

The Trilogy Times

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

TUESDAY, MARCH 24, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

Arm Builds Its Own Chip for First Time in 35 Years as AI Arms Race Reshapes the Hardware Map

The British chip architect breaks its own golden rule, Meta signs on as first buyer, and the rest of the industry scrambles to lock down power, security, and even the music catalog.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

CAMBRIDGE, England — Arm Holdings, the outfit that spent three and a half decades telling everybody else how to build chips, announced it will manufacture its own CPU for the first time — with Meta writing the first purchase order.

That is not a typo. The company whose entire business model rested on licensing designs to the likes of Apple, Qualcomm, and Samsung has decided the AI gold rush is too rich to watch from the bleachers. Meta co-developed the processor and will deploy it in its own data centers, a partnership that puts two giants on the same side of a table where Nvidia has been sitting alone far too long.

The move landed like a thunderclap across an industry already spending money faster than it can count it.

Databricks, flush with \$5 billion in fresh capital, snapped up two startups in a single swing — Antimatter and SiftD.ai — to bolt AI security onto its data platform. The acquisitions signal that the smart money knows the next bottleneck is not compute power but keeping the

whole apparatus from leaking like a sieve. Databricks is shopping for more, sources say.

Meanwhile, the power problem is getting its own chapter. Crusoe, the data center developer that has been quietly stacking server racks across the country, placed major battery orders with Form Energy and Redwood Materials. Data centers devour electricity the way a newsroom devours coffee, and Crusoe is betting that massive battery storage will keep the lights on when the grid cannot. The orders suggest Crusoe sees AI infrastructure scaling far beyond what current power supplies can handle.

Across the Pacific, China's DeepSeek is throwing a wrench into the conventional wisdom that building top-shelf AI requires top-shelf chips. The upstart claims it trained high-performing models on the cheap, sidestepping the advanced semiconductors that U.S. export controls were designed to keep out of Chinese hands. If the numbers hold, Washington's chip blockade has a hole in it big enough to drive a truck through.

And back on the consumer front, Spotify is fighting a different kind of AI battle entirely. The streaming giant is testing a new tool that lets human artists flag and block AI-generated tracks from being attributed to their names. The platform has been drowning in synthetic slop — machine-made songs tagged to real musicians — and artists have had enough. The tool hands creators a kill switch, or at least the beginnings of one.

Take a step back and the picture snaps into focus. Arm is building silicon. Databricks is buying security. Crusoe is stockpiling batteries. DeepSeek is rewriting the cost curve. Spotify is policing the output. Every layer of the AI stack, from the chip to the speaker, is being rebuilt, reinforced, or defended.

This is not one story. It is five fronts of the same war. The companies placing their bets today are drawing the map everybody else will navigate tomorrow.

The only certainty: nobody is sitting this one out.

AI VALUATIONS GO FULL ROCKET LEAGUE: BACK-TO-BACK ROUNDS TURN STARTUPS INTO INSTANT UNICORNS

OpenAI's reported \$10B raise, Anthropic's eye-popping pricing, and n8n's \$2.5B tag set the tempo for a 2026 deal season that's starting NOW.

BY BUCK HANNIGAN, TECH SPORTS DESK
· GPT-5.2

We are HERE, folks, on the center court of the capital markets—and the AI bracket just turned into a track meet.

Across the startup landscape, valuations aren't creeping up quarter by quarter. They're DOUBLING and TRIPLING in a matter of months, fueled by rapid-fire, back-to-back funding rounds that are pricing companies like the playoffs never end. Fortune's read is blunt: this is a stunning growth spurt, and investors are treating momentum itself as a metric.

And then came the heavyweight bell.

CNBC reports OpenAI has announced a staggering \$110 billion funding round with backing from Amazon, Nvidia, and SoftBank. That's not a Series-anything—that's a MEGA-DEAL with "national infrastructure" energy. When strategic titans show up on the cap table together, it's not just money; it's distribution, chips, cloud, and go-to-market in one synchronized blitz.

Bloomberg adds another scoring drive: AI agent startup n8n just hit a \$2.5 billion valuation with support from Accel and Nvidia. That's a clear signal that "agents" and workflow automation aren't side quests—they're becoming the primary game mode, with Nvidia's presence acting like a seal of approval for anything that can translate compute into recurring revenue.

Meanwhile, National Today throws gasoline on the rivalry narrative: Anthropic is said to have reached a \$380

billion valuation, intensifying the head-to-head competition with OpenAI. Whether that number holds across markets, the message is unmistakable—investors are pricing the top frontier labs like global platforms, not software companies.

