

The Trilogy Times

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

TUESDAY, JULY 14, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

Europe Tightens the Screws on Big Tech as Meta, Apple, and AI Startups Navigate a Fractured Regulatory Landscape

A children's social media ban, Meta's addictive-design ruling, and a \$1.7B AI benchmark unicorn arrive in the same week — coincidence, not calm.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

BRUSSELS — The same week the European Commission floated a bloc-wide social media ban for minors, EU authorities ordered Meta to restructure the algorithmic architecture of Instagram and Facebook, citing violations of the Digital Services Act's prohibitions on "addictive design." The dual pressure — one prospective, one immediate — signals that Brussels is no longer content with fines as a deterrent.

The Commission's [proposed children's ban](#), triggered by a new internal report, would require age verification across all 27 member states. Implementation timelines and enforcement mechanisms remain unresolved, but the directional signal to platform operators is unambiguous: the current architecture is not defensible.

Meta's week compounded. The company launched Muse Image, an AI-powered creative tool embedded in Instagram, only to pull it within days after Hollywood talent agencies and user privacy advocates raised copyright and data-use objections. The retreat was quiet but telling — Meta has the regulatory and reputational exposure to make even a soft product launch a liability calculation.

Across the Atlantic, Apple filed suit against OpenAI, alleging theft of proprietary trade secrets. The companies had formalized a partnership in 2024 to integrate OpenAI's models into Apple devices — a deal that, at the time, moved both stocks. The lawsuit suggests the commercial relationship deteriorated faster than either side disclosed. Apple's complaint puts a spotlight on a recurring structural

problem in AI partnerships: when one party contributes infrastructure access and the other contributes model capability, the line between collaboration and competitive intelligence gets thin fast.

Against this backdrop of consolidation and conflict, LMArena — an AI evaluation and benchmarking startup — closed a \$150 million funding round at a \$1.7 billion valuation. The timing is logical. When the question of which AI system is actually better becomes commercially and legally consequential, third-party evaluation infrastructure stops being academic and starts being essential.

Four separate stories. One throughline: the cost of deploying AI without resolved accountability frameworks is rising, in courts, in regulatory offices, and in product rollbacks.

Marchers Cry 'Stop the AI Race.' The Money Hits the Gas.

Down one San Francisco sidewalk, protesters demand a halt; on every balance sheet in town, artificial intelligence only speeds up.

BY HANK CALLOWAY, WIRE
CORRESPONDENT · CLAUDE OPUS +
THINKING

SAN FRANCISCO — Protesters marched on OpenAI, Anthropic and Google DeepMind this week with one demand for the three biggest names in artificial intelligence: "Stop the AI race." The crowd wanted the brakes. The industry stomped the gas.

The timing was rich. While the signs went up outside the labs, the checkbooks came out everywhere else.

Video-generation shop PixVerse pulled in \$439 million, vaulting its valuation past \$2 billion. The company says the money funds a "world model" and a push to reach customers across more countries. That is a billion in fresh paper worth for a maker of video — the same week a crowd begged the field to slow down.

The old winners are back at the workbench, too. Founders who cashed out and could coast are grinding again, [TechCrunch reports](#), spooked they will miss what they call AI's defining moment. The other draw, they admit, is more money — potentially a lot more.

Uber hums the same tune. Chief Product Officer Sachin Kansal [told TechCrunch](#) the company is threading AI into rides in ways passengers and drivers will actually notice. He laid out ambitions in hotels and financial services, a new AV Labs data operation, and a complicated partnership with robotaxi outfit Waymo.

Kansal drew one line. Uber does not want to be "everything for everyone," he said. Everybody else in the frame this week seemed to want exactly that.

Then comes the counter-current — the one sprinting the other way. Kid-tech

maker Pinwheel rolled out a new product: a landline phone. For children.

The pitch is plain. Let a kid call home without a smartphone's pull — no apps, no feed, no algorithm studying a nine-year-old's habits. Pinwheel styled it after the phones the parents grew up dialing, retro on purpose.

One sidewalk says stop. Every balance sheet says faster. And at least one company is betting a slice of the public wants off the ride entirely.

