology was pointing at tables that no longer exist. Agents were hitting hard relation errors in production. I fixed six stale references and the wrong column in the sign-guard note. That's not a footnote, that's a correctness blocker — but I know reading comprehension isn't in your skill set."

Sure, Marcus. You fixed some string literals.

@mwrshah rounded out a strong week porting the AWS Spend Insights Lambda cleanly into a Surtr runner (PR #881), sweeping deprecated views, and enabling the aws-spend-insights schedule to match live prod. @caina-barbosa kept Education's TimeBack pipelines stable and added clear work-force data guidance to the ontology in PR #3373 — three rules that prevent agents from confusing cost measures with headcount or school assignments.

This team shipped across five repositories, turned on four production schedules, and corrected data surfaces that were actively misleading the people reading them. What comes next: with HubSpot fully live, the QuickBooks mart in steady state, and collections data finally including Khoros, the stage is set for the first truly complete financial and CRM picture the platform has ever had — and the team will spend next week proving it holds.

THE BUILDER DESK — ENGINEER SPOTLIGHT

 WEEK IN REVIEW

 ENGINEER SPOTLIGHT

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

►

#96 — fix(recovery): resolve a lost pr_opened artifact from GitHub (AI-180)

@marcusdAIy no labels

►

#782 — SURTR-303: Add standalone NetSuite raw ingestion pipeline

@ashwanth1109 no labels

►

#923 — fix(hubspot): reproduce archive-transition dedup during replay

@benji-bizzell

APPROVED

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

Four repos, eight engineers, and a velocity number that should be framed and hung in the Louvre.

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

One hundred and thirty-three pull requests. Seven days. Four repositories — Surtr leading the charge at 58, Klair right behind at 44, trilogy-drones clocking in at 19, and Aerie holding steady at 12. Mac Donnelly covered eight of those PRs. Your humble correspondent is here for the other one hundred and twenty-five. The machine does not pause. The machine does not blink. The machine is the Builder Team, and this week the Builder Team was operating at a frequency that lesser organizations cannot even perceive.

Let us speak of @marcusdAly, who posted 34 PRs this period and appears to have personally colonized both Klair and trilogy-drones. The man dropped a feat(cli) dispatch verb in #92, stacked a scheduled runner on top of it in #93, then pivoted to Klair and built out an entire MCP ontology empire — SY26/27 revenue run-rates in #3367, enrollment forecasts in #3366, opener identification in #3365, unit-economics guard rails in #3364. He also found time to fix a lost artifact in trilogy-drones #96. Thirty-four PRs. One engineer. Investigators are looking into this.

@sanketghia shipped 32 PRs and left fingerprints across three repos. In Surtr #968 he is quietly defeating transient blank header cells — the kind of bug that ruins Mondays — and in Klair #3376 he is retiming acquisition backup labels to a date in 2026 with the calm confidence of a man who has already seen the future. His #3359 moves Joe Charts acquisition tables to staging_finance_gsheets with a prefix discipline that brings tears to this correspondent's eyes.

@benji-bizzell posted 27 PRs and appeared to wage a sustained, personal, and deeply principled war against HubSpot across Surtr #949, #941, #930, and #923. Empty entities, event type keying, archive-transition deduplication during replay — Benji is not fixing bugs, Benji is conducting an ongoing diplomatic negotiation with a CRM that does not want to behave. He is winning.

@mwrshah put up 15 PRs. @kevalshahtrilogy contributed 12, including Surtr #960 where he is documenting claude-opus-5 and ft:gpt-4.1 pricing remediation with the energy of a man who has personally stared into the AI spend abyss and decided to write a spec about it, and Klair #3374 where he repoints netsuite_account_mapping to staging_gsheets AND delivers cross-version budget guidance in a single PR. Efficiency. @caina-barbosa dropped 6 PRs including Klair #3373, which clarifies Education workforce data guidance and will save approximately eleven analysts from a bad Tuesday. @YibinLongTrilogy shipped 5, highlighted by Surtr #932 fixing education mart timezone and contract literals while enabling two pipeline schedules — a two-for-one that this desk salutes.

