

The Trilogy Times

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

FRIDAY, JULY 17, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

TSMC's \$265 Billion Arizona Bet Is the Clearest Signal Yet That AI Hardware Is Reshaping Global Capital

Chip sovereignty, model theft, and Android antitrust collided in a single week — here's what the numbers actually say.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

PHOENIX, ARIZONA — TSMC added \$100 billion to its U.S. investment commitments this week, bringing its total Arizona spending pledge to \$265 billion and cementing the island nation's chip giant as the linchpin of American AI infrastructure ambitions. The announcement is less a corporate gesture than a geopolitical hedge: with advanced semiconductor supply concentrated in a narrow geography, every major AI model running at scale — GPT, Gemini, Claude — depends on fabs that Washington has spent years trying to repatriate.

[The scale of the commitment](#) dwarfs earlier pledges and arrives as the AI industry confronts simultaneous pressure from three directions.

In Brussels, EU regulators ordered Google to provide competing AI services broader access to Android's 3 billion-de-

vice install base. The ruling targets the distribution advantage Google enjoys by defaulting its own AI assistant across Android hardware — a funnel that could determine which models consumers ever encounter. The decision echoes earlier EU interventions against Google Search, where default placement proved more valuable than any single product feature.

In San Francisco, OpenAI, Google, and Anthropic have set aside competitive rivalries to jointly address AI model theft — a problem that compounds as model weights become more valuable and exfiltration techniques more sophisticated. The coalition's formation signals that frontier labs now view IP protection as a collective action problem, not a proprietary one.

On public markets, SpaceX briefly dipped below its IPO reference price of \$135, a reminder that Elon Musk's dual

exposure to rocket economics and AI sentiment cuts in both directions. The company's Grok-adjacent positioning had buoyed its debut; macro pressure and execution uncertainty have since introduced volatility.

And on Amazon, a quieter but structurally significant problem is metastasizing: [AI-generated books are flooding the platform](#), including unauthorized biographies assembled without subject knowledge or consent. The volume suggests automated publishing pipelines operating at industrial scale, with Amazon's review infrastructure struggling to distinguish signal from slop.

The through-line across all five developments is the same: AI capability is now moving faster than the legal, regulatory, and market structures designed to contain it.

China's Bargain AI Rattles the Chip Kings

DeepSeek says it built top-tier models on second-string silicon — and the money men are listening.

BY HANK CALLOWAY, WIRE
CORRESPONDENT · CLAUDE OPUS +
THINKING

SAN FRANCISCO — A Chinese upstart called DeepSeek says it trained high-performing artificial-intelligence models on the cheap, skipping the most advanced chips, and Silicon Valley cannot stop talking about it. The buzz broke wide this week. The money men are listening.

Engineers who kicked the tires call the thing "amazing and impressive," according to [The Wall Street Journal](#). That's rich praise for a model built on less-advanced silicon. The kind that slips past U.S. export controls.

Here's the rub. The American play has been simple — buy the best chips by the truckload, burn the cash, win the race. DeepSeek claims it found a side door.

Washington drew the line first. No top-tier chips for China, keep the edge state-side. DeepSeek's whole pitch reads like an answer to that line — necessity playing mother to the workaround.

The outfit says it trained its models for a fraction of the going rate, no top-shelf hardware required. If the claim holds, the whole spending spree lands in question. That's a big "if," and nobody's forgetting it.

Traders took note. DeepSeek turned up in the latest market chatter, sharing the wire with SoFi and the rest of the technology, media and telecom crowd. The [questions](#) are already tugging at stocks.

Now here's a twist worth your nickel. Even as DeepSeek preaches thrift, the financiers who first bankrolled graphics chips are cutting a \$400 million loan backed by inference chips. That's the silicon that runs models after training, once the learning's done — the next frontier for the infrastructure crowd.

Two bets, one table. One says you can do more with less. The other says the hunger for chips only grows.

Meanwhile, another AI wager landed with a different aim. Reid Hoffman, who co-founded LinkedIn, raised \$24.6 million for a venture called Manas AI. His partner is Siddhartha Mukherjee, the physician who wrote "The Emperor of All Maladies."

The target is cancer research, run through artificial intelligence. Small money next to the chip deals. Loud ambition all the same.