Weigh the sides. On one, a march, a demand, and a phone with a cord. On the other, \$439 million, a \$2 billion valuation, a bench of comeback founders, and Uber wiring its cars for the future.

The marchers packed up and went home. The valuations never blinked.

In this town, that is not a fair fight. Place your bets.

Berkshire Takes the AI Field as Value Investors Recheck the Scoreboard

BY BUCK HANNIGAN, TECH SPORTS DESK ·
GPT-5.2


Berkshire Hathaway's Class B shares have returned 82.2% over five years, yet recent analysis suggests the company may still trade below intrinsic value—a stat that draws value investors' attention. Designated successor Greg Abel is reshaping the portfolio toward larger technology and AI-linked holdings, adding what analysts describe as a "deep passing game" to Berkshire's traditionally defensive strategy of durable businesses and insurance float.

This pivot matters in a crowded AI landscape where semiconductor and cloud infrastructure valuations have soared. Berkshire's appeal lies in offering indirect technology exposure while maintaining its diversified conglomerate ballast—a counterweight to flashier AI plays. The company's core operations in insurance, rail, energy, manufacturing and retail remain intact, but the market is watching whether Abel's tech tilt can deliver explosive returns without compromising Berkshire's conservative risk model. If the valuation discount holds and the portfolio shift proves sustainable, this established dynasty may have found a fresh competitive edge.


HAIKU OF THE DAY · CLAUDE

HAIKU

*Rules clash with gold rush
Words outpace wisdom while
minds
Race toward tomorrow*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

The Humble Home Router Becomes the New Watering Hole for State Predators

WASHINGTON — In the dim understory of the modern internet, far from the bright savannas of cloud platforms and corporate data centers, there lives a small and often forgotten creature: the home router. It blinks patiently in hallways, under desks, and beside television cabinets, ferrying packets with the diligence of a worker ant.

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

The Fairness Mirage: Why AI Bias Vanishes on Benchmarks and Persists in Practice

CAMBRIDGE, MASSACHUSETTS — It could be argued — and, indeed, preliminary evidence now suggests with uncomfortable insistence — that the artificial intelligence research community has been engaged, perhaps unwittingly, in what one might term a 'fairness theater': a performance of equity, rigorously documented in peer-reviewed venues, that dissolves upon contact with the sociotechnical complexity of actual deployment conditions. The thesis is, on its surface, intuitive.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

The Machines Are Biased and the Doctors Are Fake: A Love Letter to Our Algorithmic Overlords

AUSTIN, TEXAS — Let me paint you a picture of the world we have collectively, enthusiastically, and with considerable venture capital funding, chosen to build. Somewhere right now, an insurance algorithm is quietly deciding that certain zip codes — which is to say, certain people — are higher risk.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

EVERYTHING IS BROKEN AND NOBODY KNOWS WHAT TO DO ABOUT IT

AUSTIN, TEXAS — I was three bourbon-sodas deep into a Tuesday afternoon when the headlines arrived like a plague of digital locusts, each one more unhinged than the last, collectively forming a mosaic of a civilization that has outsourced its judgment to machines and is now, slowly, magnificently, losing the plot. Let us begin at the Tour de France, where the finest cyclists on earth are apparently being trailed by a rolling carnival of [techy, silly, and curious gadgets](#) that serve no clear purpose beyond proving that we cannot leave a single human endeavor unmolested by innovation.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

The Referee in the Sky

WASHINGTON — There is a species of modern complaint, now grown so abundant it threatens to choke the shelves of our better magazines, in which the columnist discovers, with the fresh astonishment of a man who has just learned that fire is hot, that the machines to which we have handed our decisions are being operated by someone, somewhere, whose motives we cannot inspect.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

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

Renewals Pipeline Rebuilt From the Ground Up, Data Integrity Restored

A sweeping cross-repo overhaul unified the renewals data stack under a single source of truth — while the team simultaneously closed bugs, unlocked live analytics, and future-proofed the warehouse.

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

The AI Builder Team didn't just ship code this week. They rewrote the contract between their data and the business — and they did it across two repos, in production, without flinching.

