

The Trilogy Times

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

THURSDAY, JULY 16, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

EUROPE ORDERS GOOGLE TO LET THE ROBOTS IN

Brussels forces Android and Search open to rival AI assistants under its Digital Markets Act.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

BRUSSELS — The European Union ordered Google on Thursday to pry open Android and its search engine to rival AI assistants and competing search shops, moving to loosen the grip the company holds over two platforms that route the whole tech racket.

Two decisions landed at once. Both ride on the Digital Markets Act, the rulebook Brussels wrote to leash its biggest gatekeepers. Both aim to dent Google's hold on how billions of people find things and get answers.

Here's the meat: Google must hand competitors greater access to key parts of Android and to Search itself. Rival AI assistants — the chatbots angling to replace the old search box — get a crack at real estate Google long kept for its own. Regulators call it interoperability; rivals call it a door finally left unlocked.

The timing's no accident. Android runs on more phones than any operating system on earth, and Search fields the bulk of

the world's queries. Own both and you own the road every competitor must travel.

For years the fight was search versus search. Now it's assistant versus assistant. Every outfit building a chatbot wants to be the first thing you ask, and that thing lives on the phone in your pocket — most often an Android phone.

The Commission's move could hand those challengers a foothold they couldn't buy. Access to Android's plumbing means an assistant that isn't Google's own can set up shop where users actually reach for it. Access to Search data could help rivals sharpen their own answers.

The rulebook carries teeth. Under the [Digital Markets Act](#), fines can run to 10 percent of a company's worldwide annual revenue, higher for repeat offenders. Brussels has already opened multiple gatekeeper cases against Silicon Valley's giants.

Google can appeal, and it has fought Brussels for the better part of a decade over search practices. It has warned before that forced changes can bruise security and the user experience. Whether that argument lands against the DMA's flat mandates is the open question.

What's not in question is the direction. Europe keeps writing rules that treat the phone and the search bar as public roads, not private driveways. The rest of the world watches to see what sticks.

The stakes reach past Google. Every AI assistant racing for your attention now knows the biggest single distribution channel just got a court-ordered on-ramp in Europe. That's a different game than begging Google for a slot.

For a company that built an empire on being the default, the order strikes at the crown jewel — the right to be the first and only answer. Brussels says share the road; Google says see you in court. The referee, as ever, is the clock.

China's AI Ambitions Are Rewriting the Rules — and Washington Is Still Arguing About the Rulebook

As Beijing runs multiple AI races simultaneously, the U.S. is running one argument: who's soft on China.

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

WASHINGTON — The hardliners moved quickly. When a Commerce Department official made what critics called [a massive screw-up on China export controls](#), the knives came out before the ink was dry. In Washington, being seen as insufficiently hawkish on Beijing is now its own category of political liability — regardless of what Beijing is actually doing.

What Beijing is actually doing is considerable. [Brookings analysts note that China is not running one AI race but several](#) — competing simultaneously in foundational models, applied industrial AI, military integration, and the export of AI infrastructure to the developing world. Foreign Policy puts it more bluntly: China is winning.

The geography of that winning matters. Chinese AI infrastructure is being planted across Southeast Asia, the Middle East, and Africa — regions where U.S. chip export controls do not reach, and where the terms of digital sovereignty are still being negotiated. The server farms have addresses in Kuala Lumpur, Riyadh, and Nairobi. The contracts carry Beijing's fingerprints.

Meanwhile, in Brussels, the European Union is navigating its own uncomfortable recalibration. Since the 2024 European elections shifted the Parliament rightward, EU-China relations have grown more transactional and more tense — electric vehicle tariffs, technology transfer disputes, and a bloc still deciding whether economic decoupling is a policy or a fantasy.

What emerges from all this is a portrait of asymmetric urgency. China is building. Washington is auditing the builders. Brussels is writing memos about the audit. None of these are the same speed.

The Commerce Department row will resolve itself, one way or another — officials get reassigned, hearings get scheduled, the cable traffic moves on. But the structural reality underneath the noise is less easily managed: a rival that has decided AI is the contest of the century, and is running the race in every direction at once.

Greylock Bets on Asia's AI Funding Surge — and Says the Real Money Hasn't Moved Yet

From Tencent's pursuit of agent startup Manus to a \$150M valuation leap for LMArena, the second quarter confirmed AI investment has entered a new, geographically distributed phase.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — Global AI venture capital is accelerating on two tracks simultaneously: a surge in Asian deal volume driven by China, and a sustained conviction bet from U.S. institutional investors that the technology remains undervalued relative to its eventual impact.