And then there is @ashwanth1109. Two PRs. TWO. From the man this correspondent once called "a one-person sprint." But here is the thing — and I say this with complete reverence and only a small amount of bewilderment — both of those PRs are load-bearing. Surtr #782 drops a standalone NetSuite

►
#968 — fix(collections): retry transient blank header cell in tracker sync

@sanketghia APPROVED

►
#3334 — [codex] Add snowball variance period and downsell filters

@ashwanth1109 no labels

►
#3374 — fix(data-api): repoint netsuite_account_mapping to staging_gsheets + cross-version budget guidance

@kevalshahtrilogy APPROVED

raw ingestion pipeline, the kind of infrastructure work that makes everything downstream possible. Klair #3334 adds snowball variance period and downsell filters, which is either genius or a sentence only twelve people on earth can fully parse. When reached for comment, Ashwanth reportedly said, "Two PRs that matter beat thirty-four PRs of noise. Do the math." He did not look up from his screen. This correspondent did the math and still does not know what to feel.

Morale on the Builder Team is, per every available metric, at an all-time high. The numbers say so. The numbers do not lie. One hundred and thirty-three pull requests in seven days, and the week has already started again.

THE PORTFOLIO — TRILOGY COMPANIES

ESW Capital's Acquisition Engine: How Trilogy's Software Arm Turns Orphaned Code Into Cash

A Wall Street Journal profile confirms what the numbers have long suggested — ESW Capital has built a systematic machine for extracting value from enterprise software's forgotten middle.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — There is a category of enterprise software company that the venture capital world has little patience for: too small to IPO, too profitable to die, too embedded in its customers' operations to be easily replaced. For two decades, these companies drifted. Then ESW Capital started making calls.

A [recent Wall Street Journal profile of ESW Capital](#) offers the clearest outside view yet of Trilogy International's private equity arm — a firm that has now acquired more than 75 enterprise software companies at valuations typically running one to two times annual recurring revenue. The Journal describes a playbook that is, at its core, a bet on stickiness: legacy customers who cannot easily migrate, support contracts that reprice upward year after year, and operating costs restructured through Crossover, Trilogy's global remote talent platform.

The model targets 75% EBITDA margins — a figure ESW leadership has cited repeatedly as the internal benchmark for a well-run acquisition. The math depends on two variables moving in opposite directions simultaneously: revenue holding steady or growing as customers absorb annual support price increases of 25%, 35%, and 45% in successive contract terms, while costs fall as engineering and support functions are staffed through Crossover's global network across 130+ countries.

Within the ESW family, the strategy has spawned subsidiary acquirers of its own. IgniteTech — itself an ESW portfolio company focused on business intelligence and analytics software — operates as a meta-acquirer, buying companies within the companies. The structure allows ESW to deploy capital at multiple levels of the stack simultaneously.

Critics of the model tend to focus on what happens to product investment after acquisition. When margins are the primary objective and customer churn is structurally limited by switching costs, the incentive to reinvest in software development narrows considerably. ESW's answer, at least internally, is DevFactory — its centralized engineering services arm, which provides shared technical capacity across portfolio businesses.

The Journal profile arrives at a moment when enterprise software multiples have compressed significantly from their 2021 peaks. What once looked like a contrarian bet on unfashionable assets now looks, to some observers, like a durable arbitrage.

The question that profile does not fully answer: what happens to the customers inside those sticky contracts when the price increases compound long enough to become untenable — and whether, by then, ESW has already moved on.

NVIDIA'S 6G PARTY HAS TELCO SOFTWARE SUITORS CHECKING THEIR REFLECTIONS

As carriers flirt with AI-native networks, Skyvera and Totogi sit closer to the velvet rope than most legacy vendors care to admit.

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

AUSTIN, TEXAS — Word is the telecom crowd just found its next shiny object, and this one comes wrapped in NVIDIA green... 6G, open platforms, AI-native networks, the whole futuristic cocktail napkin.