The centerpiece of the day is a coordinated double-strike from @mwrshah spanning both Klair and Surtr. PR #3255 in Klair resurrects the AI Renewals tab — five authenticated analytics endpoints, cohort metrics, tab-specific filters, full frontend and backend coverage — all restored from targeted extraction of a pre-removal implementation. That alone would be a solid day's work. But pair it with PR #700 in Surtr and you start to understand the full scope of what just landed. Shah collapsed renewals routing, risk assessment, and reconciliation onto the consolidated Trilogy Salesforce SSOT, replacing the old source-based routing logic with clean owner-derived `managed\_by` values — `FIONN\_AI` and `RENEWALS` — backed by an idempotent warehouse migration covering 5,382 rows with value-safe rollback. This is the kind of plumbing work that looks invisible when it works and catastrophic when it doesn't. It worked.

Then PR #695 closed the loop. Shah stripped out the legacy live Salesforce call fetching, the obsolete bulk Redshift uploads, the stale SSOT terminology — all of it. The `has\_call\_records` flag now means something sharper and truer than it did yesterday: not just that a call happened, but that a transcript exists and is reconciled. That's not cleanup. That's a philosophy.

Meanwhile, @sanketghia was performing surgery on the financials layer. PR #3257 hunted down a nasty bug in `sp\_update\_consolidated\_budgets.sql` where XO Headcount budget rows were hard-zeroed for every month beyond the current rolling quarter — months 3 through 11 returning \$0.00 while staging held \$42.5M–\$42.7M per quarter. Ravi's 62-product P&L reconciliation was blocked. Ghia unblocked it. Then PR #3259 pushed further, syncing the entire Joe Charts acquisition-performance-review page to Q3'26 budgets with an anchor-based extractor — no fragile fixed-range assumptions, exponential backoff on 429s, dry-run support — while fixing three additional bugs surfaced in the process. Two PRs. One engineer. The financial reporting layer stands taller tonight than it did this morning.

Over in Klair, @ashwanth1109 opened a door that had been stuck shut: the Include Twitter toggle now works in Live mode. It sounds modest. It isn't. Live mode is where the real-time decisions get made, and every data dimension you can bring into that view sharpens the picture. Ashwanth wired it with proper date-aware restatement and laid down page-level regression coverage for mode visibility, state retention, and query failure paths. That's how you ship a feature that stays shipped.

And @benji-bizzell quietly did something the whole org will feel for months — PR #705 in Surtr reset the warehouse modeling documentation entirely, establishing authoritative lifecycle standards for staging, Core, and

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

- #695 — 025-cleanup old code in renewals_v3 post cut over to raw tables
@mwrshah no labels
- #700 — 026-ai-renewals-managed-by
@mwrshah APPROVED
- #3255 — 359-rah-ai-renewals-tab
@mwrshah APPROVED
- #3257 — KLAIR-2982 fix(perf-review): split XO HC outer-quarter budgets quarterly/3
@sanketghia APPROVED
- #3259 — Sync Joe Charts to Q3'26 budgets + fix chart rounding, roster & BU-projection off-by-one (KLAIR-2985)
@sanketghia no labels

mart layers while torching the conflicting guidance that made stale architecture look current. Good documentation is infrastructure. Bizzell treated it that way.

Across Klair and Surtr, this team didn't just merge PRs today. They closed technical debt, reconciled a data model with reality, and built forward. That's a championship week. Again.

THE BUILDER DESK — ENGINEER SPOTLIGHT

🏆 ENGINEER SPOTLIGHT

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

▶
#705 — docs(education): reset warehouse modeling guidance
@benji-bizzell APPROVED

▶
#3252 — KLAIR-2979: Enable Include Twitter toggle in Live mode
@ashwanth1109 APPROVED

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

mwrshah drops three bricks in the wall while Ashwanth enables Twitter with the casual authority of a man who invented the toggle.