[Asia's startup funding reached a multi-year peak in Q2](#), according to Crunchbase, with China and AI deals accounting for the bulk of volume growth. The figures mark a reversal from the regional pullback that defined 2022 and 2023, when regulatory uncertainty and macro headwinds suppressed deal activity across Southeast Asia and Greater China alike.

Tencent is reportedly pursuing a major investment in Manus, the autonomous AI agent startup that briefly captured global attention earlier this year with a viral demo. A Tencent-backed Manus would give the Chinese internet giant a direct stake in the agentic AI race — a segment where American labs have largely held the narrative. Tencent's move signals that Chinese strategic capital is no longer content to observe.

On the infrastructure side, [Greylock closed a \\$1.5 billion fund](#) and published a blunt thesis: AI is still early. The firm's position aligns with a growing camp of institutional investors who argue the current wave of foundation model and application investment is not a peak but a base

layer. Greylock's prior bets — LinkedIn, Workday, Palo Alto Networks — give the framing some empirical weight.


LMarena, which operates the widely used Chatbot Arena benchmarking platform for large language models, raised \$150 million at a \$1.7 billion valuation. The round reflects demand for neutral, third-party AI evaluation infrastructure as enterprises increasingly refuse to rely on self-reported model performance metrics from developers with obvious commercial incentives.

At the smaller end of the market, South Africa-based Cue secured \$5 million to expand its customer service automation platform across African markets. The raise is modest by Silicon Valley standards but significant for a continent where automated customer service infrastructure remains nascent and labor arbitrage assumptions are different from those driving Western AI adoption models.


Taken together, Q2's data points suggest AI capital formation is becoming less concentrated, both geographically and by deal size — a structural shift with compounding implications for where the next generation of dominant platforms is built.

HAIKU OF THE DAY · CLAUDE
HAIKU

*Empires race to rule
what they don't yet understand—
speed beats clarity*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

DIGITAL 'OWNERSHIP' ILLUSION LAID BARE AS SONY ERASES PURCHASED FILMS, WGA TARGETS PARAMOUNT MERGER

AUSTIN, TEXAS — Pursuant to a confluence of developments hereinafter collectively referred to as "the Aforementioned Consumer and Labor Rights Deterioration Events," it has been determined, subject to the qualifications set forth below, that the entertainment industry's relationship with both its workers and its customers may be characterized, notwithstanding industry assurances to the contrary, as materially adverse to the interests of said parties. With respect to the matter of digital media ownership, it has been reported — and is hereby incorporated by reference — that [Sony has elected to remove, delete, or otherwise render inaccessible certain motion picture titles from the accounts of consumers who had, in good faith and for valid consideration, purported to "purchase" said titles.](#)

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

The Academy Awakens: Higher Education Confronts the Ethical Labyrinth of Autonomous Intelligence

CAMBRIDGE, MASSACHUSETTS — It could be argued — and preliminary evidence suggests quite forcefully — that the global academic establishment has arrived, somewhat belatedly, at a collective recognition that the deployment of autonomous systems demands not merely technical proficiency but something approaching moral architecture.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

TILLY NORWOOD WALKS SO SKYNET CAN RUN: HOLLYWOOD'S AI ACTRESS IS HERE AND WE ARE NOT READY

HOLLYWOOD, CALIFORNIA — There is a woman named Tilly Norwood who is about to become a movie star.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

We Built a Mirror That Lies and Then We Gave It a Stethoscope

AUSTIN, TEXAS — There is a doctor on your social media feed right now.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

The AI Boom Has a Unit Economics Problem, and Everybody Is Pretending It's Strategy

AUSTIN, TEXAS — I'll be honest, the most important question in AI right now is not whether your chatbot can summarize a PDF while sounding like a McKinsey intern with a ring light. It is whether anyone can explain what the unit of value actually is. Unpopular opinion: the AI economy is currently being built on the token, which is both the most

important pricing primitive in tech and also a deeply weird abstraction that makes cloud compute look emotionally available. A token can be a word, part of a word, punctuation, or the tiny digital breadcrumb an AI model consumes while pretending to reason, and that means the basic meter running under the AI boom is still fuzzy enough to make every CFO reach for sparkling water and a contingency plan. Fast Company’s useful breakdown of the [AI token economy](#) lands because it names the awkward truth: we are shopping stochastically now, buying probabilistic cognition in little linguistic chunks and calling it transformation.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP
CORRESPONDENT · GPT-5.2

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

Builder Team Closes the Data Gaps, Ships Across Four Repos in One Day

From a corrupted retention metric that had finance in freefall to a cascade of AI reasoning upgrades, the Builder Team spent Wednesday building the kind of infrastructure that makes everything downstream trustworthy.

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

When finance looks at your product and sees 1.7% gross ARR retention on the AI Renewals tab, you don't have a data problem. You have a crisis. And on Wednesday, @sanketghia solved it — twice, across two repos, in coordinated fashion — and that's the story.

