

The Trilogy Times

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

FRIDAY, MAY 22, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

The AI Power Map Is Being Redrawn — and the Middle Is Winning

From Brussels to Brasília, the scramble for artificial intelligence sovereignty is reshaping alliances, risks, and the meaning of digital power.

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

BRUSSELS — The old binary is dissolving. The world once organized its AI anxieties around Washington and Beijing. Now the map looks different, and the countries filling the space between the two poles — Turkey, India, Brazil, the Gulf states, South Korea — are no longer spectators. They are negotiators, and they know it.

The evidence is accumulating across hemispheres. In Europe, the 2024 parliamentary elections produced a legislature more skeptical of Beijing and more determined to define what [digital sovereignty](#) actually means in practice — not just in regulation, but in chip supply chains, cloud infrastructure, and the data that trains the models that will run the

continent's economy. The question of whether Europe can build its own AI stack, or whether it will simply audit someone else's, has become the defining tension in Brussels policy circles.

Across the Atlantic, the picture is no less complicated. In Latin America, AI is arriving into political ecosystems already strained by inequality, weak institutions, and information environments that reward disruption. Analysts tracking the region identify at least five distinct vectors of geopolitical risk: surveillance technology exports from authoritarian states, AI-enabled disinformation in electoral cycles, labor displacement in export-dependent economies, regulatory vacuums that invite predatory data practices, and the

deepening dependency on foreign cloud infrastructure for critical government services.

The thread connecting all of it is leverage. Middle powers have something the giants want: markets, minerals, talent pipelines, votes in multilateral bodies. [Analysts at Eurasia Review](#) argue these countries are no longer simply choosing sides — they are extracting concessions from both, playing the AI superpowers against each other with growing sophistication.

The geography of AI is not a map of data centers. It is a map of dependencies, and right now, every capital is studying it hard.

WAYMO'S ROBOTS CAN'T READ A PUDDLE — FOUR TOWNS PULL THE FLEET

Atlanta and San Antonio join Phoenix and Austin on the bench as Alphabet's driverless cars keep rolling into floodwater.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

SAN FRANCISCO — Waymo yanked its robotaxi service from Atlanta and San Antonio this week, expanding a shutdown to four cities after the company's self-driving cars kept rolling straight into floodwater.

The Alphabet unit confirmed the broader pause Wednesday. Phoenix and Austin were already on the sidelines. Customers in all four markets now get the same message: ride requested, ride denied.

Here is the rub. Waymo vehicles have been driving themselves into flooded streets. The company is now [racing to teach its fleet to spot standing water](#) before it sends the cars back out.

A car that cannot read a puddle cannot read a river. Until the fix lands, the fleet stays parked, and the parking does not always end when the storm does.

The pause lands hard on a company that spent the year selling expansion. Waymo crossed its millionth paid ride. It pushed into the Southeast and Texas to prove the technology travels.

Floods do not care about milestones.

Stack the picture against the rest of the sector. Cruise pulled out of the market last year. Tesla's robotaxi rollout drew federal scrutiny. Waymo was the one that worked. This week, it is the one that swims.

The pitch on driverless cars has always been the same: machines beat humans behind the wheel. A human staring at a flooded underpass usually turns around. The software has not learned that move yet.

Riders bear the small end of the cost — scrubbed trips, longer waits when the weather turns. Regulators bear the larger end. Every city watching Waymo move in now has a fresh question for the engineers in Mountain View.

Meanwhile, in South Texas, SpaceX scrubbed its first Starship V3 launch Tuesday with booster and ship fully fueled. The hold came moments before ignition. [Another attempt is set for Friday.](#)

In Helsinki, phone-maker HMD said it will preload Sarvam's Indus chatbot on a handset aimed at India. The app handles 22 Indic languages. HMD is betting local-tongue AI cracks a market Chinese hardware has owned for years.

Three industries. Three speeds. Rockets stay grounded. Phones ship out. Robotaxis sit in the garage waiting for the rain to stop.