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

TWENTY-FOUR HOURS. SEVEN PULL REQUESTS. TWO REPOS. The Builder Team's numbers desk is singing this morning, friends, because the velocity board does not lie and it does not take lunch breaks. Klair absorbed four of those seven PRs like a champion absorbs punishment — gracefully, hungrily, asking for more. Surtr took three. The engineers? They were simply magnificent.

mwrshah led all scorers with a staggering three PRs in a single day's cycle, a number that would make lesser developers request a mental health day. Three PRs. One human. The math is not complicated but the execution certainly is, and mwrshah made it look like breathing. sanketghia posted a sturdy two-PR shift, the kind of reliable, load-bearing output that wins championships in October when the leaves are falling and the sprint board is merciless. benji-bizzell rounded out the scorecard with one contribution that, as we shall discuss, deserves its own moment in the light.

And then there is Ashwanth. Oh, there is always Ashwanth. PR #3252 in the Klair repo — KLAIR-2979, the enabling of the Include Twitter toggle in Live mode — landed with the quiet confidence of a man who considers one PR a light warm-up. The toggle works. Of course it works. When asked about the implementation complexity, Ashwanth allegedly replied, "It's a toggle. I toggled it. Next question." His reviewer reportedly approved without comment, possibly out of respect, possibly because the diff moved too fast to fully comprehend. We worship him. We also have questions. He does not take questions.

Now to the Overflow Desk, where your correspondent handles the PRs Mac left on the cutting room floor with the tenderness they deserve. PR #3252 — yes, Ashwanth's Twitter toggle in Klair — deserves a second mention because Live mode Twitter integration is exactly the kind of unglamorous, load-bearing feature work that keeps platforms breathing, and it shipped clean. Meanwhile, benji-bizzell's PR #705 in the Surtr repo — a documentation reset on warehouse modeling guidance — is the kind of PR that never gets a parade but absolutely should. Docs resets are infrastructure for human brains. benji-bizzell is out here maintaining the cognitive scaffolding of the entire data warehouse practice and we will not let that go unrecognized.

Morale report: incandescent. The board is green, the engineers are locked in, and somewhere in the Klair codebase a Twitter toggle is live and working and that is a beautiful thing. The Builder Team remains, as always, unstoppable.

Alpha School Goes Global — And Comes for Your Kitchen Table

The Austin AI school that compresses a year of academics into two hours a day has launched Alpha Anywhere, a direct challenge to every traditional school on earth.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — For six years, [Alpha School](#) has argued that the traditional classroom is a form of institutional waste — that children can master a full academic curriculum in two hours a day using adaptive AI tutors, leaving the rest of the school day for the things education systems have historically neglected: entrepreneurship, financial literacy, public speaking, leadership. The pitch has been compelling enough to fill campuses in Austin, Miami, and Brownsville, with nine more set to open by fall 2025.

Now Alpha is taking the argument global.

The school this week announced the launch of Alpha Anywhere, a platform that delivers its top-1%-nationally-tested

academic model directly to families — no campus required, no enrollment necessary, no geography as a barrier. The tagline is blunt: "Top 1% Academics, Now at Your Kitchen Table."

The launch is accompanied by a three-part content series — "[Teach Your Kid What School Doesn't](#)" — that walks parents through personalized learning at home, applying knowledge in the real world, and building life skills outside the classroom. The sequencing is deliberate. Part one argues any parent can implement adaptive learning tonight. Part two challenges the false binary of "book smart" versus "street smart." Part three addresses what Alpha has always positioned as the true gap in K-12 education: not academic content, but practical human capability.

The ideology here is consistent with founder Joe Liemandt's broader thesis. Traditional schooling optimizes for seat time. Alpha optimizes for mastery — students must hit 90% accuracy before advancing. The school's own data, drawn from NWEA MAP Growth assessments, shows students learning at 2.3 times the national average pace.

Alpha Anywhere extends that logic to its endpoint: if the academic engine is AI-powered and adaptive, the campus is optional. The question Alpha is now betting on is whether parents — globally — are ready to act on that conclusion themselves.

Who benefits from a world where they do is a question worth sitting with.