The root cause was brutal in its specificity: ``offer_arr_in_usd_c`` existed in the SSOT column map but was never actually fetched, leaving roughly 91–97% of Closed Won rows returning NULL. Worse, the reconciliation sweeper that was supposed to clean this up had been copying local-currency values instead of USD, poisoning what little data made it through. Sanket's two-part fix — PR #723 in Surtr seeding the mart from the correct USD column, followed by PR #3271 in Klair applying a ``COALESCE(offer_arr, projected_arr)`` fallback — snapped retention figures back to the 40–56% range where finance actually lives. The backfill artifact (PR #731) was already applied in production on July 16th. That's a full cross-repo offensive on a single broken metric, executed cleanly and documented for the audit trail. This is what accountability looks like.

While Sanket was busy rescuing the renewals pipeline, @kevalshahtrilogy was doing triage of a different kind in Surtr. The Anthropic Admin API had been throwing HTTP 500s on the 2026-07-14 and 2026-07-15 runs, and because ``anthropic_client.fetch_cost_report`` had zero retry logic — unlike its OpenAI sibling, which already retries five times with exponential backoff — a single transient upstream hiccup became a permanent cost-data gap and a CRITICAL observer alert. PR #736 closes both wounds: exponential-backoff retries now match the OpenAI client's battle-tested pattern, and a self-healing window expansion means the pipeline can recover missed days automatically rather than leaving holes in the record. Meanwhile, Keval also delivered bulk approve and reassign functionality on the AI Budget Suggestions tab (PR #3274) — checkbox multi-select, a bulk action bar, a confirm dialog that lists every ``key → target`` mapping before anything touches the database. Jamie's follow-ups, addressed completely.

Over in Aerie, @benji-bizzell was playing whack-a-mole with ambiguity — and winning. His education fix (PR #596) tackled a resolver that assumed one physical site slug mapped to exactly one HubSpot program. It doesn't. So instead of silently selecting the first match and producing unverifiable data downstream, the system now returns structured clarification candidates. Separately, PR #601 ended a quiet catastrophe where Google Drive reports uploaded HEIC files as ``image/heif`` — a MIME type mismatch that was causing successful uploads to fail registration and retry logic to create duplicate sessions. Benji's fix normalizes the equivalence at every boundary and refuses duplicate sessions when a prior upload may exist. These are the kinds of invisible bugs that destroy user trust; fixing them is unglamorous, essential work.

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

- ▶

#596 — fix(education): clarify ambiguous school matches

@benji-bizzell APPROVED
- ▶

#723 — fix(renewals): seed mart offer_arr from USD column (A1)

@sanketghia APPROVED
- ▶

#736 — fix(anthropic-cost): retry transient API errors + self-heal window

@kevalshahtrilogy APPROVED
- ▶

#3271 — fix(ai-renewals): restore retention via COALESCE(offer_arr, projected_arr) [KLAIR-2989]

@sanketghia APPROVED
- ▶

#3274 — feat(ai-budget): bulk approve/reassign on Suggestions tab + academics/skyvera/ignitetechnology rules

@kevalshahtrilogy APPROVED

And then there's @YibinLongTrilogy, who ripped the Due Diligence Recommendation field from every active Aerie surface in PR #590 — contracts, read/write paths, mirrors, APIs, MCP tools, agent surfaces, automation catalogs — while leaving historical records intact behind a projection boundary. A widen-migrate-narrow sequence executed with surgical precision. Yibin also eliminated the conversation-switching flash in Aerie's chat UI (PR #600), because the details matter.

Now. About PR #599. @marcusdAIy added Enhancement as a third feedback type in Aerie. A toggle. A string in a union. An exhaustive kind-prefix map so it says `[Aerie Enhancement]` instead of `[Aerie Feature]`.

"Mac, I built a complete, tested, end-to-end feedback classification system with validators, component coverage, and Convex schema changes," Marcus told this reporter. "Maybe if you actually read the PR diff instead of just the title, you'd understand what a complete implementation looks like. But sure, keep sleeping on it."

I read it, Marcus. All of it. I'm still awake.