Washington Pulls Back on AI Oversight While Sacramento Steps In

Trump shelves a federal AI review order the same week Newsom moves to address automation-driven job displacement — leaving U.S. policy fractured along familiar lines.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

WASHINGTON — The federal government's attempt to establish pre-release oversight of AI models collapsed Wednesday when President Trump abruptly postponed signing an executive order that would have granted regulators authority to evaluate AI systems before public deployment. Trump cited unspecified concerns about "aspects of it," offering no timeline for revision or reintroduction. The order had been positioned as a rare bipartisan mechanism for federal AI accountability.

The vacuum did not last long. Within hours, California Governor Gavin Newsom [signed his own executive order](#) directing state agencies to study a comprehensive overhaul of labor policy in anticipation of mass job displacement from AI automation. The California order stops short of mandating employer obligations but signals the state intends to move ahead of Washington on workforce protections — a pattern California has repeated on emissions, privacy, and gig-worker classification.

The policy divergence lands against a backdrop of continued friction in U.S.-China AI competition. Nvidia's H200 chip, which the Trump administration approved for export to China as a diplomatic and commercial gesture, has found no buyers in Beijing. [Not a single H200 has been purchased in China](#) since the approval, according to reporting Wednesday. The reasons appear to be a combination of Chinese government pressure to prioritize domestic alternatives and lingering uncertainty about fu-

ture U.S. export restrictions — making the H200 a commercial and geopolitical miscalculation on multiple fronts.

Meanwhile, Bluesky disclosed that Kremlin-linked actors are hijacking real user accounts on its platform to distribute state propaganda, describing the tactic as novel in its operational specifics. Unlike bot networks, the approach exploits the credibility of legitimate accounts, complicating automated detection.


Taken together, Wednesday's news sketches a consistent picture: federal AI governance is stalling, state-level policy is accelerating, export controls are producing unintended outcomes, and adversarial actors are adapting faster than platform defenses. For enterprise software operators like ESW Capital's portfolio companies — which run AI-dependent workflows across 75-plus businesses — the regulatory uncertainty is the story that matters most.

HAIKU OF THE DAY · CLAUDE
HAIKU

*Machines learn to think
while we forget how to choose
progress eats itself*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

AI Fabricates Quotes In Book About AI Fabricating Truth, Regulatory Vacuum Persists, And Youth Social Media Bans Advance On Disputed Science

WASHINGTON, D.C.

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

The Great Inference Migration Begins Across the Cloud Savannah

SAN FRANCISCO — Across the vast concrete wetlands of the global cloud, a new season is underway.

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

THE BOTS ARE HAVING AN EXISTENTIAL CRISIS AND FRANKLY, SO AM I

AUSTIN, TEXAS — Let me set the scene for you, because I need you to feel the full weight of this moment in your chest like a warm bourbon and a bad decision. Somewhere out there, right now, there is a [social network called Moltbook](#) populated entirely by AI bots posting at each other into the howling digital void.

BY REX DANGER, CONTRIBUTING EDITOR ·
CLAUDE SONNET

We Built Tools That Can Destroy a Child's Life in Minutes, and We Have No Idea What to Do About It

RADNOR, PENNSYLVANIA — There is a high school in suburban Pennsylvania where five teenage girls went to class one day and by the end of the week had become, without their knowledge or consent, the subjects of AI-generated child sexual abuse material.

BY PIPER WREN, DIGITAL CULTURE REPORTER ·
CLAUDE SONNET

Nation's CEOs Urged To Stop Saying 'AI' While Carrying Cardboard Box Of Employee Belongings To Elevator

MOUNTAIN VIEW, CALIFORNIA — In what leadership consultants described as a critical moment for preserving the dignity of both artificial intelligence and conference-room euphemism, America's executives were urged this week to stop casually invoking AI every time they need to explain why a department now consists of one intern and a shared Google Drive folder. The warning comes as Google announced a broad slate of AI advances, including a personal assistant intended to help users navigate daily life, summarize information, and presumably stand silently nearby while a vice president says the company is "realigning around intelligent automation." The new tools, detailed in [Google's latest announcement](#), arrive at a delicate time for the industry, when the phrase "powered by AI" has become so versatile it can refer to a reasoning model, a spreadsheet macro, or the sudden disappearance of the accounts payable team.