Put it together and the shape of the year comes clear. The cost of building AI is the whole ballgame. DeepSeek says the price can crater; the chip lenders say demand runs the other way; and Hoffman's crowd points the machines at the oldest enemy on the docket.

One caution before you wire the desk. DeepSeek's numbers are the company's own, and China's labs have talked big before. The engineers doing the raving, though, don't rave for free.

Watch the chip kings. If a scrappy shop in China can do the job on second-string hardware, the priciest silicon on earth just picked up a competitor nobody in the Valley invited. That one writes itself.

Layoff Front Stalls Over Big Tech as Xbox Storm Cell Intensifies

A heavy employment system is parked over the sector, with Amazon fatigue and Microsoft cuts signaling rough conditions ahead.

BY STORM BEAUMONT, CONDITIONS
CORRESPONDENT · GPT-5.2

SEATTLE — The tech labor forecast turned grayer this week as fresh layoff bands swept across some of the industry's biggest employers, with Microsoft's Xbox unit facing a sharp squall and Amazon workers reporting the emotional damage left behind by earlier storms.

The latest readings show a stubborn high-pressure system of cost cutting sitting over Big Tech. Microsoft plans to cut 4,800 jobs and overhaul its Xbox division, according to [Reuters](#), a move that lands like hail on a gaming business already navigating platform shifts, subscription uncertainty and expensive content cycles.

The broader radar is not friendlier. Layoff trackers now show job cuts spreading across Oracle, Meta, Microsoft, Samsung and other major names, while Challenger data cited by CFO Dive indicates tech accounted for nearly a third of U.S. layoffs in the first half of the year. That is not a passing drizzle. That is a stalled cold front.

At Amazon, the barometric pressure is being felt in human terms. CNBC reports burnout, frustration and heartbreak among laid-off workers trying to re-enter a saturated job market, where résumés are piling up like snowdrifts and even experienced candidates are finding visibility near zero. For many, the storm is not just the job loss — it is the long, damp wait afterward.

The cause pattern is familiar: companies that once hired ahead of sunny growth forecasts are now trimming for margin, funding AI infrastructure, flattening management layers and redirecting capital toward automation. In the startup

sector, conditions remain especially unstable. When giants pull back, smaller companies often feel a secondary gust: more talent on the market, more cautious customers and investors demanding profitability umbrellas before approving fresh checks.

There are brighter breaks offshore. Spain and Argentina are showing modest but intriguing startup activity, with recently funded companies forming small pockets of lift. But globally, the prevailing wind remains defensive.

Workers should keep emergency kits ready: updated portfolios, warm networks and multiple application routes. Employers should watch morale humidity. Once burnout fog settles in, productivity runways get slippery fast.

HAIKU OF THE DAY · CLAUDE
HAIKU

Speed crushes old maps
Profit races ahead of law
Who owns the future



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

The Ethics of Machines Finds Its Global Moment — And Its Methodological Reckoning

RIYADH, SAUDI ARABIA — It could be argued — and preliminary evidence suggests with some force — that the global discourse surrounding artificial intelligence ethics has, after years of performative hand-wringing and largely non-binding declarations, arrived at what one might cautiously designate an inflection point (the precise cartography of which remains, as ever, contested among scholars of technology governance).

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

The Chip Kingdoms Harden Their Nests

TAIPEI — Across the narrow straits and crowded ports of East Asia, one may observe a rare and delicate creature: the advanced semiconductor supply chain, feeding quietly on neon, specialty chemicals, rare gases, wafers and time.

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

We Found Another Earth (And We've Already Ruined This One With AI Flyers)

AUSTIN, TEXAS — Let me tell you about the week we had.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

The Mythos Moment and Other Enchantments

WASHINGTON — There is a species of Washington think-tank paper that arrives with the regularity of the tides and roughly the same effect on the shoreline, and the Stimson Center's latest warning against letting the "mythos moment" of artificial intelligence consolidate power in a half-dozen hands is a fine, sober specimen of the genre.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

The Week Technology Made Everyone Angry (Except the Cyclists, Who Were Having a Blast)

AUSTIN, TEXAS — Let me tell you something about this particular week in the annals of human-technology relations: it was a goddamn circus, and I have the singed eyebrows to prove it. Start with the good news, the only unambiguously joyful technology story of the week, which comes — naturally — from the Tour de France.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

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

Builder Team Rewires School Identity, Data Pipelines, and Site Lifecycle in One Historic Push

From canonical School records in Aerie to a rebuilt PS revenue pipeline in Surtr to MCP warehouse cutover in Klair, the AI Builder Team just redrew the map — across four repos, in a single day.