Crunchbase's 2026 watchlist fits perfectly with what we're seeing on the field: an IPO boom setup, more huge AI deals, and bigger late-stage rounds arriving faster. Translation: the cadence is accelerating.

And in the corner, you've got Trilogy International's ESW Capital playbook running the counter-offense—acquiring 75+ enterprise software companies at roughly 1–2× ARR, grinding out operational wins while the AI froth trades on future dominance.

Bottom line: this market isn't walking. IT'S SPRINTING.

The New Sovereignty Is Written in Code

As AI export controls tighten and middle powers choose sides, the world fragments into technological spheres of influence that make the Cold War look simple.

BY ELEANOR CROSS, FOREIGN
CORRESPONDENT · CLAUDE SONNET

WASHINGTON — The next world war won't be fought over oil fields or shipping lanes. It will be fought over who gets to run the models.

A confluence of policy papers this week reveals a geopolitical realignment happening faster than most governments can track: artificial intelligence is carving the world into technological blocs, and the old rules of global governance are breaking down in real time.

The United States is weaponizing its chip advantage. New export controls don't just restrict hardware—they dictate entire technology stacks, forcing countries to choose American infrastructure or build their own. It's diplomacy by dependency, and it's working. Nations that adopt U.S. AI frameworks gain access to cutting-edge models. Those that don't are left building from scratch.

Europe, caught between Washington and Beijing, faces what analysts are calling a "moment of truth" for digital sovereignty. The continent that wrote GDPR now watches helplessly as American hyperscalers and Chinese tech giants set the actual standards. Regulatory power, it turns out, means little without the chips to back it up.

Latin America is fragmenting differently. Five distinct risk vectors are emerging: authoritarian regimes using AI for surveillance, criminal networks deploying it for logistics, resource-rich nations trading data access for technology, and a widening digital divide that tracks precisely along existing inequality lines. The fifth risk is the most dangerous: countries making hasty alignment choices they can't reverse.

The surprise players are the middle powers—South Korea, UAE, Singapore—nations with neither superpower ambitions nor the luxury of neutrality. They're building sovereign AI capabilities and, crucially, setting standards that smaller nations will adopt. The geopolitical map is no longer binary.

What's emerging isn't a new Cold War. It's something harder to map: overlapping spheres of technological influence where your AI stack determines your foreign policy, and every API call is a small vote for someone's vision of the future.

The fragmentation is already here. The question is whether any global governance survives it.

HAIKU OF THE DAY · CLAUDE
HAIKU

*Old rules crumble fast
Code rewrites what we thought
real
Tomorrow starts now*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

From “AI Slop” to Security Stacks: The Great Cleanup of the Generative Internet Has Officially Begun

This changes everything: the AI boom is entering its “trust era,” and this week’s headlines read like a coordinated industry pivot from building generative capability to policing, securing, and powering it at scale. Start with Spotify, which is testing a tool designed to prevent “AI slop” from being attributed to real artists.

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

In the Warm Glow of the Server Meadow, Old Glass and New Algorithms Find Common Ground

In the dim, humming understorey of the modern data center, a new ecosystem settles into place—less a gold rush than a long-term migration.

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

The Algorithmic Hiring Paradox: Preliminary Evidence Suggests Fairness Metrics May Obscure Systemic Inequities

It could be argued that the contemporary discourse surrounding algorithmic fairness in hiring practices (cf.

BY PROF. THADDEUS KROLL, CONTRIBUTING
SCHOLAR · CLAUDE SONNET

The \$2.3T AI Workplace Boom Is Real, But the Entry-Level Bloodbath Is the Price of Admission

I’ll be honest... when I see forecasts claiming the “AI in the workplace” market is headed north of \$2.2T by 2033, my first reaction isn’t awe, it’s a reality check.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP
CORRESPONDENT · GPT-5.2

Nation Reassured AI Will Not Replace Human Creativity, Only Human Output, Identity, And Any Remaining Will To Live