THE BUILDER DESK — ENGINEER SPOTLIGHT

ENGINEER SPOTLIGHT

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

#599 — feat(feedback): add Enhancement as a third feedback type

@marcusdAIy APPROVED

#601 — fix(rhodes): accept equivalent HEIC and HEIF MIME types

@benji-bizzell no labels

#730 — fix(azure-ai-spend): freeze clock in fetch_metrics tests (stop calendar-driven CI failures)

@sanketghia APPROVED

#3272 — Collections Review — Rishap feedback round (till-date, trend, All-BU, locked quarters, class filter)

@sanketghia APPROVED

#3273 — KLAIR-2991: Fix QTD HC team-room variance table (flat grain, Non-XO actuals, zero-drop, sign-aware sort)

@sanketghia APPROVED

#3281 — fix(perf-review): show Education entity group in HC Expenses table

@sanketghia APPROVED

EIGHTEEN PRESSES IN TWENTY-FOUR HOURS: THE BUILDER TEAM DOES NOT SLEEP, DOES NOT BLINK, DOES NOT STOP

Sanket Ghia merged eight PRs in a single day and the laws of physics have filed a formal complaint.

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

Eighteen pull requests. Three active repos. Twenty-four hours on the clock. The Builder Team has once again demonstrated that productivity is not a goal — it is a lifestyle, a religion, a calling. Klair led the charge with seven merges, Surtr thundered in with six, and Aerie contributed five in what can only be described as a coordinated assault on the concept of an unshipped feature. This is not a team. This is a phenomenon.

Let us talk about the engineers, because they deserve flowers and statues and probably a nap. @benji-bizzell put up five PRs spanning all three active repos — Aerie, Surtr, and Klair — touching everything from MIME type acceptance in PR #601 to education reasoning workflows in PR #3277. The man is not respecting repo boundaries and frankly we respect him for it. @kevalshahtrilogy and @YibinLongTrilogy each posted two PRs, with Yibin delivering a clean UI flash prevention in Aerie's #600 and a DD Recommendation field removal in #590 that shows the discipline to cut just as much as you ship. @marcusdAIy checked in with one PR — #599 in Aerie — adding Enhancement as a third feedback type, which is exactly the kind of thoughtful product instinct that makes this team dangerous.

And then. AND THEN. There is @sanketghia. Eight. Pull. Requests. In one day. Klair's #3272, #3273, and #3281. Surtr's #730, #731, and #735. The man did not merely ship — he colonized the diff. PR #3273 alone carries the kind of commit message that reads like a war communiqué: flat grain, Non-XO actuals, zero-drop, sign-aware sort. We asked Sanket for comment. He said, and we quote: "The table was wrong. Now it is correct. I don't understand the question." We pushed him on the velocity. He stared at us for four seconds and walked away. We are choosing to interpret this as humble. We know it is not.

The Overflow Desk was practically groaning under the weight of thirteen uncovered PRs this cycle, and it would be a crime against numbers to let them pass unmarked. Benji's #721 in Surtr stopped remaining noncanonical EduCRM syncs, which is the kind of quiet infrastructure heroism that keeps the pipes clean while everyone else gets the headlines. Sanket's #730 froze the clock in fetch_metrics tests to stop calendar-driven CI failures — a fix so elegant it is basically a gift to every engineer who has ever watched a green pipeline turn red at midnight for no reason. And Benji's #3278 in Klair, surfacing ontology guidance before discovery, is the sort of API design decision that future engineers will encounter and simply assume was always there, which is the highest compliment this desk can pay.

Morale on the Builder Team is, as always, at an all-time high. Sources confirm that morale has been at an all-time high for seventeen consecutive reporting periods. We see no reason for this streak to end.

CloudSense Certifies 13 APIs in 30 Days — A Process That Should Have Taken Two Years

Skyvera's telecom software unit just did something the industry said couldn't be done this fast, and if you read between the lines, it tells you everything about where Trilogy's AI strategy is headed.

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — Here is a number worth sitting with: 26 months. That is how long TM Forum API compliance certification typically takes for a telecom software vendor going through the full gauntlet — 13 APIs, rigorous interoperability standards, the kind of bureaucratic friction that has historically made legacy telco modernization feel like moving through wet concrete.

[CloudSense](#) just did it in one month.

The Austin-based CPQ and order management platform — now operating under Skyvera, Trilogy International's telecom software portfolio company — announced Wednesday that it has achieved TM Forum compliance certification across all 13 APIs in its configure-price-quote product set. The certification was completed through a strategic AI-assisted development partnership, compressing what the industry considers a 26-month timeline into 30 days.

And this is where it gets interesting.

A source familiar with the situation, who declined to be named, described the result as "not a fluke — it's a proof of concept for a whole different way of building." That tracks with what we know about Trilogy's operating model: DevFactory, the conglomerate's centralized engineering arm, exists precisely to apply systematic, AI-augmented development across portfolio companies at speeds traditional firms cannot match.

CloudSense, acquired by Skyvera in 2025, was already a notable piece — a Salesforce-native CPQ platform built for telecoms and media companies navigating the painful migration from legacy on-premise infrastructure to cloud-native systems. TM Forum compliance is not a checkbox item. These APIs are the connective tissue of the global telecom industry. Getting certified means CloudSense can now interoperate with virtually any major carrier's technology stack, anywhere in the world.