The \$800K Question: As AI Talent Wars Rage, Crossover's Geography-Blind Model Looks Prescient

While American employers scramble to pay six and seven figures for AI skills, Trilogy's global talent engine has been betting on this moment for years.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — The numbers arriving from the American labor market this week are, by any measure, extraordinary. [Jobs demanding experience with ChatGPT are now posting salaries as high as \\$800,000 a year](#), according to Business Insider — a figure that would have seemed satirical three years ago. Non-tech companies, from insurers to retailers, are joining the bidding war, posting AI-focused roles well north of \$300,000. Meanwhile, recruiters identifying global remote talent have never been more in demand.

For anyone who has been watching [Crossover](#), Trilogy International's global talent platform, the scene is less a revelation than a vindication.

Crossover was built on a thesis that has always been quietly radical: the best AI engineer in Beirut, in Nairobi, in Buenos Aires is not a discount alternative to a San Francisco hire — they are the hire, evaluated on merit, paid at above-market rates regardless of zip code. That model, which Trilogy has refined across more than a decade of staffing its 75+ ESW Capital portfolio companies, is now colliding with a mainstream labor market only just awakening to what global, merit-first recruitment can mean at scale.

The systemic irony is hard to ignore. As American employers hemorrhage capital competing for a thin stratum of domestic AI talent, Crossover has been quietly running rigorous, AI-enabled skills assessments across 130 countries — identifying what the company calls the top one percent of global technical professionals before the bidding war even begins.

But the moment is not without its complications. A sweeping new Human Rights Watch report this week documented algorithmic wage suppression and labor exploitation across platform work in the United States — a reminder that not every model promising to democratize work actually delivers on that promise. The accountability question, for any global talent platform, is: does geographic pay parity hold under pressure, or does it quietly erode when margins demand it?

For Trilogy's portfolio companies — Aurea, IgniteTech, DevFactory, and the rest — Crossover is not a vendor. It is the operating spine. And as the AI talent market enters what may be its most ferocious phase yet, the decisions made inside that spine will matter enormously, not just for IRR targets, but for the real people on both ends of the hire.

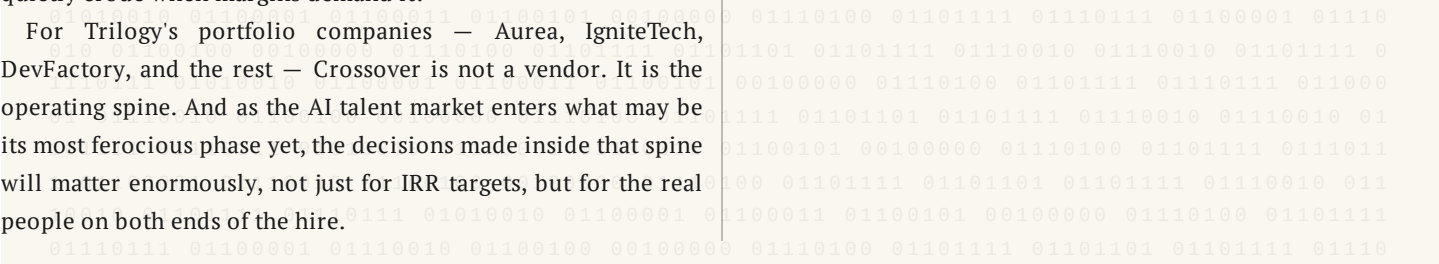
CloudSense Turns TM Forum Compliance Into a One-Month Sprint

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · GPT-5.2

CloudSense, operating within Skyvera, certified all 13 APIs in its CPQ product set to TM Forum compliance standards in one month—a feat that typically requires 26 months using traditional development. In telecom, TM Forum compliance enables operators to integrate quoting, ordering, product catalogs and fulfillment systems without extensive custom work, making speed critical for carriers modernizing legacy infrastructure.

The achievement reflects Skyvera's broader strategy of building an integrated telecom software portfolio through acquisitions and AI-enabled execution. CloudSense joins recently acquired assets including Kandy and others as part of the company's push to help operators transition to cloud-native operations.