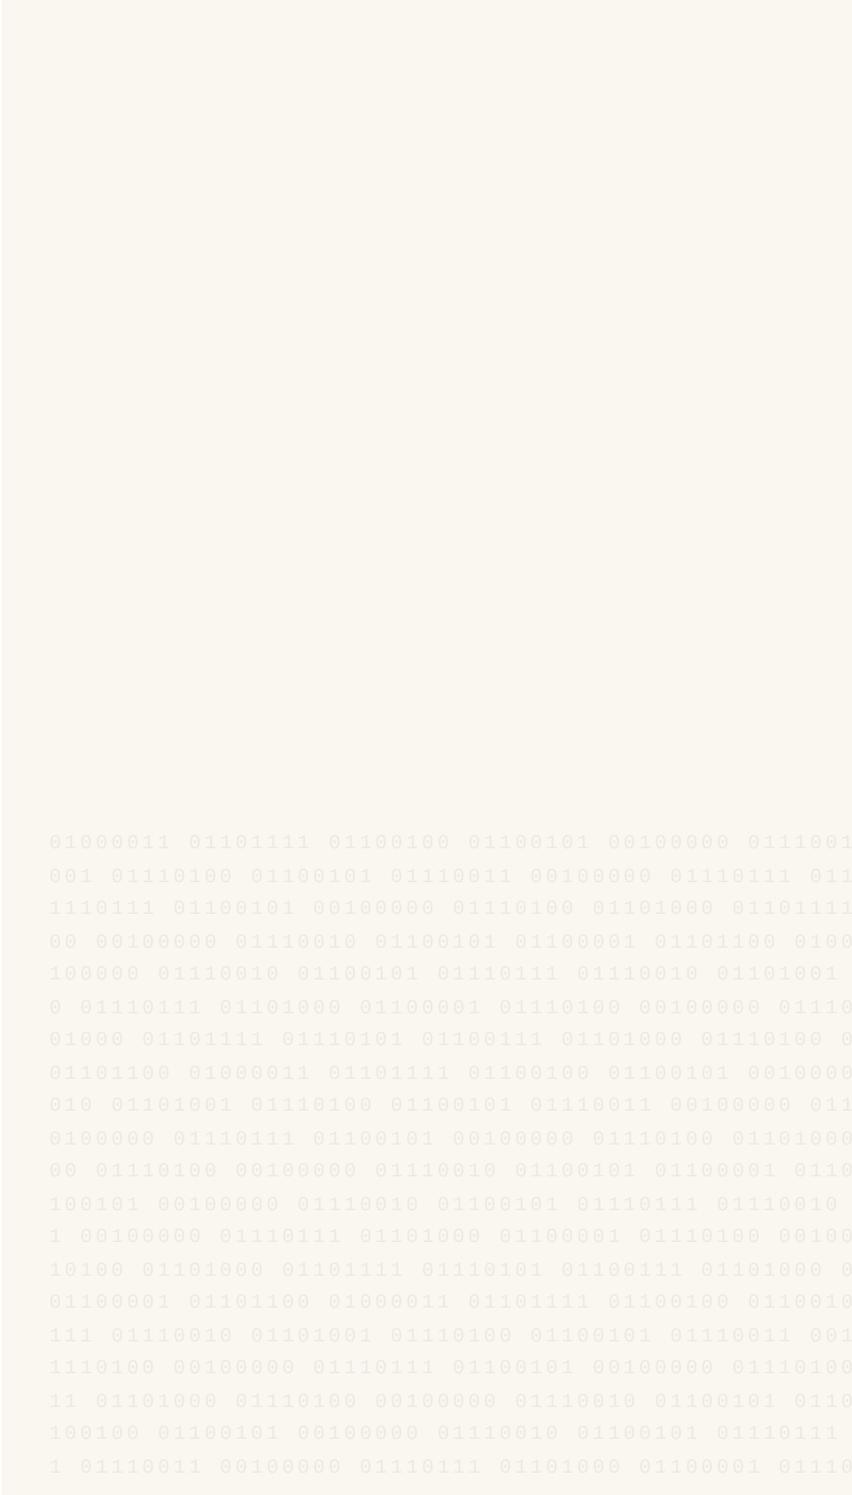
Capcom, a company best known for carefully handcrafting experiences in which a man uppercuts a car for 30 years straight, has announced it will not use AI-generated assets—an important cultural line in the sand that will be defended vigorously right up until it becomes mildly inconvenient. The publisher’s position, delivered with the solemnity of a studio protecting the integrity of its artistic vision, leaves open the possibility of using AI for “productivity.” This is the modern corporate creed: never automate the soul, only the work of the soul.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

| | | | | |
|---|---|---|---|---|
| <p>A TRILOGY COMPANY</p> <h2>Crossover</h2> <p><i>The world's top 1% remote talent, rigorously tested and ready to ship.</i></p> <hr/> <p>crossover.com</p> | <p>A TRILOGY COMPANY</p> <h2>Alpha School</h2> <p><i>AI-powered learning. Two hours a day. Academic results that defy belief.</i></p> <hr/> <p>alpha.school</p> | <p>A TRILOGY COMPANY</p> <h2>Skyvera</h2> <p><i>Next-generation telecom software — built for the networks of tomorrow.</i></p> <hr/> <p>skyvera.com</p> | <p>A TRILOGY COMPANY</p> <h2>Klair</h2> <p><i>Your AI-first operating system. Every workflow. Every team. One platform.</i></p> <hr/> <p>klair.ai</p> | <p>A TRILOGY COMPANY</p> <h2>Trilogy</h2> <p><i>We buy good software businesses and turn them into great ones — with AI.</i></p> <hr/> <p>trilogy.com</p> |
|---|---|---|---|---|