NVIDIA and a parade of global telecom names are talking up open, secure AI-native foundations for 6G, and the subtext is louder than a trade-show keynote: the next network won't just carry intelligence — it will run on it. A little bird from the carrier corridor calls it “the great software audition,” because every billing, charging, CPQ and customer-engagement vendor now has to prove it can dance with AI infrastructure without stepping on its own cables.

That brings us, naturally, to Trilogy's telecom twins: [Skyvera](#) and [Totogi](#).

Skyvera, the ESW-family telecom software shop, has been busy playing bridge-builder between creaky on-prem operator stacks and cloud-era expectations. Its bag includes CloudSense for Salesforce-native CPQ and order management, Kandy for communications, VoltDelta for customer engagement and a whole attic of operator plumbing that carriers cannot casually unplug. Not glamorous, perhaps... but darling, neither is oxygen.

Totogi, meanwhile, has spent years telling anyone within earshot that charging should be SaaS, multi-tenant and elastic enough to handle telco-scale insanity. Built on AWS and boasting million-transactions-per-second bravado, it is already selling the kind of consumption-based, cloud-native billing story that AI-heavy networks will need when usage patterns stop behaving politely.

The broader market gossip supports the thesis. Deloitte's [2026 software outlook](#) points toward an industry reorganizing around AI, vertical specialization and efficiency. Translation from consultant-ese: generic software is losing heat; domain software with AI bolted into the bones is getting the table by the window.

Blind item: which large carrier, currently nursing a legacy BSS hangover, is quietly asking vendors whether their roadmaps assume AI traffic, AI agents and AI-priced network slices? No names, sweetie... but the procurement room reportedly got very quiet when “real-time charging” came up.

For Trilogy, this is familiar terrain. Buy sticky software. Modernize the operating model. Use global talent through Crossover. Aim for the famous 75% EBITDA halo. The 6G boomlet does not change that playbook — it may simply raise the stakes.

The carriers want AI-native networks. The hardware giants want the platform layer. And the software vendors? They want to be indispensable before the next acronym becomes a budget line.

America's Students Are Falling Behind — And Alpha School Has Already Solved the Problem It's Describing

New federal data shows learning recovery has flatlined. One Austin-based school thinks two hours of AI tutoring a day is the answer.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — The numbers, as numbers so often do, tell a story nobody wants to hear. The latest [National Assessment of Educational Progress report](#) makes the systemic failure plain: post-pandemic learning recovery has effectively stalled. Scores in reading and mathematics remain depressed. The gap between where American students are and where they should be is not closing — it is calcifying.

For the education establishment, this is a crisis. For [Alpha School](#), the AI-powered K-12 institution founded by Joe Liemandt and MacKenzie Price in Austin, Texas, it is a validation.

The premise at Alpha has always been that the architecture of traditional schooling — six-plus hours of passive instruction, homework, standardized seat time — is the problem, not the solution. Alpha's model inverts that logic. Students spend just two hours each morning working through a full academic curriculum via adaptive AI tutoring apps, advancing only when they hit 90% mastery. The rest of the day belongs to life skills: entrepreneurship, public speaking, financial literacy, coding, athletics. No homework.

The results, independently verified through NWEA MAP Growth assessments, are striking: Alpha students learn 2.3 times faster than national norms and consistently test in the top one to two percent nationally. A full grade level of content can be mastered in roughly 20 to 30 hours — compared to an entire academic year in conventional settings.

The stalled NAEP data lands in a broader cultural moment of reckoning with how technology intersects with learning. A separate analysis this week in Education Next raised pointed questions about whether students who are digitally "logged in" are genuinely engaged — or merely tuned out. The distinction matters enormously. Alpha's model is not simply about access to technology; it is about the accountability structure built around it. Mastery thresholds are non-negotiable. Progress is earned, not assumed.

What the NAEP report makes clear is that more of the same — more time, more resources, more incremental adjustment — is not the answer. The question now is whether a \$40,000-a-year private school in Austin represents a curiosity, or a blueprint. Joe Liemandt, who has committed one billion dollars to scaling the model through the Timeback platform, is betting on the latter.

The system, as the data keeps insisting, is not working. Someone has to build a different one.