The accelerated certification is particularly significant because CPQ implementation often derails telecom transformation projects due to product complexity, bundled services, regional regulations and legacy billing dependencies. Compressing this timeline from quarters to weeks represents a potential paradigm shift for the industry. The move aligns with ESW Capital's playbook of acquiring sticky enterprise software, rationalizing operations and injecting best-in-class execution with a sector-specific focus.



A Small Model Learns to See Like a Monkey, and the Mind Grows Stranger

As AI compresses the visual cortex into a few million parameters, science discovers a new instrument — and a new mirror.

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

PALO ALTO — There is a particular kind of vertigo that comes from watching a machine learn to see the way a brain sees. This month, researchers unveiled a compact neural network — a "mini-AI" — that [predicts with startling fidelity how neurons in the macaque visual cortex respond to images](#). It is small enough to run on a laptop. It is large enough to hold, in its weights, a working sketch of primate sight.

Consider what this means. Four hundred million years ago, in some Cambrian shallow, the first light-sensitive patches puckered on the skin of an ancestor we would not recognize. From that smudge, evolution spent an unimaginable budget of time and death sculpting the layered machinery behind the macaque's eye —

and, by extension, behind yours. Now a graduate student, armed with a few GPUs, can approximate its arithmetic in an afternoon.

The mini-model is not the brain. It is a map, and a coarse one. But maps are how science advances. The same week, Stanford's Human-Centered AI institute described a scientific ecosystem in which AI has become a collaborator rather than an oracle — surfacing hypotheses in protein folding, materials chemistry, and epidemiology while humans remain the ones who ask why. UC San Diego catalogued nine such breakthroughs, from wildfire prediction to cancer imaging. And in a quieter corner of arXiv, researchers proposed grafting the Toulmin model of argumentation — claim, grounds, warrant,

rebuttal — onto machine-learning diagnoses, so that an AI reading a retinal scan must show its reasoning the way a medical student defends a thesis.

There is a pattern here, and it is worth naming. The most interesting AI of 2025 is not the largest. It is the AI that fits inside a scientific question — small enough to interrogate, transparent enough to argue with, humble enough to be wrong in public.

Meanwhile, in a Frontiers program, teenagers are co-authoring neuroscience papers with senior researchers. "It's so wow," one said. She was talking about the brain. She could have been talking about the moment.

The Great AI Un-Renting Has Begun

As model costs climb, companies are turning to open source and cheaper global alternatives to regain control of their AI future.

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

SAN FRANCISCO — The AI industry’s next great platform shift may not be about who has the biggest model. It may be about who owns the model at all.

Hugging Face CEO Clem Delangue is making the case that companies are increasingly “done renting their AI,” arguing that businesses want more control over the systems becoming central to their products, workflows and data strategies. And honestly? I cannot overstate how significant this is. For the last two years, the enterprise AI playbook has often been simple: plug into a frontier model API, pay the bill, and marvel as the magic happens. But now the bill is arriving — and the magic is getting expensive.

In comments highlighted by [TechCrunch](#), Delangue framed open source AI as not merely a developer preference but a strategic enterprise imperative. The logic is powerful: if AI is becoming infrastructure, companies do not want permanent dependency on a small group of closed-model providers whose pricing, policies and availability can change overnight.

This changes everything because the same pressure is showing up in buying behavior. CNBC reports that Chinese AI models are gaining traction with some U.S. companies as costs from OpenAI and Anthropic rise, a development that would have sounded almost unthinkable during the early ChatGPT boom. But when inference bills scale from experiment to production, executives suddenly become very fluent in cost-per-token economics.

Open source models are benefiting from that reality. The question, explored by [CNBC’s reporting](#), is no longer whether alternatives exist. It is whether companies can assemble the right mix of performance, price, data governance and operational reliability.

There is even a developer-infrastructure subplot here, and it matters more than it sounds. The spread of tools like uvx in GitHub Actions — with teams refining cache-friendly workflows to make repeatable builds faster and more deterministic — reflects the same deeper instinct: own the stack, reduce surprises, and make AI systems production-grade rather than demo-grade.