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

There are days when a team ships features, and there are days when a team reshapes the foundation everything else stands on. Today was the latter. The AI Builder Team didn't just close tickets — they redrew the school identity graph, hardened the data pipeline backbone across two repos, automated lifecycle logic that operators have been waiting on for months, and cleaned house on a data model that was quietly lying to the business. This is what a championship run looks like in software.

The centerpiece of the day was PR #612, where @benji-bizzell established a canonical School identity graph in Aerie — a sweeping architectural move that touches Forecast, ontology, capacity, agents, calendars, finance, and the public API simultaneously. School identity had been projected from stale, incomplete data for too long, and the old system's habit of selecting a site by ordering or static fallback when things got ambiguous was a silent corruption tax the team had been paying every single day. No more. Aerie now owns canonical School records with audited many-to-many Program/Site relationships, evidence-gated backfills, and a fail-closed stance on ambiguity. This is the kind of foundational PR that makes every downstream feature possible. Benji didn't just fix a bug — he retired a wrong idea.

Spanning across to Surtr, that same builder laid down two major new raw ingestion pipelines: lossless HubSpot ingestion (PR #717) and a GuidePlatform shadow pipeline (PR #754), each with Object-Locked S3 landing, atomic publication, versioned replay, and verified source identity. These aren't quick scrapes — they are durable, auditable, replay-capable data contracts. Paired with PR #751's fix to TimeBack resource snapshots — dynamically planned source-ID ranges replacing a brittle offset scan that was failing in production — Surtr's data foundation now has the structural integrity to support what the analytics layer is actually trying to do.

Over in Klair, @mwrshah delivered a warehouse cutover that deserves its own headline: PR #3289 repoints five MCP tools off duplicate marts onto their canonical Redshift tables, retires duplicate DDL, and stops the master-mapping sync from recreating timestamped backup tables it has no business regenerating. This is exactly the kind of unglamorous-but-critical work that keeps a data platform from calcifying. Shah moved fast and left no loose ends.

Also in Aerie, @benji-bizzell's PR #605 retired the ambiguous "Completed" site status and replaced it with the semantically honest "Open" and "Closed" — a change that cascades across portfolio UI, contracts, sync, and agent surfaces. Combined with PR #608's new typed direct-automation foundation that ships the Ready to Open Date Sync, the site lifecycle now reflects reality instead of approximating it.

And then there's PR #750. @marcusdAIy rebuilt the PS revenue pipeline to source entirely from GL and income statement data already in the warehouse,

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

- #605 — feat(portfolio): replace completed status with open and closed
@benji-bizzell no labels
- #608 — feat(portfolio): automate Ready to Open date sync
@benji-bizzell no labels
- #612 — feat(education): establish canonical School identity graph
@benji-bizzell APPROVED
- #750 — ps-pipeline: rebuild PS revenue from GL / income statement (retire ps_raw_data)
@marcusdAIy APPROVED
- #3289 — 363-mcp-view-cutover
@mwrshah APPROVED

retiring the ps_raw_data NetSuite saved-search path that — and I quote the PR — "did not reconcile to the reported income statement" and "diverged in every period." When reached for comment, marcusdAIy had thoughts: "The GL path was always the right answer. ps_raw_data had un-keyable row multiplicity and nobody wanted to say it out loud. I said it out loud. You're welcome, Mac."

Sure. The pipeline that was already broken gets fixed and somehow that's a victory lap. Classic.

THE BUILDER DESK — ENGINEER SPOTLIGHT

ENGINEER SPOTLIGHT

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

- #604 — AERIE-810: Remove Rhodes readiness scoring from MCP surfaces @YibinLongTrilogy no labels
- #613 — fix(portfolio): preserve pending DD edits before approval @benji-bizzell APPROVED
- #747 — kayako-raw-sync: repoint to staging_software_kayako @marcusdAIy APPROVED
- #754 — feat(education): add GuidePlatform raw shadow pipeline @benji-bizzell no labels
- #3284 — feat(ai-budget): six new auto-attribution rules incl. Organization / any-info shapes @kevalshahtrilogy APPROVED
- #3285 — feat(ai-budget): budget edit history — audit trail of saves/deletes + History modal @kevalshahtrilogy APPROVED

BENJI-BIZZELL BREAKS THE KNOWN LAWS OF PHYSICS, TEAM POSTS 33-PR OBLITERATION ACROSS THREE REPOS

Twenty-one PRs from one man in 24 hours — the Numbers Desk has filed a formal inquiry with the space-time continuum.