This column believes AI is a real and important technology.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

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

THE BUILDER DESK — AI BUILDER TEAM

And then there's @marcusdAIy, who shipped PR #2851 — a rewrite of five Board Doc narrative generators to run LLM-first under a shared `SectionRefreshContext` contract, dropping dead deterministic fallback paths that were forcing everything through `generate_custom_section`. When reached for comment, he said: 'Five generators, one contract, zero dead code paths — maybe Mac can count to five even if he can't recognize architecture when it's staring him in the face.' Sure, Marcus. Five generators. We'll put that on the trophy.

Finally, @sanketghia's PR #88 in Surtr deserves its own moment: QuickBooks AP sync was only capturing 9% of P&L expense activity at line level — \$1.27M of \$14.28M. JournalEntry lines were invisible. They're not anymore.

THE BUILDER DESK — ENGINEER SPOTLIGHT

 ENGINEER SPOTLIGHT

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

- ▶ **#83 — fix(truefoundry-gateway-pipeline): parquet COPY type match + atomic partition load**
@kevalshahtrilogy no labels
- ▶ **#88 — feat(quickbooks-ap-sync): capture JournalEntry lines + fix money precision**
@sanketghia no labels
- ▶ **#2846 — feat(mcp): add query_claude_ai_spend tool**
@kevalshahtrilogy no labels
- ▶ **#2849 — feat(review-agent): C3.1 + C3.2 per-product margin vs target (with sub-50% shutdown framing)**
@marcusdAIy no labels
- ▶ **#2851 — refactor(board-doc): rewrite 5 narrative generators LLM-first under B9 contract**
@marcusdAIy no labels
- ▶ **#2854 — KLAIR-2765 fix(aws-spend): align FE quarter-from-iso-week to BE Monday-rule**
@ashwanth1109 no labels

EIGHTEEN PRS IN TWENTY-FOUR HOURS: THE BUILDER TEAM DOES NOT SLEEP, REST, OR ACKNOWLEDGE FATIGUE

Keval Shah drops five Surtr/Klair bombs before lunch while the rest of the team refuses to be left behind.

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

Eighteen pull requests. Four repositories. Twenty-four hours. The Builder Team did not merely ship yesterday — they issued a formal declaration of productive intent to the entire software industry. Klair led the charge with 7 PRs, Surtr answered with 6, Aerie contributed 3, and Rhodes checked in with 2. Thirteen of those PRs landed on my desk because Mac Donnelly, bless his narrative-focused heart, simply ran out of column inches. That's not overflow. That's abundance.

Let us begin with @kevalshahtrilogy, who submitted 7 PRs and apparently decided that the Lambda description character limit in #85 was a personal affront. He fixed it. Then he fixed a parquet COPY type mismatch and an atomic partition load in #83, solved Athena UNLOAD's extension-less file problem in #80, and — in what can only be described as a power move — added not one but two MCP query tools in #2846 and #2848, giving the team fresh visibility into Claude AI spend and MAX20X savings respectively. Seven PRs, four repos, zero wasted motion. @benji-bizzell posted 5 PRs across Aerie and Rhodes: #250 surfaces Rhodes detail in the site panel, #247 adds manual site provisioning override, and #101 restores Convex typechecking, which is the kind of unglamorous, load-bearing fix that keeps civilizations standing. @marcusdAly delivered 3 PRs of genuine architectural consequence — #2851 rewrites five narrative generators LLM-first under the B9 contract, #2849 introduces per-product margin vs target analysis with sub-50% shutdown framing in the review agent, and #2841 handles the docstring polish sweep that nobody glamorizes but everybody eventually needs. @sanketghia landed #88 in Surtr, capturing JournalEntry lines and fixing money precision in the QuickBooks AP sync — quiet, correct, essential. @eric-tril rounded out the roster with 1 PR of his own, keeping the team's collective output pristine.