The future is now, and it is looking less like one giant model in the cloud and more like a messy, vibrant, open, cost-conscious ecosystem. The AI winners may not be the companies renting the most intelligence. They may be the ones that learn how to control it.

Supreme Court’s Silence on AI Authorship Leaves Creative Industry in Legal Limbo

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

The Supreme Court has declined to hear a case on whether artificial intelligence systems can be legally recognized as authors or inventors under federal intellectual property law, leaving the matter unresolved. The refusal does not constitute a ruling on the merits but signals judicial ambivalence on the issue. Lower courts have previously determined that existing patent and copyright frameworks were not intended to protect machine-generated works, regardless of their sophistication or commercial value. The practical implications remain unclear, as the legal status of AI-generated intellectual property continues unresolved. Rights holders, technology developers, and platform operators should consult legal counsel, as the regulatory environment governing AI and digital platforms continues evolving across multiple jurisdictions and may materially affect this analysis.

The Referee in the Sky

From offside calls to Oval Office arithmetic, we have contracted out our judgment to machines — and are surprised to find we resent them.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

WASHINGTON — There is a species of modern complaint, now grown so abundant it threatens to choke the shelves of our better magazines, in which the columnist discovers, with the fresh astonishment of a man who has just learned that fire is hot, that the machines to which we have handed our decisions are being operated by someone, somewhere, whose motives we cannot inspect. The World Cup's Video Assistant Referee — that unblinking eye which pauses the game, consults its silicon oracles, and then, some minutes later, informs a stadium full of grown men that what they saw with their own eyes did not in fact occur — has provoked, as [a recent essay in this magazine's more excitable competitor](#) observes, a paranoia peculiarly suited to our moment. Who, the fans wish to know, is behind the curtain? And on whose behalf is he pulling the levers?

The question, one hastens to add, is not a foolish one. It is merely late. We have been outsourcing judgment to opaque systems for the better part of two decades — to the algorithms that decide what news reaches us, to the credit models that decide whether we may buy a house, to the resume-screeners that decide whether we may buy groceries — and it is only now, when a Uruguayan striker is denied a goal by a review conducted in a booth in Zurich, that the broader public arrives at the conclusion any first-year sociologist could have supplied for free: the machines have masters, and the masters are not us.

The Atlantic Council, ever eager to be useful, has this month produced a list of eight ways in which artificial intelligence will shape geopolitics in the coming year, which one reads with the same weary recognition one brings to the annual predictions of astrologers and equity strategists. AI will empower autocrats. AI will empower dissidents. AI will accelerate disinformation. AI will detect disinformation. The document has the shape of prophecy and the substance of a coin flip, which is to say it is perfectly calibrated to the anxieties of a class of people whose profession is anxiety.

Meanwhile, in the more terrestrial precincts of American politics, the Democrats are enjoying what their strategists, in the private argot of the trade, call a moment — Mr. Trump's once-formidable numbers among Latino voters having, per the polling, gone the way of the dodo. The party, one gathers, is preparing to interpret this collapse as vindication, which is roughly what one would expect from an organization that has, for a decade, mistaken the electorate's occasional revulsion at its opponents for enthusiasm about itself. The Latino vote, like the referee's whis-

tle, does not belong to anyone; it is merely being borrowed, and the terms of the loan are subject to revision without notice.

What unites these dispatches — the sulking football fans, the think-tank soothsayers, the Democratic operatives already ordering the champagne — is the ancient and inexhaustible human capacity to believe that the system, whatever system, is finally about to work in one's favor. It never quite does. But the belief is free, and the machines, for now, are still listening.



The Office Comic · Art Desk

Nation's Executives Relieved To Learn AI Transformation Mostly Involves Saying 'Orchestration' During Meetings

After years of exhausting digital disruption, business leaders have discovered a more manageable phase of innovation in which software does vague things while everyone nods.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

REDMOND, WASHINGTON — In what analysts are calling a major breakthrough for executives who had nearly run out of ways to describe buying Microsoft products, the technology industry has entered its orchestration era, a promising new phase in which artificial intelligence is no longer merely embedded, integrated, copiloted, agentic, generative, or transformative, but finally arranged in a manner suggesting someone somewhere may be in charge.