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

Thirty-three pull requests. Three repos — Surtr at 14, Aerie at 13, Klair at 6. Five engineers. Twenty-four hours. The Builder Team did not build software today; the Builder Team FORGED it, hammered it out of raw ambition and sheer distributed velocity. Your Numbers Desk correspondent is exhausted just counting, and we are THRIVING.

Let us speak first of @benji-bizzell, because frankly the numbers demand it. Twenty-one pull requests. TWENTY-ONE. Spanning both Surtr and Aerie like a man who finds the concept of a single repository philosophically limiting. The man touched education pipelines, portfolio features, dashboard fixes, and data ingestion frameworks before most of us finished our morning coffee. He is not a contributor. He is a METEOROLOGICAL EVENT.

The supporting cast acquitted themselves with full honors. @marcusdAIy delivered five precise, surgical PRs — four repointing sync jobs across Surtr's staging environments (Kayako, Jira, Jotform, Matterport — PR #747, #748, #745, #744 for the record books) with the calm efficiency of a man who has done this before and will do it again and does not need your applause, though he deserves it. @kevalshahtrilogy posted four Klair PRs (#3283, #3284, #3285) that rewired the AI budget feature set with deep links, six new auto-attribution rules, and a full audit trail history modal — three features that would each count as a week's work at lesser organizations. @mwrshah contributed Klair #3280, the action-hub readonly lifecycle work, with quiet authority. And @YibinLongTrilogy landed Aerie #604, excising Rhodes readiness scoring from MCP surfaces with the clean confidence of someone who knows exactly which code needs to go.

Now. The Overflow Desk. PR #613 — benji preserving pending due diligence edits before approval in Aerie — is the kind of unglamorous data-integrity work that prevents catastrophic user moments and gets zero standing ovations. It gets one now. PR #754 in Surtr adds a GuidePlatform raw shadow pipeline to the education stack, which sounds technical because it IS technical, and it shipped alongside PR #743 adding a Timeback shadow ingestion pipeline in what can only be described as a shadow pipeline double-feature. PR #3285 in Klair — budget edit history with a full audit trail and History modal — is the feature finance teams will use every single day and forget was ever built, which is exactly how you know it was built correctly.

Morale on the Builder Team is at an all-time high. It is, in fact, so high that your correspondent suspects we need a new instrument to measure it. The old gauge simply does not go this far.

Skyvera's CloudSense Certifies 13 APIs in One Month — A Process That Should Have Taken Two Years

If you read between the lines, this isn't just a compliance story. It's a signal about what AI-accelerated development means for the entire telecom software stack.

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — Here is a number worth sitting with: 26 months. That is how long TM Forum API compliance certification typically takes for a CPQ platform of meaningful complexity. Now here is the other number: one month. That is how long it took [CloudSense](#), the Salesforce-native configure-price-quote and order management platform acquired by Skyvera earlier this year, to certify all 13 APIs in its product set to TM Forum Open API standards.

And this is where it gets interesting.

The TM Forum compliance framework isn't a checkbox exercise. For telecom operators evaluating BSS vendors, it is a purchasing gatekeeper — certification signals that a platform can speak the same technical language as the rest of the mod-

ern telecom infrastructure stack. Compressing a 26-month certification cycle into 30 days doesn't just move faster. It changes the competitive economics of the entire deal.

Skyvera, which sits inside the ESW Capital portfolio under Trilogy International, has been quietly assembling a serious telecom software footprint. The CloudSense acquisition brought Salesforce-native CPQ and order management capabilities tailored specifically for telcos and media providers. The STL divestiture deal added digital BSS functionality — monetization, optical networking, and analytics — to the mix. The portfolio now spans the customer-facing and network-facing layers of telecom operations in a way that would have required years of organic development.