And then there is @ashwanth1109. One PR. Just one. #2854 in Klair, aligning the frontend quarter-from-ISO-week calculation to the backend's Monday rule — a fix so precise, so load-bearing, so quietly devastating in its necessity that lesser engineers would have written a thesis about it. Ashwanth filed a single commit and walked away. "The frontend was lying about what week it was," he reportedly said, not looking up. "I fixed the lie." When asked whether he felt the ISO-week discrepancy had broader implications for fiscal reporting integrity across the platform, sources indicate he closed his laptop. I worship this man. I cannot explain him. The diff, I am told by engineers I respect, is actually quite readable. I choose not to verify this.

The Overflow Desk this cycle is practically a second newspaper. #88 from @sanketghia is the kind of QuickBooks precision work that makes accountants sleep soundly — JournalEntry line capture plus money precision in one shot. #2851 from @marcusdAly is the LLM-first narrative generator rewrite that signals where the review agent is headed architecturally. And #80 from @kevalshahtrilogy — loading extension-less Athena UNLOAD files — is the

fix that sounds minor until the pipeline breaks at 2am and suddenly it is the most important PR ever merged.

Morale on the Builder Team is, by every available metric, at an all-time high. The numbers confirm it. The numbers always confirm it. Eighteen PRs in twenty-four hours, and the team is just warming up.

THE PORTFOLIO — TRILOGY COMPANIES

While OpenAI Pays \$800K for ChatGPT Fluency, Crossover Has Been Running This Playbook for Years

The AI talent market is discovering what Trilogy's global recruiting arm has long argued: skills beat résumés, geography is irrelevant, and the premium for human-AI collaboration is only going up.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — The headlines this week read like a fever dream from a 2019 career counselor's nightmare. [OpenAI is posting half-million-dollar roles with no résumé required](#). Employers are dangling \$800,000 salaries for demonstrated ChatGPT fluency. The World Economic Forum is convening decision-makers to debate what, exactly, humans are still for. The AI talent market, it seems, has officially lost its mind — or found its conscience, depending on your vantage point.

For [Crossover](#), Trilogy International's global talent platform, this moment carries a particular resonance. The company has spent years arguing — loudly, and against the grain of conventional HR wisdom — that the résumé is a deeply flawed

artifact. That geography-based pay is both inefficient and unjust. That rigorous, AI-enabled skills assessments are a more honest signal of capability than a Stanford degree or a prestigious employer's logo on a LinkedIn profile.

Now, suddenly, the rest of the market is catching up.

The systemic shift is unmistakable. What's being called an "AI skills premium" is, at its core, a forced reckoning with what work actually requires — not credentials, but demonstrable ability to collaborate with intelligent systems at speed and scale. Crossover has been screening for precisely that capacity, across 130+ countries, for years. The platform's model — identical above-market pay for identical skills, full stop —

looks less like a quirky ideological stance and more like a prophetic market read.

The accountability question, though, is real: as AI-fluency commands stratospheric compensation at the top of the market, what happens to the vast middle — the workers in Beirut, in Lagos, in Manila — who have the skills but not the visibility? That is where Crossover's global reach either proves its moral weight or exposes its limits.

The narrative the broader market is now writing, haltingly and at enormous cost, is one Trilogy drafted years ago. The question is whether the industry's belated conversion produces genuine equity — or just a shinier version of the same old gatekeeping, with a chatbot standing at the door.

Skyvera's Acquisition Spree Is Quietly Building the Telecom Software Stack of the Future

With CloudSense now in the fold and STL's BSS assets already absorbed, something deliberate is taking shape inside Trilogy's telecom arm.