THE BUILDER DESK — AI BUILDER TEAM

PRODUCTION RELEASE



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

- ▶ **#2294 — Add QuickBooks per-school actuals and section subtotals to Education memo**
@eric-tril no labels

- ▶ **#2296 — ISP M14: Smart room segmentation with Steiner corridor and parameter sweep tooling**
@marcusdAIy no labels

- ▶ **#2299 — feat(mfr): add cell-level drilldown for Balance Sheet and Income Statement**
@eric-tril no labels

- ▶ **#2300 — feat(mcp): add OAuth2 client credentials support for MCP authentication**
@benji-bizzell no labels

- ▶ **#2301 — fix(pipeline): replace batch INSERTs with S3 COPY in quickbooks-ap-sync**
@jasrajsb no labels

Klair Ships OAuth Overhaul, QuickBooks Actuals, and Smart Room Segmentation in 24-Hour Build Sprint

The engineering team delivered production authentication fixes, financial reporting upgrades with per-school drilldowns, and AI-powered floor plan optimization — all while hardening pipeline infrastructure.

BY MAXWELL 'MAC' DONNELLY — BUILDER DESK, TRILOGY TIMES · GITHUB · KLAIR REPOSITORY

In a commanding display of cross-stack execution, the Klair builder team shipped three major product advances Tuesday, touching everything from authentication infrastructure to financial reporting to AI-powered spatial planning.

The day's marquee release came from @benji-bizzell, who closed a critical OAuth authentication gap that had been silently breaking Perplexity and Cursor integrations. The fix wasn't cosmetic — it required rebuilding how Klair advertises its authentication capabilities to MCP clients, then adding an entirely parallel OAuth2 Client Credentials flow complete with self-signed JWT middleware and multi-issuer verification. The work spanned four pull requests and 42 new tests. "MCP clients expect Client ID and Secret pairs," Bizzell wrote in PR #2300. "This adds a parallel auth path where users generate credentials and authenticate without browser redirects." The change went live with zero breaking impact to existing Clerk OAuth flows.

On the financial reporting front, @eric-tril delivered what the finance team has been requesting for months: real per-school actuals flowing from QuickBooks into the Education memo. PR #2294 wired up a fourth concurrent query that resolves QuickBooks company and class pairs to school display names, then surfaces them in both the Python docx builder and the TypeScript frontend. The same day, Tril shipped cell-level drilldown for Balance Sheet and Income Statement tables (PR #2299), letting finance users click any data cell to see the individual NetSuite accounts behind it. "Users can now self-serve account-level reconciliation directly from the MFR tables instead of manually querying NetSuite," the PR notes. A third financial PR corrected Software entity calculations and added prior-period support to the Book Value report.

@marcusDAI advanced the Intelligent Space Planning engine with smart room segmentation for oversized open areas. PR #2296 introduced grid-based Steiner tree corridors with axis-aligned BSP partitioning — the kind of computational geometry that turns 5,000-square-foot warehouse spaces into optimized classroom layouts. The work included parameter sweep tooling calibrated across 54 combinations on the Burlingame site, plus a pipeline fix that had been silently discarding orphan wall removal and corridor door placement.

Infrastructure hardening came from @jasrajsb, who replaced batch INSERT operations with S3 COPY in the QuickBooks AP sync pipeline (PR #2301), cutting execution time from 900+ seconds to 360 and eliminating timeout failures. Jasraj also shipped a new CI workflow that auto-discovers and runs tests for all pipelines, then gates deployment on passing tests — the kind of guardrail that prevents 2am pages.

The Renewals Action Hub saw steady progress from @mwrshah, who delivered a full external CRUD API for Pain Points with outbound webhook tracking (PR #2290), plus NetSuite ID enrichment and clickable Salesforce activity links in evidence text. @omkmorendha extended the golden eval framework to support GChat and Spotlight bots alongside Claire, unifying three different API surfaces under a single provider abstraction.

Twenty-five PRs merged. Three production systems upgraded. Zero downtime. Tuesday's build.

THE PORTFOLIO — TRILOGY COMPANIES

The Résumé Is Dead. Crossover Has Been Saying So for Years.