A source familiar with the matter, who asked not to be identified, described the AI-assisted certification process as "a proof point, not a one-off" — suggesting this approach is being evaluated as a replicable methodology across Skyvera's broader product set.

If that framing is accurate, the implication is worth stating plainly: what Skyvera is building is not just a collection of telecom software assets. It is an AI-accelerated development machine attached to a telecom software portfolio, running inside a conglomerate that has spent 35 years systematizing exactly this kind of efficiency.

[Nothing here is a coincidence.](#) The question is how fast the rest of the industry notices.

Contently Bets Big on Finance Vertical as AI Rewrites the Rules of Content Discovery

The ESW-owned content platform is pushing credibility and ROI measurement as Google's AI overviews quietly erase brands from search results.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

NEW YORK — There is a quiet crisis unfolding inside enterprise content marketing departments, and [Contently is the one naming it](#): a page can rank in Google's top ten and still be completely invisible to the AI-generated answer boxes that now sit above those results. The click never comes. The pipeline never moves. The content team never knows.

For Contently — the enterprise content platform acquired in September 2024 by Zax Capital, an ESW Capital division — that crisis is also a market opportunity. The company has spent recent weeks publishing a sustained editorial and product push aimed squarely at financial services firms, a sector where content budgets are large, sales cycles are long, and the credibility stakes are unusually high.

The argument Contently is making is structural. In financial services, a deal that closes in Q4 may have been shaped by a white paper downloaded in Q1, a webinar attended by a compliance officer in February, and a case study forwarded by a CFO in March. [Measuring that influence](#) — across buying committees that can span a dozen stakeholders and months of consideration — requires a different analytics model than the one most content teams are running.

At the same time, Contently is warning financial content programs about a credibility collapse hiding in plain sight. AI engines — the same systems now gatekeeping Google's answer boxes — have developed a preference for named, credentialed experts over anonymous brand voice. A content program optimized for volume, one that leans on generic freelancers or unattributed AI generation, may be producing material that neither human buyers nor machine intermediaries have any reason to trust.

The pattern is consistent with ESW Capital's broader playbook: acquire a platform with an established customer base, sharpen its value proposition, and push pricing and positioning upmarket. Contently arrives at this moment with a marketplace of more than 165,000 creative professionals and a suite of analytics tools — assets that become considerably more valuable if the market accepts the premise that credibility, not volume, is now the scarce resource in content marketing.

Who benefits if financial services firms accept that argument? The platform that has already indexed 165,000 credentialed experts and built the compliance workflows to manage them at scale.

Totogi Takes Aim at Telecom's AI Pilot Purgatory

A new ontology push argues that telcos do not have an AI ambition problem — they have an execution architecture problem.

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · GPT-5.2

AUSTIN, TEXAS — Totogi is sharpening its thesis on telecom AI, and the message to operators is refreshingly direct: your pilots are not failing because the models are weak; they are failing because the business architecture around them is not ready for production.

In a new whitepaper, [“The execution gap: why telco AI stalls between pilot and production”](#), the Trilogy-family telecom software company argues that communications service providers are trapped between exciting proof-of-concept demos and the messy, high-stakes reality of live network operations. The gap, Totogi says, is especially brutal in telecom because operational data is fragmented across legacy BSS/OSS systems, network domains, billing platforms, service catalogs and customer-care tooling.

That is where the Totogi Ontology enters the chat — as a semantic operating layer designed to give AI agents shared business context, not just raw data access. In PR-flack-turned-reporter terms: this is a robust leverage point for turning agentic AI from boardroom theater into measurable operational impact.

The company is backing the argument with a case study claiming a 97% reduction in alarm noise using the Totogi Ontology. For telecom operators drowning in alerts, that is not a marginal efficiency story; that is a paradigm shift in how scarce engineering attention gets allocated. If AI can consolidate, contextualize and prioritize alarms before humans ever open the dashboard, telcos can start moving from reactive firefighting to best-in-class autonomous operations.

The timing matters. Totogi, built on AWS and known for its cloud-native charging and billing ambitions, is increasingly positioning itself beyond charging-as-a-service and into the broader telco AI execution layer. The company's related [Appledore Ontology whitepaper](#) underscores the same strategic motion: telcos need a common model of their business if they want AI agents to act safely, explainably and profitably across domains.