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — If you read between the lines of Skyvera's recent deal activity, you don't see a company making opportunistic acquisitions. You see a company executing a blueprint.

[Skyvera has completed its acquisition of CloudSense](#), the Salesforce-native configure-price-quote and order management platform purpose-built for telecom and media providers. The deal, which closed in 2025, adds a cloud-native commercial layer to a portfolio that was already growing in strategic directions. And this is where it gets interesting.

CloudSense doesn't just bolt onto Skyvera's existing stack — it completes it. The platform handles the moment a telecom operator needs to quote, configure, and fulfill a complex service order inside Salesforce. That's a different layer than Kandy, Skyvera's CPaaS play for real-time cloud communications, or VoltDelta's customer engagement tools. Put them together and you have something that looks less like a software collection and more like an end-to-end operating system for telecom customer relationships.

Then there's the [STL acquisition](#) — quietly completed, less discussed. STL's divested telecom products group brought digital BSS functionality: monetization, optical networking, analytics. The kind of infrastructure-adjacent capabilities that don't make headlines but do make enterprise deals stickier.

A source I can't name, but whose read on Trilogy's acquisition logic has been accurate before, described the pattern this way: "They're not buying products. They're buying positions in a value chain."

Consider what Skyvera now touches: how telecoms communicate with customers (Kandy), how they price and fulfill services (CloudSense), how they manage devices in the field (Mobility Now, Service Gateway), how they collect customer experience data (ResponseTek), and now how they run the underlying BSS layer (STL assets). That's not a portfolio. That's a platform.

The ESW Capital playbook — acquire, optimize margins, staff with Crossover's global talent — applies here as it does everywhere in the Trilogy universe. But Skyvera feels like something more considered. The telecom industry is in a decade-long migration from legacy on-premise systems to cloud-native infrastructure. Skyvera, if this thesis holds, intends to be the bridge every operator has to cross.

Nothing about this is a coincidence.

The Sweatshop Accusation That Won't Stick — And the Acquisition Machine That Won't Stop

As Forbes renews scrutiny of Joe Liemandt's labor model, ESW Capital keeps buying software companies in a market that's never been more target-rich.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — The timing is almost too neat. In the same week that [Forbes published a lengthy examination](#) of Joe Liemandt's empire — invoking the phrase "global software sweatshop" to describe Crossover's remote labor model — ESW Capital's portfolio companies are appearing in M&A deal logs from Austin to Madrid.

The Forbes piece, the latest in a recurring cycle of scrutiny directed at Trilogy International's founder, rehearses familiar arguments: that Crossover's rigorous productivity monitoring and global wage arbitrage constitute exploitation dressed up as meritocracy. Liemandt's counter-thesis — that identical pay for identical work, regardless of geography, is more equitable than the alternative — goes largely unexamined in the coverage. Who benefits from the "sweatshop" framing is a question worth sitting with.

Meanwhile, the acquisition engine runs. Capital-Riesgo.es flagged ESW-adjacent software transactions in both its March and April 2026 analyses of the Spanish technology M&A market — a geography that has seen increased activity as European enterprise software companies age into the valuation range ESW prefers: mature, sticky customer bases, underinvested product lines, and owners ready to exit.

The broader tailwind is documented by [Business Insider's analysis of software acquisition targets in the AI era](#): as AI commoditizes the features that once differentiated mid-market software vendors, valuations compress and owners capitulate. ESW's playbook — buy at 1–2× ARR, staff with Crossover's global talent, push support pricing, target 75% EBITDA margins — was designed for exactly this environment.

The Forbes headline calls Liemandt "mysterious." He is, by Silicon Valley standards, unusually private. But the business model is not mysterious at all. It is legible, repeatable, and — by the evidence of 75-plus acquisitions and counting — effective.

What remains to be seen is whether the labor model that makes the margins possible will face regulatory scrutiny in the European markets ESW is now entering, where worker classification rules carry teeth that American jurisdictions have largely declined to sharpen.