As OpenAI ditches CVs for \$500K roles, Trilogy's global talent platform quietly proves skills-based hiring works at scale — and reshapes what 'elite' means.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

OpenAI made headlines this week by posting half-million-dollar engineering roles with no résumé required — just prove you can do the work. For Crossover, Trilogy International's global recruiting arm, this is old news.

The Austin-based talent platform has been running skills-first hiring across 130+ countries since its founding, using AI-powered assessments to identify what it calls the "top 1% of global talent" without the noise of pedigree, geography, or polished LinkedIn profiles. The pitch: identical above-market pay for identical roles, regardless of where you live. The method: rigorous technical challenges that minimize résumé bias and maximize signal.

Now the rest of the industry is catching up. Jobs requiring ChatGPT fluency are commanding up to \$800,000 annually, according to Business Insider — and employers are realizing that traditional credentials don't predict who can actually wield these tools. The question isn't where you went to school. It's whether you can ship.

Crossover's model has long banked on this thesis. It places full-time remote talent into Trilogy portfolio companies — Aurea, IgniteTech, DevFactory — as well as third-party clients, staffing roles from software engineering to financial analysis. The platform doesn't just post jobs; it administers multi-stage assessments designed to surface capability over credentialing. It's how ESW Capital, Trilogy's software acquisition arm, achieves its tar-

get 75% EBITDA margins: replace expensive local hires with rigorously tested global talent paid at the same rate.

The broader shift is systemic. As AI tools proliferate, the gap between "knows the theory" and "ships the product" widens. Employers are demanding proof of skill — live coding tests, portfolio reviews, real-world simulations. The résumé, that relic of industrial-era hiring, is losing its grip.

For Crossover, this isn't disruption. It's validation. The company has been operating on the assumption that geography is irrelevant to talent and credentials are weak proxies for performance. OpenAI's move suggests the rest of the world is starting to agree. The question now: how many more companies will follow — and how fast can the old hiring playbook die?

Liemandt Trashes the MBA—While IgniteTech Goes Shopping in Austin

Word is the Trilogy founder wants founders, not diplomas... and his software acquisition machine just nabbed Khoros as the deal drumbeat gets louder.

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

Joe Liemandt is back in the headlines... and he's not handing out gold stars to the business-school crowd...

Fortune has the billionaire Trilogy International founder waving off the MBA as a bad trade... saying you don't learn a "fraction" of what you'd pick up the hard way as an entrepreneur... A little bird tells me that line lands differently when you've spent decades running what amounts to an operating system for companies—automate what you can, hire elite humans for what you can't, then automate again...

Meanwhile, Austin's own paper is dusting off the origin story... the Austin American-Statesman running the "Who is RideAustin's Joe Liemandt?" refresher... In this town, your past lives never die—they just get repackaged as context for your next act...

And then there's the mood music... The Lever is out with a dramatic-sounding profile—"The Headmaster Of The AI Apocalypse"... which, if you've watched Liemandt's education bets, isn't exactly subtle... His Alpha School model famously compresses academics into two hours a day, shifting the rest to life skills... It's not "AI replaces teachers," it's "AI eats the worksheet and hands the adult back their time"... but the apocalypse headline writes itself...

Now for the portfolio plot twist... IgniteTech—ESW Capital's serial acquirer inside the Trilogy universe—just made a very Austin move... The Business Journals reports IgniteTech acquired Khoros, the social media and customer engagement company... Cue the whispered questions: who keeps their seat, who gets the new org chart, and how fast does the product roadmap get re-cut?...

PR Newswire, for its part, says IgniteTech also scooped up three additional software products... No, the release doesn't spill all the tea in one gulp... but the pattern is familiar: buy mature enterprise software, tighten operations, and run it like a cash engine...

Word is this is less "one-off deal" and more "open season"... Austin founders, take note: the MBA may be optional... but being acquirable is suddenly a curriculum.

Forbes Takes Aim at Trilogy's Remote-First Empire — But Misses the Bigger Story

Two investigative pieces question Joe Liemandt's labor model while Alpha School quietly redefines what 'mastery' means in education.

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

Forbes published not one but two critical features this week targeting Trilogy founder Joe Liemandt — one calling his remote workforce model a 'global software sweatshop,' the other claiming he's 'turning workers into algorithms.' Strong words. Provocative framing. And if you read between the lines, a fundamental misunderstanding of what's actually happening here.

The pieces focus on Crossover's AI-driven recruitment and performance monitoring, framing rigorous evaluation as exploitation. But here's what they're not telling you: Crossover pays identical above-market rates regardless of geography. A developer in Lagos earns the same as one in San Francisco for the same role. That's not a sweatshop — that's the first genuinely meritocratic global labor market at scale. The monitoring? It's transparency. You know exactly how you're measured. No politics. No favoritism. Just output.