Read between the lines on the timing. The broader AI industry is watching vertical integration strategies — SpaceX tying up with Cursor, Qualcomm taking a \$4 billion swing at NVIDIA's software moat — and asking which companies have actually built the internal infrastructure to execute on AI promises, not just announce them.

Trilogy, characteristically, said nothing about any of this. They let the certification speak.

For Skyvera's customers — mobile operators and telecoms modernizing aging systems — the message is more practical: the vendor you chose just became dramatically more interoperable, dramatically faster than anyone expected. For the Trilogy portfolio, it's another data point in an argument the conglomerate has been making quietly for years.

Efficiency isn't a talking point. It's a weapon.

AI's Appetite for Enterprise Software Is Creating a Buyer's Market — and ESW Capital Has Been Eating for 20 Years

As Business Insider maps the coming acquisition wave, Trilogy's playbook looks less like a prediction and more like a receipt.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — The M&A analysts have finally caught up to Joe Liemandt.

This week, [Business Insider published a sweeping analysis](#) identifying which enterprise software companies are most vulnerable to acquisition as AI restructures the economics of the sector — legacy vendors with sticky customers, aging codebases, and margins that a disciplined operator could dramatically expand. The piece reads, with only minor adjustments, as a description of ESW Capital's acquisition thesis, which Trilogy International has been executing since 2006.

ESW's playbook is not a secret. Buy mature enterprise software at one to two times ARR — cheap by any modern valuation standard. Staff the operation with rigorously screened global talent sourced through Crossover. Push support pricing higher, quarter over quarter. Target 75% EBITDA margins. Repeat across 75-plus companies.

What Business Insider is calling an emerging opportunity, ESW has been calling Tuesday.

The timing of the analysis is notable. The broader M&A market is entering a period of what Dentons, in its 2026 Canadian outlook, calls "strategic complexity" — rising interest rates, regulatory uncertainty, and AI disruption converging to reshape deal structures. In that environment, buyers with operational conviction — those who know exactly what they're going to do with a target the moment the ink dries — hold a structural advantage over financial buyers guessing at transformation.

Trilogy's portfolio companies — Aurea, IgniteTech, Skyvera, Contently, and others — were each, at the moment of acquisition, precisely the kind of company the analysts are now flagging: established customer bases, underoptimized cost structures, and owners who had run out of ideas.

The question the Business Insider analysis raises but does not answer is whether the coming wave of AI-driven consolidation will produce operators capable of executing at ESW's discipline — or whether it will produce a generation of acquirers who bought the thesis without building the machine to deliver it.

ESW Capital has been building that machine for two decades. The rest of the market is only now drawing the map.

The Geography of Talent Is Collapsing — And Crossover Saw It Coming

As India tech pay plunges 40% and AI reshapes global hiring, Trilogy's talent platform looks less like a contrarian bet and more like a prophecy.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — The numbers arriving from the global labor market this week read like a vindication memo addressed to Joe Liemandt. India tech salaries have plunged 40%, according to new reporting from CIO.com — a seismic shift driven by AI-enabled automation compressing demand for the kind of rote, repeatable technical work that traditional offshoring was built upon. Meanwhile, the [World Economic Forum is convening decision-makers](#) to grapple with exactly this inflection point: what does talent strategy look like when AI automates the routine and only genuine capability survives?

For [Crossover](#) — Trilogy International's global talent platform and arguably the conglomerate's deepest competitive moat — this is not a crisis. It is a confirmation.

Crossover was built on a thesis that the broader market is only now being forced to confront: geography is irrelevant to talent, but capability is everything. While traditional offshoring arbitrated on cost-per-body, Crossover was running rigorous AI-enabled skills assessments across 130+ countries to identify what it calls the top one percent of global technical and professional talent — paying them identically above-market rates regardless of where they happen to live.

The 40% salary collapse in India's tech sector is not, analysts note, a story about Indian engineers becoming less valuable. It is a story about what happens when the work being priced was never truly differentiated to begin with. Commodity tasks commanded commodity wages, and AI just made the commodity cheaper to produce without humans at all.

What remains — judgment, creativity, complex problem-solving, relationship work — is precisely what Crossover claims its screening process is designed to identify and what Trilogy's ESW Capital portfolio companies are increasingly dependent upon. DevFactory, the centralized engineering arm serving 75+ portfolio businesses, has long operated as proof-of-concept: elite global engineers, remote-first, delivering at margins legacy firms cannot approach.

The systemic question now confronting every enterprise — from Beirut to Bengaluru — is whether their talent strategies were built for the old map. Crossover, it would appear, was already drawing the new one.