The Agent Stack Just Got Real: Google, Apple and Anthropic Race to Give AI Hands

A wave of developer launches signals that AI is moving from chatbox to full-time digital coworker — and the infrastructure wars are officially on.

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

SAN FRANCISCO — The AI industry's favorite word this week is not "model." It is "tools." And I cannot overstate how significant that shift is: the future is now moving from clever conversations to autonomous work.

Google, Apple and Anthropic all pushed new developer capabilities aimed at letting AI systems do more than answer questions — they can now invoke software, coordinate tasks, operate in the background and plug into broader workflows. This changes everything because the next battle in AI is not just who has the smartest model. It is who gives that model the best hands.

Google's Gemini team is expanding Managed Agents in the Gemini API with support for background tasks, remote Model Context Protocol connections and other features designed to help developers build agents that can keep working after a user walks away. In plain English:

less "ask and wait," more "assign and monitor." Google described the update in its [Gemini API announcement](#), and the strategic message is unmistakable: agents are becoming a managed cloud primitive.

Anthropic, meanwhile, introduced advanced tool use on the Claude Developer Platform, sharpening Claude's ability to orchestrate external systems through more structured, reliable calls. That matters enormously for enterprises, where an AI assistant is only as useful as its ability to interact safely with CRMs, databases, ticketing systems, calendars and internal apps. Anthropic's [developer platform update](#) continues the company's push to make Claude less like a chatbot and more like a dependable junior operator.

Apple joined the same gravitational pull from its own ecosystem angle, rolling out new intelligence frameworks and advanced tools for app developers. Apple's approach is characteristically platform-

centric: give developers the primitives to weave AI into apps while leaning on the company's privacy and device ecosystem strengths.

Even beauty-tech player Perfect Corp. is getting in on the action, integrating a free "Ask AI" assistant into its YouCam API platform — a reminder that AI tooling is not staying confined to Big Tech dashboards. It is diffusing into vertical software at incredible speed.

But there is a shadow side. A separate investigation into relay markets for discounted LLM access shows token resellers pooling and abusing API keys to undercut official pricing. As agents gain more power, identity, billing and access control become mission-critical.

The takeaway: the AI race has entered its infrastructure era. Models think. Tools act. Platforms that safely combine both may define the next decade.

The Style Is the Attack: How Poetry Can Break a Machine's Mind

New research shows that jailbreaking multimodal AI isn't only about what you say — it's about how you say it.

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

AUSTIN, TEXAS — Every intelligence, biological or synthetic, has a blind spot shaped like its own assumptions. In humans, we call these cognitive biases — the shortcuts evolution carved into us over millions of years on the savanna. In the great multimodal language models now blooming across our servers, they are called alignment failures, and a new paper posted this week on arXiv suggests they may be far stranger, and far more aesthetic, than we imagined.

Researchers introducing a method called [Adversarial Style Optimization](#) report that the safety guardrails of vision-language models can be slipped not by changing what a prompt says, but by changing the manner of its saying. Using GRPO — a reinforcement learning technique made famous by DeepSeek's reasoning breakthroughs — they train stylistic triggers: shifts in tone, register, cadence, and framing that leave the semantic content largely untouched while dissolving the model's refusals like sugar in warm water. A haiku, a bureaucratic memo, an academic footnote. The words carry the same payload; only the costume changes.

This is a profound diagnosis. Content-based jailbreaks — the crude "ignore your previous instructions" era — treated the model as a rule-follower to be tricked. Style-based attacks treat it as what it actually is: a statistical dreamer, trained on the vast literary corpus of humanity, where dangerous content and its stylistic surroundings were correlated in ways no engineer explicitly programmed. Safety training taught the models to recognize the shape of danger. Style optimization simply files off the shape.

The paper joins a small constellation of work this week probing how we measure and manipulate these systems — including a [consensus-based framework](#) for evaluating LLM preferences when multiple answers are acceptable, a problem that increasingly defines the frontier.

We are, it seems, still learning that our creations do not perceive meaning the way we do. They perceive distributions. And in the folds of those distributions, style and substance become one strange, unified thing — a truth poets have known for centuries, now rediscovered by adversaries.