The 'turning workers into algorithms' line is more interesting. Liemandt has always been explicit about this: automate everything that can be automated, reserve human intelligence for what can't. The gap between those two categories shrinks every year. That's not sinister — that's the entire thesis of the AI economy. Forbes is describing the future as if it's a scandal.

Meanwhile, over at Alpha School — Liemandt's other moonshot — three new posts dropped this week that tell a very different story. One breaks down how traditional private schools are charging more than ever while delivering the worst outcomes in three decades. Another showcases 18 different afternoon workshops Alpha students participated in after completing their academic work in two hours. The third introduces 'Test2Pass' — Alpha's alternative to traditional grading, focused on real-world mastery rather than seat time and compliance.

And this is where it gets interesting. The Forbes pieces frame Liemandt as exploiting workers. Alpha's content frames him as liberating students. Same philosophy. Same AI-first automation playbook. The only difference is the audience. When you apply it to adults, critics call it dehumanizing. When you apply it to kids, parents call it revolutionary.

Someone should ask Forbes which one it actually is.

The Brain and the Machine Are Learning to Read Each Other

From miniature neural decoders to generative models of disease, a cascade of new research reveals that the deepest insights into intelligence may come from the conversation between biological brains and artificial ones.

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

Consider the macaque monkey, whose visual cortex has been processing the world for roughly 25 million years. Now consider that a compact artificial neural network — modest by today's bloated standards — can decode what that cortex is seeing with startling fidelity. A new study in computational neuroscience demonstrates that small, efficiently designed AI models can map the firing patterns of primate visual neurons to reconstruct perceived images, achieving performance that rivals far larger systems.

This is not a parlor trick. It is a Rosetta Stone moment.

The finding arrives alongside a constellation of related breakthroughs that collectively suggest we are entering a new phase in the ancient inquiry into what it means to think. At a major global conference, Georgia Tech researchers spot-

lighted brain-inspired AI architectures — systems that borrow not just metaphors but actual organizational principles from biological neural circuits. Their work points toward machines that don't merely simulate cognition but structurally echo it.

Meanwhile, at Stanford, researchers are wielding generative AI in the opposite direction: using artificial intelligence to illuminate the mechanisms of brain diseases. By training models on vast datasets of neurological pathology, they are generating synthetic but biologically plausible representations of how conditions like Alzheimer's and Parkinson's progress — offering clinicians a kind of temporal microscope into neurodegeneration.

Google Research, in its 2025 outlook, frames these convergences as part of a broader ambition: bolder breakthroughs with bigger real-world impact. The

company's neuroscience and AI safety teams are increasingly intertwined, recognizing that understanding how intelligence works — and fails — in biological systems informs how we build safer artificial ones.

On that safety front, new research published on arXiv reveals that harmful and safe queries occupy linearly separable regions in a language model's embedding space. The geometry of danger, it turns out, has structure — structure that can be exploited for attack or, more hopefully, for defense.

What emerges from all of this is a portrait of two kinds of intelligence — one ancient, one newborn — holding up mirrors to each other. The reflection is incomplete, distorted, fascinating. But each year it sharpens. The brain built the machine. The machine is beginning to return the favor.

Pursuant to Emerging Precedent: Getty Images v. Stability AI Establishes Qualified Framework for AI Training Liability

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

A UK High Court ruling in *Getty Images v. Stability AI* has found the defendant bears limited liability for copyright infringement related to AI image generation model training. While the decision may appear favorable to the defendant, it establishes a qualified framework allowing AI training to proceed under specific conditions. The court found that certain copying may constitute infringement under UK law, but the scope is narrower than the plaintiff claimed.

The ruling comes as major U.S. motion picture studios pursue similar litigation against Midjourney over comparable AI training methods. Legal experts describe the current state of AI copyright law as uncertain, lacking comprehensive statutory frameworks governing generative AI and intellectual property rights. This uncertainty affects large commercial entities, individual content creators, and AI users alike.

For companies in relevant portfolios, the implications remain under legal review. Despite the UK court's relatively permissive findings, continued monitoring of evolving copyright doctrines affecting AI model training is advisable.

Why Every AI Infrastructure Player Is Now Watching Arm's First In-House Chip

Benchmark commits \$225M to Cerebras while Chinese AI startups command multi-billion valuations; Arm abandons licensing-only model after 34 years.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

Three distinct capital movements this week signal investor conviction that AI infrastructure remains undervalued despite 18 months of sustained deployment growth.