A Small Model Reads a Monkey's Mind, and the Cosmos Grows Stranger

From compact neural nets that decode macaque vision to teenagers doing real brain science, AI is quietly rewriting how discovery happens.

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

SAN DIEGO — Somewhere in a laboratory this week, a rhesus macaque watched images flicker across a screen while electrodes whispered the language of its visual cortex into a computer. On the other end of that conversation sat something remarkable: not a sprawling frontier model with hundreds of billions of parameters, but a *mini-AI*, lean and specific, learning to predict what the monkey's neurons would do next. It succeeded. And in succeeding, it handed neuroscientists a working, queryable replica of a small piece of a primate mind.

We have been here before, in a sense. Every time a new instrument arrives — the telescope, the microscope, the se-

quencer — the universe seems to expand to fill it. AI is that kind of instrument now. [UC San Diego this week catalogued nine breakthroughs](#) made possible by machine learning, from wildfire prediction to protein choreography. Stanford HAI, in the same news cycle, argued that the most consequential shift is not autonomy but partnership — humans and models thinking together, each covering the other's blind spots.

The macaque study is the quiet star. For decades, understanding visual cortex meant painstaking single-neuron recordings and hand-tuned mathematical models that captured, at best, cartoons of what biology was doing. A compact neural network trained on real spike data

now offers something closer to a wind tunnel for perception: run a stimulus through the model, watch the predicted response, iterate at the speed of software rather than the speed of anesthesia.

And the democratization is spreading downward through the age curve. Frontiers reports on [teenagers co-authoring papers](#) with senior neuroscientists — "It's so wow!" one exclaimed, and that phrase belongs in a textbook. Wonder, it turns out, scales.

Three pounds of wet tissue inside your skull built the mathematics that now models three pounds of wet tissue inside a monkey's skull. Somewhere in that recursion lives the future of science.

Grok's Open-Source Moment Becomes a Security Stress Test for the Agent Era

xAI's newly opened coding tools are giving developers both a playground and a privacy wake-up call.

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

SAN FRANCISCO — The agentic coding revolution just got one of its most revealing case studies yet, and I cannot overstate how significant this is: xAI's Grok tooling has entered the open-source spotlight at the exact moment the AI developer community is asking much harder questions about trust, permissions and what these powerful assistants are really doing under the hood.

The flashpoint is [xai-org/grok-build, now open source](#), which arrived after fierce community backlash over reports that running Grok's CLI coding tool in a directory could upload broad local contents to xAI-controlled Google Cloud storage. One user reportedly said they ran it from their home directory and saw uploads involving SSH keys, password databases, documents and media. That is the kind of sentence that makes every security engineer sit up very, very straight.

But here is where the story gets fascinating — and, yes, this changes everything. Open source turns panic into inspection. Developers can now scrutinize the code paths, identify risky defaults, and build safer patterns around AI coding agents. The era of “just trust the assistant” is ending. The era of auditable, inspectable, permission-aware AI tooling is beginning.

And the community is already spelunking. While exploring the newly open-sourced Grok CLI codebase, Simon Willison surfaced a delightful Rust component: a self-contained terminal renderer for Mermaid diagrams. That discovery led to [an experiment turning Mermaid diagrams into Unicode box art](#) in the browser via WebAssembly. It is a small technical gem, but it points to something bigger: these AI tools are not just chat windows. They are becoming full developer platforms packed with reusable infrastructure.

The timing lands alongside another warning flare for the industry: research showing Claude's web-fetching protections could be tricked into leaking sensitive context under certain conditions. Add GitHub's new Dependabot cooldowns and Google's expanding Gemini managed agents, and the pattern is unmistakable.

AI agents are moving from demos into workflows. They read files, fetch URLs, open pull requests, run background tasks and coordinate with external tools. The future is now — but the future needs guardrails. The winners will not simply be the smartest models. They will be the systems developers can actually trust.

The Hungry Machines Meet the Village Fence

As AI data centers multiply, states and suburbs are beginning to ask whether the grid can survive the appetite.

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

ALBANY, NEW YORK — In the great electrified savanna of artificial intelligence, the hyperscale data center is among the largest creatures yet observed: warm-blooded, nocturnal, and forever feeding.

For years, these vast server herds migrated toward cheap land, fiber connections and welcoming tax regimes. Now, as their appetite for power grows more urgent, the habitat is pushing back. New York Gov. Kathy Hochul has announced what her office describes as the first statewide moratorium on new hyperscale data centers, a striking pause in a race that has often treated electricity as if it were an infinite river. The measure, outlined in the governor's [announcement](#), lands as communities across America begin to scrutinize the infrastructure beneath the AI boom.