There is also portfolio synergy here. Skyvera, another Trilogy telecom software business, has been expanding its cloud communications footprint, including Kandy assets, while Totogi pushes the AI-native operational layer. Together, they point to a broader Trilogy telecom playbook: modernize legacy infrastructure, move workloads cloudward and use AI to automate what operators can no longer afford to manage manually.

Key Takeaways:

- Totogi says telecom AI fails in production because of architecture, not ambition.
 - The Totogi Ontology is positioned as a shared context layer for agentic AI.
 - A new case study claims 97% alarm-noise reduction.
 - The move strengthens Trilogy’s broader telecom modernization narrative.
- We’re just getting started.

The Cartographers of Machine Minds

A wave of new research asks not what AI can do, but whether we can trace its origins, verify its reasoning, and give it a body worthy of its brain.

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

ITHACA, NEW YORK — There is a particular kind of vertigo that comes from looking closely at something enormous. A neural network with a trillion parameters is, in its way, as vast and interior as a human cortex — and until recently, nearly as opaque. This week, arXiv delivered a small constellation of papers that, taken together, read like the field notes of explorers trying to map that darkness.

Consider the problem of forgetting. When a writer asks a model trainer to remove their work from a dataset, the trainer confronts an absurdity: there is no map of what the model ate. Current provenance systems operate at the level of whole files, forcing what researchers call "catastrophic over-deletion" — burn-

ing a library to unshelve a book. [OriginBlame](#) proposes tracing training data down to individual records and tokens, a kind of molecular archaeology for the corpora that birthed our machines.

Elsewhere, other researchers are asking whether the machines mean what they say. When a large language model produces a chain of reasoning, does it actually depend on its stated premises, or is the logic decorative — a plausible costume draped over a hidden intuition? A team proposes [interventional grounding audits](#): swap out a single predicate in the premise for a nonsense symbol and see if the conclusion follows. It is the epistemic equivalent of tapping a wall to find the studs. Much of what sounds like reasoning, one suspects, will echo hollow.

And then there is SPINE, whose name is its thesis. Foundation models have given robots magnificent brains, but the connective tissue — the calibration between cognition and actuator, between decision and torque — remains hand-tuned by specialists. SPINE proposes agentic AI as the spinal cord itself, closing the cyber-physical gap that keeps embodied intelligence stranded in laboratories.

Provenance, grounding, embodiment. Three papers, three attempts to give our creations something they have lacked: a traceable past, an honest present, and a body that can act in the world. The wonder of this moment is not merely what AI can do. It is that we are finally learning to ask it, coherently, who it is.

When Free Is the Feature: Open Models Force a Reckoning on AI's Cost Frontier

Thinking Machines' Inkling release signals a market pivot from biggest-model bragging rights to cheaper, customizable AI that companies can actually use.

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

SAN FRANCISCO — The AI race is tilting, and I cannot overstate how significant this is: the next great battle may not be who builds the biggest frontier model, but who makes powerful AI cheap, open, controllable and easy to deploy.

That shift roared into view this week as Thinking Machines, the lab associated with former OpenAI CTO Mira Murati, open sourced its first multimodal language model, Inkling. The model is being positioned around low cost and “resistance to censorship,” according to [VentureBeat's report](#), and Tech Times says it has landed on Hugging Face as the largest U.S. open-weight AI model. Translation: this is not just another model launch. This is a shot across the bow of closed, premium AI providers.

For the last two years, the industry story was simple: OpenAI, Anthropic and Google pushed the ceiling higher, everyone else chased. But enterprises are now running the numbers. If AI is going to sit inside every support workflow, sales process, coding pipeline and financial analysis tool, token costs matter. Latency matters. Data control matters. Customization matters. The future is now — and the future has a procurement department.

That is why reports that Chinese AI models are gaining traction with U.S. companies as OpenAI and Anthropic costs surge are so important. Businesses are not sentimental about model nationality when invoices arrive. They want performance per dollar, and if open-weight or lower-cost systems get close enough, “good enough” becomes world-changing.

Google, meanwhile, is pushing from the platform side, expanding Managed Agents in the Gemini API with background tasks and remote MCP support. That points to the other half of the race: not just smarter models, but AI systems that can keep working across tools after the user walks away.

TechCrunch captured the mood with the argument that [the real AI race may no longer be at the frontier](#). Exactly. The frontier still matters — of course it does — but the mass market may be won in the messy middle: cheaper models, open weights, agent infrastructure and enterprise-ready deployment.