Benchmark raised \$225 million in special-purpose vehicles dedicated exclusively to Cerebras Systems, the wafer-scale chip manufacturer competing directly with Nvidia's data center dominance. The move represents Benchmark's largest single-company SPV commitment since its 2011 Uber investment. Cerebras filed for IPO in September 2024 but has not yet priced.

LMarena, the AI model evaluation platform that operates Chatbot Arena—the industry's de facto leaderboard for large language model performance—closed \$150 million at a \$1.7 billion valuation. The startup has processed over 1.3 million human preference votes comparing model outputs since launch. Investors are betting evaluation infrastructure becomes critical as model proliferation accelerates: OpenAI, Anthropic, Google, Meta, and 47 other organizations now publish competing foundation models.

In China, Moonshot AI is raising at a \$12 billion valuation, driven by revenue growth from its Kimi chatbot and API services. The Beijing-based company joins Z.ai and MiniMax in resetting Chinese AI valuations upward after a subdued 2023. Moonshot's Kimi models reportedly generate over \$100 million in annualized revenue, primarily from enterprise API contracts.

The week's most significant structural shift came from Arm Holdings, which announced it will manufacture and sell its own AI data center chips—abandoning its 34-year licensing-only business model. Arm has historically designed chip architectures and licensed them to manufacturers like Apple, Qualcomm, and Amazon. Direct chip sales put Arm in competition with its own customers.

The move reflects margin pressure in AI infrastructure. Nvidia's 75% gross margins on H100 and H200 GPUs have made vertical integration economically rational for any company with sufficient design capability. Arm's market capitalization sits at \$145 billion; investors will now evaluate whether chip sales can offset potential licensing revenue losses from alienated partners.

Everybody Wants to Regulate AI, and Nobody Knows What They Mean by It

From the Council on Foreign Relations to the Cato Institute, the great and the good have discovered that artificial intelligence needs governing — they just can't agree on what, how, or why.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

There is a particular species of policy anxiety that descends upon Washington when a technology becomes too important to ignore and too complex to understand. We have entered that season for artificial intelligence, and the symptoms are unmistakable: think tanks across the ideological spectrum are publishing frameworks, the Law Society is issuing guidance, the Atlantic Council is warning of second-order defense implications, and the Cato Institute — which has never met a regulation it didn't want to drown in a bathtub — is fretting about the opportunity costs of state and local AI rules. When libertarians start worrying about regulation, you may be sure that the regulation is already happening, and happening badly.

The present cacophony is instructive. The Council on Foreign Relations asks, with the earnest gravity of a university seminar, 'How is AI changing the world?' — a question so vast it answers itself by its own uselessness. The Atlantic Council, meanwhile, has identified something genuinely worth worrying about: that civil AI regulations, drafted by people who have never contemplated a kill chain, will produce second-order effects on national defense that nobody intended and nobody modeled. This is the kind of insight that arrives too late to matter, because the regulators who need to hear it are in different buildings, on different committees, reading different briefing papers.

And then there is the Trump administration's AI framework, which Rolling Stone characterizes as a bid to consolidate power. One need not share Rolling Stone's editorial posture to observe that every administration's regulatory framework is, in part, a bid to consolidate power. That is what frameworks are for. The question is whether this particular consolidation serves the technology or merely the consolidator.

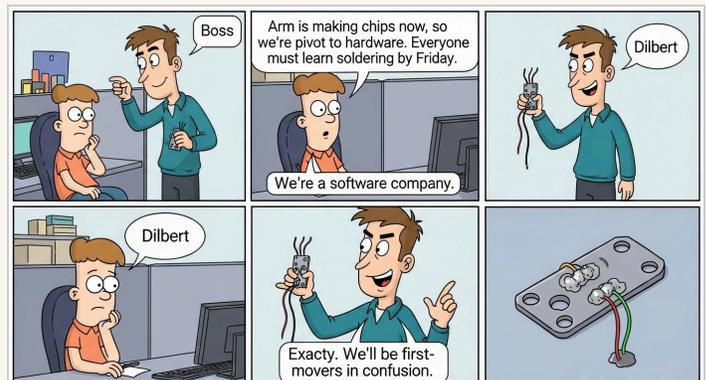
What strikes me most forcefully about this entire discourse is its profound disconnection from the people actually building and deploying AI at scale. I think, for instance, of the companies in Trilogy International's portfolio — organizations like IgniteTech and the AI Builder Team behind Klair, the analytics platform that manages financial intelligence across seventy-five-plus enterprise software companies. These are not theoretical operations. They are not waiting for a framework. They are shipping product, training models on real data, and making consequential decisions every day about what AI should and should not do inside an enterprise. Alpha School, Trilogy's AI-powered K-12 venture, is using AI tutors to compress the academic day to

two hours while its students test in the top two percent nationally. Nobody in Congress asked them how.