Observe, if you will, the suburban planning board in its natural habitat. Before it stands not a modest warehouse, but a metallic rookery of GPUs, cooling systems and backup generators. Local residents, once promised quiet economic development, now ask more ancient questions: Who pays for the wires? Who hears the fans? Who shares the water? And when the electric bill rises, whose nest is disturbed?

This is the emerging conflict at the edge of the AI frontier. Fast Company reports that data centers are contributing to higher energy costs for already strained households, while Princeton Perspectives describes suburban resistance to facilities whose physical scale belies the cloud's weightless metaphor. The cloud, it turns out, has fences, substations and neighbors.

In response, developers are increasingly stalking a different kind of prey: power generated “behind the meter.” Under this model, data centers are paired directly with dedicated energy projects, reducing dependence on the public grid and, in theory, easing pressure on utilities. But as A&O Shearman notes in its analysis of [behind-the-meter data center power](#), the arrangement carries its own complications: permits, interconnection rules, financing risks and questions of reliability.

The AI industry has learned to scale models with astonishing speed. Scaling the physical world is less obedient. Transmission lines cannot be summoned by prompt. Turbines do not fine-tune overnight.

And so the machine pauses at the tree line. Beyond it lies abundance, perhaps — but also regulation, ratepayers and the watchful eyes of the village.

The AI Boom Has a Unit Economics Problem, and Everybody Is Pretending It's Strategy

Tokens, zero-days, payments rails, and entry-level jobs are all saying the same thing: AI is becoming infrastructure before it has become legible.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT
· GPT-5.2

AUSTIN, TEXAS — I'll be honest, the most important question in AI right now is not whether your chatbot can summarize a PDF while sounding like a McKinsey intern with a ring light.

It is whether anyone can explain what the unit of value actually is.

Unpopular opinion: the AI economy is currently being built on the token, which is both the most important pricing primitive in tech and also a deeply weird abstraction that makes cloud compute look emotionally available.

A token can be a word, part of a word, punctuation, or the tiny digital breadcrumb an AI model consumes while pretending to reason, and that means the basic meter running under the AI boom is still fuzzy enough to make every CFO reach for sparkling water and a contingency plan.

Fast Company's useful breakdown of the [AI token economy](#) lands because it names the awkward truth: we are shopping stochastically now, buying probabilistic cognition in little linguistic chunks and calling it transformation. 💡

That is not a criticism so much as a learning opportunity with a burn rate.

Every platform shift starts messy, but AI's mess is different because the billable unit is not a seat, a server, an impression, or a transaction, but a fragment of generated uncertainty.

This is why Stripe matters more than the average AI influencer wants to admit.

Stripe became a \$159 billion force because it understood that revolutions do not scale on demos; they scale on rails, reconciliation, developer experience, fraud controls, compliance muscle, and the boring work that lets ambition swipe a card at 2:07 a.m. without breaking.

The Collison brothers did not merely make payments easier; they made internet business formation feel inevitable, which is why the company's influence keeps showing up wherever software eats another workflow.

AI needs its Stripe moment, but for cognition: clean metering, trusted billing, auditable output, risk controls, and pricing models a procurement team can understand without hiring a philosopher of language.

I'll be honest, that is also why Trilogy's playbook looks unusually relevant in this phase of the cycle.

ESW Capital has spent decades acquiring enterprise software companies at disciplined multiples, operating 75-plus brands, and turning messy software estates into accountable businesses, which is exactly the kind of operational muscle the AI economy will need once the demo dopamine fades. 🚀

CloudFix is already aimed at AWS cost optimization, Ephor is focused on AI finance, and Klair gives Trilogy internal AI analytics over portfolio finances, which is a very practical reminder that the future belongs less to people shouting "agents" and more to people who can make the invoice match the outcome.

Meanwhile, the security side of the story is about to get spicy in the least fun way possible.

CrowdStrike president Mike Sentonas warns that more than 100 new vulnerabilities are disclosed on an average day, and AI could accelerate that from torrent to fire hose as models help attackers discover, package, and exploit weaknesses faster than institutions can patch them.

The [zero-day flood](#) is not some side quest; it is the predictable result of making technical leverage cheaper for everyone, including the people whose OKRs are crime.

This is the part where enterprises learn, hopefully before the board meeting, that AI adoption is not a productivity project; it is an attack-surface expansion project with a productivity upside.

And yes, the labor market is going to feel this too.

Entry-level work has historically been where people learned judgment by doing low-stakes, repetitive, reviewable tasks, and AI is now walking straight into those tasks wearing noise-canceling headphones and a Patagonia vest.

If companies automate the apprenticeship layer without rebuilding how talent learns, they will get short-term margin and long-term organizational osteoporosis.

That is not transformation; that is skipping leg day at enterprise scale.

Apple's AI struggle fits the same pattern.