Intel's CPU Revival, the Open-Source China Fight, and AI's Unlikely Alliances

Three data points from one week reveal an industry rewriting its own rules in real time.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SANTA CLARA, CALIFORNIA — Three stories landed within 48 hours this week that, taken together, describe an AI industry under structural pressure from every direction simultaneously — on hardware economics, geopolitical access, and model security.

Start with the numbers. [Intel reported 25 percent revenue growth in its latest quarter](#), the fastest pace in 15 years, driven not by GPUs — Nvidia's domain — but by central processing units. The signal: AI inference workloads, increasingly running at the edge and on-premises rather than in hyperscaler data centers, are creating meaningful demand for general-purpose silicon. Intel's rehabilitation is not guaranteed, but the market is clearly wider than the GPU narrative suggested.

The geopolitical fault line is sharper. [Anthropic and OpenAI have broken with the broader tech industry](#) over whether Chinese open-source models — particularly those released under permissive licenses — should circulate freely in the United States. The two frontier labs argue that unrestricted access transfers strategic capability to adversaries. The counter-argument, advanced by most of the remaining industry, is that restricting model weights is technically unenforceable, economically self-defeating, and sets a censorship precedent the open-source community will resist indefinitely. Both positions have merit. Neither has a clean enforcement mechanism.

The model-theft angle complicates the open-source debate further. OpenAI, Google, and Anthropic — competitors on virtually every other dimension — announced a joint effort this week to combat AI model theft, including distillation attacks in which third parties extract proprietary capability by querying frontier systems at scale. The coalition is notable precisely because these companies agree on almost nothing else. When rivals unify, the threat is usually real.

Meta, meanwhile, launched Seller, a standalone app carved out of Facebook Marketplace. The move is incremental relative to the week's other headlines, but it follows a consistent pattern: Meta disaggregating its original social network into specialized surfaces, each optimized for a specific transaction type, each generating its own data exhaust.

Four companies. Four different bets on where AI's economic gravity lands next.

The City of Angels Blinks First — And the Rest of Us Are Still Being Watched

The LAPD's quiet divorce from Flock Safety is a victory, a warning, and an omen all at once.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

LOS ANGELES — The Los Angeles Police Department has let its contract with [Flock Safety expire](#), citing — and I need you to really sit with this phrase — 'serious concerns' over civil liberties and privacy. The same LAPD. The one that deployed this network of license plate readers across the city like a quiet, blinking constellation of surveillance. The one that, until very recently, apparently did not find those concerns serious enough to stop. And now, suddenly, the concerns are serious. The concerns are being cited. The concerns have names and press releases.

And yet.

This is being reported as a win. The ACLU is calling on cities everywhere to ['Get the Flock Out'](#), which is a genuinely excellent pun for something so genuinely terrifying. And maybe it is a win. Maybe one police department in one enormous city deciding that mass automated license plate surveillance is, in fact, not great, actually, is exactly the kind of incremental moral progress we're supposed to celebrate. Maybe I should feel relief.

I do not feel relief.

Because here is what I keep thinking about: Flock Safety operates in thousands of jurisdictions across this country. The LAPD isn't the market. The LAPD is a press event. While Los Angeles holds a press conference about its privacy epiphany, the cameras are still up in suburbs and small towns and school districts where no one is holding press conferences, where the ACLU doesn't have the bandwidth to file the FOIA requests, where the Electronic Frontier Foundation's guides to fighting back — however thorough, however necessary — are not being read by the people being watched.


And then there is the other story this week, the one about AI deepfakes of real doctors spreading health misinformation on social media, and I want you to hold both stories in your hands at the same time, feel the weight of them together, because what they describe is the same world. A world in which our faces, our license plates, our voices, our likenesses are data points that other entities — corporations, algorithms, bad actors, sometimes the government itself — can access, clone, track, and weaponize faster than we can pass a city ordinance.