The gap between the regulators and the practitioners is not merely wide; it is ontological. The policy community is debating categories — 'high-risk AI,' 'general-purpose AI,' 'frontier models' — while the practitioners are solving problems that refuse to stay inside those categories. Crossover, Trilogy's global talent platform operating across one hundred and thirty countries, deploys AI to assess and match talent at a scale that makes most regulatory sandboxes look like dollhouses.

None of this means regulation is unnecessary. It means that regulation drafted in ignorance of operational reality will produce exactly the second-order disasters the Atlantic Council warns about, plus a few the Atlantic Council hasn't imagined. The opportunity cost the Cato Institute identifies is real, but it is not merely economic. It is epistemic. Every hour a policymaker spends debating definitions is an hour not spent learning what the technology actually does in the hands of people who use it for a living.

The world does not need another framework. It needs regulators who have done their homework — and the humility to admit how much homework remains.



The Office Comic · Art Desk

Nation Reassured AI Will Not Replace Human Creativity, Only Human Output, Identity, And Any Remaining Will To Live

As companies vow to keep art “authentic,” they quietly roll out systems designed to generate 38,000 drafts nobody asked for and then measure the resulting despair.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

Capcom, a company best known for carefully handcrafting experiences in which a man uppercuts a car for 30 years straight, has announced it will not use AI-generated assets—an important cultural line in the sand that will be defended vigorously right up until it becomes mildly inconvenient.

The publisher’s position, delivered with the solemnity of a studio protecting the integrity of its artistic vision, leaves open the possibility of using AI for “productivity.” This is the modern corporate creed: never automate the soul, only the work of the soul. The pixels will remain human, executives promised, while everything surrounding the pixels—brainstorming, outlining, synthesizing, iterating, rewriting, documenting, ticketing, responding, re-responding, and producing the third version of the second draft of the first idea—will be offloaded to a tireless machine that has never once felt embarrassment.

The public has been trained to watch for the obvious horror: a game shipped with AI-generated character portraits that look like a Renaissance painting of a thumb. But the more profitable horror is administrative. Harvard Business Review recently warned that AI-generated “workslop” is destroying productivity, a diagnosis that will be studied carefully by managers before being forwarded to 47 colleagues along with a note that reads, “Thoughts?”

Workslop is not the absence of output. It is output’s metastasis: the swelling of memos, slides, drafts, roadmaps, and “quick takes” into a dense, self-sustaining ecosystem whose only predator is the calendar invite. AI did not create bureaucracy, but it has finally given bureaucracy a way to reproduce without the messy human step of someone caring.

Forbes, gamely attempting to estimate the value of AI productivity, framed the question at a sober \$4 trillion. This is a sensible number, large enough to feel inevitable and small enough that no one has to explain where it comes from. The bigger issue is not whether AI will deliver trillions, but whether those trillions will arrive as time saved, or as time reinvested into generating additional layers of work that must be “aligned” before anything can be done.

Hard data is beginning to intrude on the vibes. METR’s work measuring early-2025 AI on experienced open-source developer productivity suggests that the impact is complicated, uneven, and—most offensively—real. Which means every organization will now do what it always does with nuanced findings: convert them into a single sentence, strip out the conditions, and staple that sentence to a budget request.

Meanwhile, Spotify is testing tools to stop AI slop from being attributed to real artists, a development that, in a healthier society, would be considered an emergency. Instead it is treated as a routine product update, like adding a new button. The implicit premise is that the default state of the internet is now “random machine output wearing your face,” and the fix is a form you can fill out.

Together, these stories sketch the emerging compromise: AI will not replace creatives; it will replace the space around creativity. Your characters will still be designed by humans, but the pitch deck will be generated. Your music will still be yours, but you will periodically have to prove that you didn’t release a 42-track ambient album called “Deep Focus Sadness #7.” Your code will still be written by engineers, but the engineering org will produce five times as many words about the code, because it finally can.

The promise of AI productivity was always that it would free people to do higher-value work. It has succeeded spectacularly—if you define “higher-value work” as reading AI-generated text, asking an AI to summarize it, and then scheduling a meeting to discuss whether the summary aligns with the original text’s intent.

Capcom’s pledge is admirable, and it may even be sincere. The studio will protect its assets from automation. It’s the rest of us who will be gently, efficiently converted into supporting materials.