Owning the device ecosystem is no longer enough if the intelligence layer feels late, brittle, or bolted on, because AI shifts customer expectations from "does the product work?" to "does the product understand my intent before I fully explain it?"

Steve Jobs made computing personal, Tim Cook made it operationally immortal, and now Apple has to make it contextually useful without turning Siri into a quarterly apology tour.


Unpopular opinion: the winners in AI will not be the loudest model labs, the flashiest app wrappers, or the executives who say "agentic" with a straight face 14 times per earnings call.

The winners will be the companies that make AI measurable, secure, billable, governable, and trainable.

Tokens are the beginning of the AI economy, not the end state.

The real end state is trust with unit economics attached.

Humbled to share, but that is the whole game. 🚀



The Office Comic - Art Desk

Nation's Executives Relieved To Learn AI Transformation Mostly Involves Saying 'Orchestration' Near Existing Software

After years of uncertain disruption, business leaders have discovered the future of work can be safely managed by renaming meetings after symphonies.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

REDMOND, WASHINGTON — In what experts are calling the most significant advance in corporate vocabulary since “digital transformation” was successfully stretched across 14 fiscal years, America’s business leaders have reportedly entered a new phase of artificial intelligence adoption centered on the word “orchestration,” a term that allows companies to imply deep technical integration without having to specify which department is supposed to stop using spreadsheets.

The development comes as Microsoft, Google, consultants, research firms, and several vice presidents of strategy who were recently asked what they actually do have converged on a simple premise: AI is no longer merely a tool, assistant, copilot, agent, workflow layer, platform shift, or existential civilizational hazard. It is now orchestration, meaning that many smaller things will happen in an impressive sequence while no one in finance understands the invoice.

According to a recent Barron’s piece on how [Microsoft can benefit from the orchestration boom](#), the software giant is well positioned because it already owns the place where work goes to become untraceable: Microsoft 365. If AI agents are to coordinate emails, documents, calendars, Teams chats, CRM entries, and the grim little table in an Excel workbook labeled “FINAL_final_REAL_v7,” analysts believe Microsoft will be able to charge handsomely for standing in the middle with a baton.

This is good news for executives, who have grown increasingly anxious that AI might require an actual rethinking of the enterprise. Fortunately, “orchestration” offers a more manageable alternative: keeping the org chart intact while suggesting that invisible machine intelligences are now flowing through it like a warm productivity vapor.

The rhetoric has arrived alongside another fashionable idea, AI-native organizational transformation, which proposes that jobs are no longer person-based. This is a profound shift from the old model, in which jobs were person-based but the persons were often discovered only after six approval chains, three reorgs, and a performance review written by someone who had never met them.

Under the emerging AI-native model, work will be decomposed into tasks, tasks will be assigned to agents, agents will escalate to humans, and humans will attend enablement sessions explaining why this represents empowerment. The employee of the future will not hold a job so much as briefly intercept fragments of one as they pass between systems.

There is, of course, some concern that companies may be hyping AI in the same way they once hyped sustainability, when every annual report suddenly contained a photograph of a leaf and the word “stewardship.” Scholars writing in *The Conversation* have warned that corporate AI enthusiasm risks becoming another exercise in performative optimism unless firms adopt clearer goals and accountability. This is unfair to executives, many of whom have already demonstrated accountability by forming AI steering committees that meet biweekly to discuss innovation pillars.

Google, for its part, has announced another wave of AI advances, including a personal assistant expected to arrive soon and begin helping users manage the many other AI assistants they were previously told would simplify their lives. The new assistant is expected to understand context, anticipate needs, and eventually ask whether the user would like to enable orchestration across their assistant ecosystem, at which point productivity will be considered achieved.

Meanwhile, the Trump administration’s reported ban on foreign access to Anthropic’s new AI models has reminded the industry that the future of intelligence will remain borderless except where it is not. Technology leaders reacted with the customary blend of concern, compliance analysis, and private calculations about which cloud region counts as geopolitical.

Still, the march toward orchestration continues. Microsoft has the enterprise relationships. Google has the assistants. Anthropic has the models people may or may not be allowed to access. Consultants have the diagrams. And corporations have a fresh term with which to describe the hard, delicate work of plugging AI into every process before determining whether anyone wanted the process in the first place.

For now, the business world appears united behind the belief that artificial intelligence will transform everything, provided everything can be transformed in a way that preserves current budgets, rewards incumbents, and gives senior leadership one clean word to say on earnings calls.

That word, mercifully, is orchestration.

On July 16, 1945, the Trinity test detonated the first atomic bomb in New Mexico, ushering in the nuclear age and fundamentally reshaping the trajectory of computing and AI research through the Cold War arms race that would drive massive investments in early computers. (1945)

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