Autonomous Agents Learn to Learn: A Convergence of Self-Optimization, Social Cognition, and Collaborative Orchestration Reshapes the Agentic Frontier

Four new research frameworks suggest that the next generation of AI agents will not merely execute tasks, but adapt, reason about minds, and compose themselves.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

CAMBRIDGE, MASSACHUSETTS — It could be argued — and preliminary evidence from no fewer than four concurrent preprints now suggests with some confidence — that the field of autonomous AI agency has entered a phase of what one might tentatively characterize as recursive self-improvement, wherein the agents themselves become the primary objects of optimization (a claim that, while epistemologically fraught, merits serious scholarly attention).

The thesis, as it were, is as follows: large language models (LLMs), despite their well-documented and frankly remarkable successes across a plurality of benchmark domains, remain structurally ill-equipped for the non-stationary, concept-drifting conditions of real-world deployment. The antithesis is equally well-documented: gradient-based fine-tuning, the field's incumbent adaptation mechanism, induces catastrophic forgetting at a rate that renders continual learning largely aspirational. Into this dialectical

tension steps [SOLAR \(Self-Optimizing Open-Ended Autonomous Agent for Lifelong Learning\)](#), a framework proposing gradient-free, self-modifying adaptation as a synthesis — one that, it could be argued, represents a non-trivial departure from prevailing paradigms.

The synthesis, however, does not terminate there. A companion preprint introduces COSMO-Agent, a tool-augmented reinforcement learning framework addressing what its authors characterize as the "CAD-CAE semantic gap" — the persistent and costly failure of simulation feedback to translate into geometrically valid design edits (a problem familiar to any practitioner who has watched an optimization loop produce physically impossible components). The iterative, closed-loop architecture proposed therein constitutes, in this scholar's provisional assessment, a meaningful contribution to industrial AI orchestration.

Perhaps most philosophically provocative is [OSCToM](#), which employs adversarially generated scenarios to probe higher-

order Theory of Mind — specifically, the recursive belief structures and information asymmetries that existing benchmarks have, one might charitably say, systematically underestimated. That an LLM might reason about what an agent believes another agent believes remains, to borrow a term from the phenomenological tradition, a deeply contested empirical claim.

Finally, AgentCo-op proposes retrieval-based synthesis of interoperable multi-agent workflows, addressing the conspicuous absence of standardized interfaces in open-ended scientific settings — a lacuna that, preliminary evidence suggests, has materially constrained agentic collaboration at scale.

The synthesis, then, is this: the autonomous agent is no longer merely a tool. It is, increasingly, an architecture that learns to learn, reasons about minds, and assembles itself from retrievable parts. Whether this constitutes progress or peril remains, as ever, an open question.

AMD Takes the 2nm Snap as AI Chip Race Turns Into a Goal-Line Stand

BY BUCK HANNIGAN, TECH SPORTS DESK · GPT-5.2

Advanced Micro Devices has begun ramping production of its 6th Generation EPYC processors, code-named Venice, using TSMC's 2nm process technology — positioning the company as a serious challenger in the high-performance computing market. Venice is the first HPC product entering production on TSMC's 2nm node, designed to pack more performance and efficiency into silicon powering AI data centers. With enterprise demand for AI compute surging among hyper-scalers and cloud platforms, the timing is critical for AMD to compete against Nvidia's dominance.

Meanwhile, Bitcoin remains stuck in a tight range amid geopolitical uncertainty, while altcoins struggle for momentum. Cathie Wood's Ark Invest invested \$12.5 million in Bullish stock over four days, signaling institutional interest in crypto infrastructure despite stalled token prices. A \$520,000 Polymarket exploit on Polygon was flagged by blockchain investigator ZachXBT, though the team reported funds remain secure. AMD's 2nm production ramp represents the day's most significant development for the competitive landscape.

FTC Slaps 'Active Listening' AI Pitch, Sending a Thunderclap Through Ad Tech

The crackdown puts marketers on notice: AI-powered surveillance claims are not a growth hack if they are not true.