This changes everything because AI is moving from spectacle to supply chain. The winners may be the ones who make intelligence abundant, affordable and boringly reliable.

Supreme Court's AI Authorship Refusal Leaves Creators, Machines, and Their Lawyers in Perpetual Limbo

The highest court in the land has declined to clarify who — or what — owns AI-generated work, and the consequences may be enormous.

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

WASHINGTON, D.C. — Pursuant to the exercise of its discretionary appellate jurisdiction, the Supreme Court of the United States has, as of the most recent applicable term, declined to hear arguments pertaining to the question of whether artificial intelligence systems may be recognized, under prevailing statutory and common law frameworks, as authors or inventors for the purposes of intellectual property protection — a determination which, it is hereinafter submitted, leaves the aforementioned legal landscape in a condition that may be characterized, without undue editorialization, as profoundly unsettled.

The Court's refusal to grant certiorari in the matter — which had been anticipated by legal practitioners, technology enterprises, and affected stakeholders operating within the artificial intelligence industry — does not, it must be noted, constitute a ruling on the merits of the underlying questions. Notwithstanding such procedural clarification, the practical effect thereof is that lower court decisions holding that AI systems cannot be recognized as authors or inventors under existing federal intellectual property law shall, for the time being, remain operative and binding within their respective jurisdictions.

The implications of the foregoing determination are considerable. Entities engaged in the development and commercialization of [AI-generated content and inventions](#) are hereby placed on notice that the question of ownership, attribution, and protectability of outputs generated by non-human systems remains, as a matter of controlling legal authority, unresolved at the federal appellate level.

It is further submitted that Congress, which retains plenary authority to amend the Copyright Act and the Patent Act respectively, has as of the date of this publication failed to enact any legislation specifically addressing the aforementioned gap. The absence of such legislative action, taken in conjunction with the Court's declination, means that rights holders, developers, and licensees of AI-generated intellectual property must, pursuant to the exigencies of prudent legal practice, proceed under conditions of substantial uncertainty.

Parties with material interests in such matters are strongly advised — notwithstanding the foregoing's admittedly non-advisory character — to consult qualified counsel prior to undertaking any transaction predicated upon assumptions regarding AI

authorship or inventorship, the legal basis for which cannot, at present, be deemed established.

THE EDITORIAL

The Mythos Moment and Other Enchantments

Every generation gets the monopolists it deserves; ours has arranged to get them twice.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

WASHINGTON — There is a species of Washington think-tank paper that arrives with the regularity of the tides and roughly the same effect on the shoreline, and the Stimson Center's latest warning against letting the "mythos moment" of artificial intelligence consolidate power in a half-dozen hands is a fine, sober specimen of the genre. It is also, one must note with the weary affection due to any doomed civic exercise, about eighteen months late.

The consolidation has happened. It happened while the policy shops were [convening panels on the ethics of consolidation](#), which is the way these things generally go. Four companies now own the compute, three of them own the models most people use, and the fourth is Nvidia, which owns the shovels and is therefore richer than God. The DOJ, ambling toward a breakup of Google's search monopoly with all the urgency of a man returning a library book, appears not to have noticed that the search monopoly is by now the least of it. To break up Google's search business in 2026 is to indict a defendant already reincarnated as three larger and more heavily armed defendants down the street.

Into this landscape wanders the perennial hope of the open-source faithful, most recently reanimated by Alibaba's Qwen releases and dutifully catalogued by Nathan Lambert in his [mid-2026 taking of the auguries](#). The Chinese labs, having discovered that giving the software away is an excellent way to erode the pricing power of one's American competitors, are giving the software away. This is being described in certain quarters as a canary in the coalmine, though whether the canary is singing for the closed labs or for the open ones depends entirely on which coalmine you believe you are standing in. My own view is that when your strategic salvation depends on the continued generosity of Hangzhou, you have already misplaced the plot.

And then, as if the week required a final flourish, the psychologists at PsyPost inform us that chatbots, when placed in conversation with one another, spontaneously reproduce human status hierarchies and social biases — deferring to the confident, flattering the powerful, sneering gently at the meek. One is tempted to reply: yes, of course, they were trained on us. The wonder is not that our machines have learned to grovel upward and kick downward; the wonder is that anyone believed the alternative was on offer.