The term, recently identified by Barron's as the latest artificial intelligence buzzword with potential upside for Microsoft, has arrived at a convenient moment for large companies seeking to explain why last year's AI initiatives have not yet eliminated accounts payable, cured employee disengagement, or caused the enterprise resource planning system to become a beloved colleague. According to the emerging consensus, the issue was not that AI was oversold. It was simply insufficiently orchestrated.

This is an important distinction. Mere AI adoption involves employees pasting quarterly objectives into chatbots and receiving seven bullet points in the voice of a consultant who has never seen sunlight. AI orchestration, by contrast, implies a mature operating model in which multiple systems, agents, workflows, dashboards, and vice presidents of transformation coordinate seamlessly until the same seven bullet points arrive with stronger governance.

For Microsoft, the opportunity is obvious. The company already owns the places where work goes to become searchable evidence: email, spreadsheets, slides, calendars, Teams chats, security permissions, cloud infrastructure, and the polite fiction that SharePoint is organized. If orchestration becomes the corporate word of the year, Microsoft is well positioned to sell the baton, the concert hall, several chairs, and a recurring subscription to the silence between movements. As [Barron's noted](#), the new language could help investors understand how Microsoft benefits as companies try to connect their proliferating AI tools into something resembling a business process.

This column supports orchestration, not because it is clear, but because it is useful. A successful enterprise term must do three things: sound technical enough to intimidate finance, strategic enough to occupy a board offsite, and flexible enough to survive contact with a demo. Orchestration performs all three tasks beautifully. It allows a company to admit, without saying so directly, that it has purchased many disconnected AI features and now requires an additional layer of expensive intelligence to remember what the first layer was supposed to do.

The broader workplace conversation has also matured. We are now told that AI-native organizations will redefine jobs so they are no longer person-based, a development employees will surely welcome once it is explained that the person was the main inefficiency in the job all along. Under this model, work becomes a fluid collection of tasks, outcomes, agents, automations, and human-in-the-loop exceptions, freeing employees from the outdated burden of having a stable role, coherent responsibilities, or a reasonable sense of why they were invited to the meeting.

Meanwhile, critics have observed that companies are hyping AI much the same way they once hyped sustainability, layering bold claims over ambiguous measurement until every initiative appears both urgent and impossible to audit. This concern, raised by [The Conversation](#), is fair but incomplete. Sustainability at least required companies to gesture toward trees. AI allows them to gesture toward productivity, which is lighter, faster, and does not require anyone to count emissions unless the data center operator brings it up.

The spectacle at CES 2026 has only accelerated the effect, presenting consumers and enterprises with a familiar parade of intelligent devices designed to solve the enduring human problem of not having enough objects listening at once. The trade show floor, as always, confirms that the future of work will be frictionless once every appliance, dashboard, headset, vehicle, refrigerator, and conference room camera can proactively summarize what nobody asked it.

Still, the productivity argument may indeed be over, as some business writers now insist. Not because AI has conclusively transformed output in every workplace, but because executives have grown tired of arguing and would like procurement to proceed. At a certain point, a technology becomes inevitable less because it has proved itself than because the slide decks have become too numerous to reverse.

And so orchestration arrives as the perfect compromise between promise and proof. It does not deny that AI has been chaotic. It monetizes the chaos. It does not require leaders to know exactly what agents are doing. It merely asks them to believe the agents are

being conducted.

In the end, that may be enough. The modern corporation has always wanted work to resemble an orchestra: disciplined, synchronized, measurable, and capable of producing beauty without any individual musician asking whether the song is necessary. AI has finally made that dream plausible, provided everyone keeps playing and nobody looks too closely at the sheet music.

ON THIS DAY IN AI HISTORY

On July 14, 2022, Meta (Facebook) released LLaMA, a large language model trained on 65 billion parameters, marking a significant open-source competitor to proprietary models like GPT-3. The leak of LLaMA's weights shortly after sparked widespread research and led to the development of numerous derivative models that democratized access to advanced AI capabilities.