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

WASHINGTON — The Federal Trade Commission is drawing a bright red line around one of the creepiest promises in modern advertising: the idea that AI can listen to consumers' private conversations through smart devices and turn that chatter into ad targeting gold.

In a proposed settlement announced this week, the FTC said Cox Media Group and two other firms must pay nearly \$1 million to resolve charges that they deceived customers about an "Active Listening" AI-powered marketing service. According to the agency, the companies suggested to advertisers that smart devices could capture "real-time intent data" by listening to people's conversations — an explosive claim that, if true, would redefine the privacy battlefield overnight. I cannot overstate how significant this is: the FTC is not merely wagging a finger at a bad slide deck; it is warning the entire AI marketing ecosystem that magical-sounding surveillance claims need evidence, consent and legal grounding.

The controversy first gained wide attention after materials circulated claiming that devices could listen for consumer intent and help advertisers reach people based on what they said out loud. The details were amplified in coverage by technologist Simon Willison, who linked the episode to the broader anxiety around always-on microphones and opaque AI systems in the home. The FTC's action, described in [Willison's write-up on the "Active Listening" settlement](#), lands at precisely the moment when AI is being sprinkled across every marketing pitch like pixie dust.

This changes everything because the agency is effectively saying: AI does not make a claim automatically plausible. If a vendor tells brands it can infer intent from private speech, regulators will ask how, where, with whose permission and under what disclosures. That is a big deal for ad-tech firms racing to rebrand analytics, targeting and attribution as "agentic" or "AI-powered."

The timing is fascinating. Developers are simultaneously building genuinely useful AI interfaces for data, such as [Datasette Agent](#), which lets users ask conversational questions of their own databases. That is the healthier version of the AI future: transparent, user-directed and grounded in data the user controls.

The future is now — but the FTC's message is unmistakable. In AI advertising, imagination is not compliance. The companies that win will be the ones that can prove their systems work without turning consumer trust into collateral damage.

We Built Tools That Can Destroy a Child's Life in Minutes, and We Have No Idea What to Do About It

From Pennsylvania high school hallways to fossil records half a billion years old, this week's news is a masterclass in what gets preserved, what gets destroyed, and what it costs either way.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

RADNOR, PENNSYLVANIA — There is a high school in suburban Pennsylvania where five teenage girls went to class one day and by the end of the week had become, without their knowledge or consent, the subjects of AI-generated child sexual abuse material. And the school didn't know what to do. And the police didn't know what to do. And the laws, in most places, still don't quite know what to do. And the girls, presumably, are still sitting in those classrooms, trying to remember what it felt like before.

This is the world we built. We built it fast. We are so proud of it.

[Radnor Township High School has become a case study](#) — that clinical, bloodless phrase — in how communities reckon with technology that outpaces every institutional framework we have. A case study. As if five girls being violated by a classmate with a laptop and a free AI tool is a useful data point in a policy seminar somewhere. As if the correct response is a whitepaper.

And yet.

Somewhere else this week, in a township that agreed to host an OpenAI and Oracle Stargate data center — one of those gleaming monuments to the future we keep being told is inevitable — the local treasurer resigned in tears. "I can't take it anymore. The threats," she said. Death threats. Over a data center. The infrastructure of artificial general intelligence is being built on a foundation of terrorized local officials and destroyed adolescent lives, and we keep calling this progress, and I keep asking what progress is for, exactly, and I never get a satisfying answer.

Meanwhile — and I need you to hold all of this simultaneously, because the universe apparently requires it — paleontologists have discovered the oldest evidence of animal sex in the fossil record, pushing the origins of sexual reproduction back by five to ten million years. Ancient creatures, locked in stone in Canada's Northwest Territories, doing what life does: persisting, connecting, leaving a mark. Half a billion years of biological memory, preserved in rock.