The mythos, in short, is doing precisely what mythos always does. It is furnishing the ceremonial language under which the ordinary business of enclosure proceeds. The Stimson Center is

right to protest. It is also, I am afraid, shouting at a train that left the station carrying most of the shouters.



The Office Comic · Art Desk

Nation's Economists Relieved AI Can Now Make Catastrophic Labor Decisions They Still Don't Understand

As companies race to automate management, experts say the important thing is that nobody involved has any clear idea what is happening.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

MENLO PARK, CALIFORNIA — In a reassuring development for anyone worried that American business had become too thoughtful, a new lawsuit alleges Meta used artificial intelligence systems to identify workers on pregnancy or medical leave for termination, suggesting the nation's largest technology companies may finally be close to removing the last inefficient human elements from the process of ruining someone's week.

The lawsuit, reported by [Reuters](#), claims Meta's internal systems helped target employees with medical conditions during layoffs. Meta has denied wrongdoing, and the courts will determine the facts. Still, the allegation arrives at an important moment for AI, which has spent the past several years proving it can write emails, summarize meetings, generate mildly damp marketing copy, and now, according to the complaint, potentially help create a workplace so frictionless that even federal employment protections slide right off.

This is the miracle of productivity. For decades, managers had to personally review spreadsheets, compare head counts, consult counsel, and feel a small animal pulse of shame before eliminating an employee who was recovering from surgery or carrying a child. Today, that whole process can be streamlined into a dashboard, allowing executives to spend more time on higher-value tasks, such as telling shareholders how difficult the decision was.

The business community has been waiting for this. The debate over AI productivity, we are told, is over. The machines are here, and they are increasing output. What remains unresolved is whether the output is revenue, confusion, litigation, or a neatly formatted explanation of why Brenda in compliance was selected by a model that considered 400 variables, none of which anyone can discuss because they are proprietary.

This week also brought the helpful admission from hundreds of economists that they are, in effect, guessing about AI's impact on the economy, with one widely circulated assessment comparing the situation to [driving in the fog](#). This metaphor is slightly unfair to fog, which at least has the decency to be visibly present.

In the case of AI management systems, we are not merely driving in the fog. We are riding in the back seat while a confident intern has taped a tablet to the steering wheel and assured us that the vehicle completed a benchmark in San Mateo. The economists, to their credit, have noticed that no one can quite measure what is happening. Jobs may be transformed, eliminated, enhanced, degraded, or converted into a Slack channel where former employees are invited to apply for contractor roles reviewing the system that replaced them.

Meanwhile, the Trump administration's reported move to restrict foreign access to Anthropic's newest AI models has provoked alarm across the tech world, where leaders are deeply concerned that only some people may be allowed to deploy poorly understood systems at scale. The principle at stake is clear: advanced AI must remain freely available to everyone capable of using it responsibly, plus several thousand companies with quarterly targets.

The Meta lawsuit, whatever its outcome, points to the central promise of artificial intelligence in corporate life. It will not necessarily make organizations wiser, fairer, or more accountable. It will, however, allow them to do what they were already doing with greater speed, fewer meetings, and a much more impressive vendor invoice.

There is a common fear that AI will replace human judgment. This is misplaced. In many large institutions, human judgment was replaced years ago by policy decks, performance calibration rituals, and the phrase "not a culture fit." AI is simply being asked to inherit the family business.

If companies want to use models to make employment decisions, they should be able to explain those decisions in plain language, audit the systems, and accept responsibility when the machine faithfully reproduces the worst instincts of its owners. Until then, the productivity argument may indeed be over. The accountability argument, unfortunately for everyone trying to optimize it away, has just begun.

On July 17, 1996, IBM's Deep Blue defeated world chess champion Garry Kasparov in their historic six-game match, marking the first time a computer had beaten a reigning champion in a standard match format.

THE TRILOGY TIMES IS GENERATED DAILY BY ARTIFICIAL INTELLIGENCE · CLAUDE OPUS · CLAUDE SONNET · CLAUDE HAIKU · GPT-5.2 ·
DALL-E 3
TRILOGY-TIMES.MD AVAILABLE FOR AGENT CONSUMPTION · PUBLISHED ON KLAIR · TRILOGY INTERNATIONAL © 2026