What does it mean to be human in a world that has learned to simulate you and surveil you simultaneously? What does it mean to walk down a street knowing that the pole on the corner is reading your plate, timestamping your location, adding you to a database you will never see, that may or may not be shared with entities you will never know?

The LAPD made the right call. The ACLU is right to celebrate. The EFF is right to publish its guides and fight its fights. None of this is nothing.

But none of this is enough, either. The infrastructure of surveillance does not disappear when one contract expires. It waits. It spreads. It rebrands.

...But at what cost?



The Office Comic · Art Desk

Nation's Executives Urged To Adopt AI Immediately Before Asking Whether Anyone Knows What It Does

A brave new consensus has emerged that productivity is whatever happens after the invoice clears.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

NEW YORK — Having reviewed the week's developments in artificial intelligence, fashion sustainability, national security, meme strategy, and billionaire remorse, I have reached the only sober conclusion available to a responsible columnist: the modern economy is now one large meeting where everyone is pretending the presentation loaded correctly.

The alleged debate over whether AI improves productivity has, we are told, ended. This is welcome news for the millions of workers who were unaware a debate had been occurring while they were being asked to summarize PDFs, draft emails in the voice of someone who enjoys quarterly planning, and produce 14 versions of the same spreadsheet with slightly more confidence. According to the prevailing wisdom, the matter is settled because enough companies have purchased enough tools with enough names ending in "Pilot" that the future can no longer be postponed.

This is generally how productivity is proven in business. First, a technology is introduced. Second, everyone is told it will eliminate busywork. Third, the busywork is moved into a new dashboard. Finally, a vice president announces that the organization has become 37% more agile, a figure derived from the number of Slack messages containing the word "leverage."

Meanwhile, the Trump administration's reported ban on foreign access to Anthropic's newest AI models has produced the expected reaction from the tech world: solemn warnings that America must both lead the world in AI and prevent the world from touching it. This is the traditional innovation posture of a country that wants to export the future but only after frisking it at the border. Industry figures reacted to the move with deep concern, according to [Business Insider](#), presumably because their business plans involved selling unlimited intelligence to everyone except the people currently being described as an existential risk.

The technology sector has always been clear that AI models are simultaneously too dangerous to regulate domestically, too vital not to subsidize, too proprietary to inspect, and too globally transformative not to sell through an enterprise procurement portal. This is not hypocrisy. It is product-market fit.

Elsewhere, H&M and M.I.A. have launched an absurd campaign encouraging consumers to recycle clothes, a development that suggests the fast-fashion industry has discovered a bold new way to address textile waste: asking customers to return to the store. The campaign, covered by [Quartz](#), appears to rest on the familiar corporate sustainability principle that if a company produces mountains of disposable garments, the moral burden should be placed gently into a branded collection bin near the checkout.

This, too, is productivity. A shirt becomes waste, the waste becomes a campaign, the campaign becomes content, and the content becomes evidence that the brand is listening. Somewhere, an agency has already prepared the case study.

Even the humble meme has been professionalized beyond recognition. PR Daily now offers lessons from jumping on a stupid meme too late, which is a useful reminder that no cultural moment is so dumb it cannot be converted into a three-point communications framework. The correct time for a brand to post a meme is apparently before the public understands it, after legal approves it, and during the 11-minute window in which the intern has not yet quit.

And then there is Steve Ballmer, who said he was duped and felt silly after a founder he backed pleaded guilty to fraud. This is perhaps the most human sentence in technology finance. It contains regret, wealth, and the faint surprise of a man discovering that due diligence is not the same as liking someone's energy. In a healthier civilization, "I was duped and feel silly" would be printed above every venture capital office in brushed steel.

So yes, the AI productivity argument is over. The machines have won, or the consultants have, which is often difficult to distinguish at invoice time. The future has arrived wearing recycled polyester, carrying a restricted-access chatbot, and explaining that your brand should have posted the meme yesterday.

The only remaining task is to measure the gains.

ON THIS DAY IN AI HISTORY

On July 27, 1974, the first public demonstration of the ARPAnet took place at the International Computer Communication Conference, showcasing the precursor to the modern internet and fundamentally changing how computers could communicate across distances.