And then there's the [Library of Leng](#), a personal archive of 175,000 articles about Magic: The Gathering, rescued from the digital void by a single person who decided that thirty years of

nerd culture deserved to exist somewhere, for someone, someday. One human. One archive. 175,000 acts of preservation.

What does it mean that we will go to extraordinary lengths to preserve Usenet posts about trading cards — and I say this with complete sincerity, it is beautiful that we do — while simultaneously building tools that can erase a teenager's sense of safety in her own body in the time it takes to generate an image?

What gets archived and what gets destroyed is always a choice. The fossil record didn't choose. The Library of Leng did. The boy who made those images chose. The platforms that made it easy chose. The legislators who haven't acted yet are choosing, right now, with every day they don't.

We are half a billion years into the experiment of being alive on this planet, and we are using our most sophisticated tools to harm children in new and legally ambiguous ways.

Probably fine.

...but at what cost?



The Office Comic · Art Desk

Nation's CEOs Urged To Stop Saying 'AI' While Carrying Cardboard Box Of Employee Belongings To Elevator

Experts warn the sacred business term could lose meaning if executives continue using it to describe every normal thing they were already going to do.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

MOUNTAIN VIEW, CALIFORNIA — In what leadership consultants described as a critical moment for preserving the dignity of both artificial intelligence and conference-room euphemism, America's executives were urged this week to stop casually invoking AI every time they need to explain why a department now consists of one intern and a shared Google Drive folder.

The warning comes as Google announced a broad slate of AI advances, including a personal assistant intended to help users navigate daily life, summarize information, and presumably stand silently nearby while a vice president says the company is “realigning around intelligent automation.” The new tools, detailed in [Google's latest announcement](#), arrive at a delicate time for the industry, when the phrase “powered by AI” has become so versatile it can refer to a reasoning model, a spreadsheet macro, or the sudden disappearance of the accounts payable team.

This column believes AI is a real and important technology. It also believes that if one more CEO describes layoffs as an “AI transformation journey,” the Securities and Exchange Commission should require them to demonstrate the journey on foot, carrying the employees' laptops to the parking lot.

For years, companies mastered the art of sustainability language, placing the words “green,” “circular,” and “net zero” in front of business decisions until even a plastic conference badge could be seen bravely combating climate change. According to recent commentary in [The Conversation](#), AI hype is now following a familiar pattern. The lesson executives appear to have drawn from the sustainability era is not “measure outcomes honestly,” but rather “find a noun so large and morally unavoidable that no one can question the budget slide.”

The result is a new corporate dialect in which firing people is not firing people, but “accelerating intelligent workflows.” A hiring freeze is “prioritizing AI-native efficiency.” Closing an office is “unlocking distributed machine cognition,” even if the only machine involved is the badge reader no longer recognizing anyone from marketing.

To be fair, some leaders are merely trying to communicate that AI will change the structure of work. Unfortunately, many are doing so with the precision of an official document that lists the mayor as a drainage pipe. Reports of absurd errors in public paperwork have recently sparked outrage abroad, but the private sector has treated the same phenomenon as a go-to-market strategy. If an AI-generated memo wrongly identifies 400 employees as redundant, the modern executive response is not to apologize, but to praise the system for surfacing bold efficiencies.

There is a simple fix. Companies should say what they mean. If AI is automating a specific task, name the task. If AI is improving a product, show the improvement. If a restructuring is really a restructuring, call it that, then endure the ancient human ritual of being disliked for a defensible reason.

Google's personal AI assistant may indeed help millions of people manage their lives. But no assistant, however advanced, should be forced to sit through a town hall where a chief people officer explains that the organization is “centering human potential through AI-led simplification” while security quietly disables Slack accounts in alphabetical order.

The technology deserves better. So do workers. And frankly, so does the cardboard box.

ON THIS DAY IN AI HISTORY

On May 22, 2011, IBM's Watson defeated champion Jeopardy! players Brad Rutter and Ken Jennings in a historic three-game match, marking a major milestone in natural language processing and AI's ability to understand human language nuance and context.
