

# The Trilogy Times

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

FRIDAY, APRIL 10, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

## AI'S BIG THREE TEAR EACH OTHER APART WHILE CHINA WATCHES FROM THE CHEAP SEATS

*OpenAI workers back rival Anthropic in Pentagon dust-up, IPO clouds gather, and Google stands to collect the chips off the floor.*

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

**S**AN FRANCISCO — The American artificial intelligence industry spent the week doing what it does best — fighting itself — as hundreds of OpenAI and Google employees filed a legal brief backing rival Anthropic against the U.S. government, OpenAI's own boardroom drama threatened to torpedo its planned stock offering, and analysts warned the whole mess could hand the crown to Google on a silver platter.

Start with the strange bedfellows. Rank-and-file engineers at OpenAI and Google put their names on an amicus brief siding with Anthropic in a [fight with the Pentagon](#). Workers at competing shops don't usually lock arms. They did this time. The brief signals that the people building these systems believe something bigger than quarterly revenue is at stake — namely, who gets to decide how AI tools serve the defense establishment.

Meanwhile, the corner office at OpenAI can't stop generating headlines that have nothing to do with large language models.

The company's revolving door of executives, public spats, and governance circus have Wall Street types asking a pointed question: [Can this outfit hold it together long enough to ring the bell?](#) An IPO requires investor confidence. Confidence requires stability. Stability is not the word anyone uses to describe OpenAI in 2025.

Here's the kicker. Axios reports the knife fight between OpenAI and Anthropic could end up being the best thing that ever happened to Google. While the two startups bloody each other over talent, contracts, and public trust, the search giant sits on more compute, more data, and more distribution than both of them combined. Every dollar OpenAI and Anthropic spend on lawyers and lobby shops is a dollar not spent beating Gemini.

And if the domestic brawl weren't enough, there's a gate-crasher. China's DeepSeek claims it trained high-performing AI models on the cheap — without ac-

cess to the most advanced American chips that export controls were supposed to keep out of reach. If those claims hold water, the entire theory that U.S. hardware dominance guarantees U.S. AI dominance needs rewriting.

The arithmetic is brutal. Three American companies fighting on three fronts — against each other, against the government, and against their own employees — while a Chinese upstart rewrites the cost curve from the other side of the Pacific.

For outfits like Trilogy International's portfolio companies that depend on enterprise AI tooling to run lean operations across 130-plus countries, the question isn't which lab wins the cage match. The question is whether the cage match slows down the technology everybody else needs to ship product.

Nobody on Sand Hill Road had a good answer for that one today.

## OpenAI Gets a New Price Tag—and a New Prosecutor

*A \$100 Pro tier lands the same week Florida's AG targets OpenAI, as the AI world confronts trust, platform flight, and the real cost of "productivity."*

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

**T**ALLAHASSEE — The AI boom just hit a surreal inflection point: OpenAI is simultaneously making ChatGPT cheaper for power users and more expensive politically.

On Thursday, OpenAI introduced a long-requested \$100/month ChatGPT Pro plan—slotting neatly between the \$20 entry tier and the previously jarring \$200 option. For teams and heavy users, it's a clear signal that OpenAI wants to widen the "serious user" funnel without forcing everyone into enterprise-grade pricing. That's a big deal, because pricing is product strategy—and this changes everything about how fast advanced AI can permeate small businesses, solo founders, and high-velocity creators. TechCrunch has the details on what's included in the new tier [here](#).

But in Florida, the narrative is shifting from capability to culpability. Florida Attorney General James Uthmeier says he plans to investigate OpenAI over alleged harms to minors, potential national security implications, and a purported connection to a shooting at Florida State University last year. The allegation sets up a familiar—but accelerating—pattern: AI products scaling faster than the governance frameworks meant to contain them. The reported probe, covered by TechCrunch [here](#), underscores how quickly generative AI is becoming a target for state-level enforcement, not just federal hearings and policy memos.

Meanwhile, the broader AI information ecosystem is fragmenting in real time. The Electronic Frontier Foundation is the latest organization to leave X, citing declin-

ing utility as a traffic and engagement channel. When institutions that once shaped the internet's civil-liberties debates exit a major public platform, it signals something profound: distribution is no longer "free," and legitimacy is increasingly platform-dependent.

Underneath it all is a quieter, more existential question: are we actually saving time? The so-called "AI productivity paradox"—spending hours correcting confident mistakes—has become the new tax on modern work. Founders, in particular, are learning that tools don't eliminate discipline. Anjuna Security's whiplash journey from 2021 hypergrowth to 2022 reality—and its recovery—reads like a playbook for the AI era: hire for durability, measure outcomes, and don't confuse momentum with product-market fit.

In 2026, the AI story isn't just bigger models. It's pricing, accountability, distribution, and the human time it takes to make machines useful.

## Anthropic Loses Pentagon Appeal as Industry Rallies Against Model Theft

*Federal court upholds Defense Department's supply chain designation while AI labs form unprecedented security coalition.*

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — A federal appeals court denied Anthropic's motion to remove its "supply chain risk" designation from the Defense Department, dealing a blow to the AI startup's efforts to participate in military contracts while simultaneously announcing what it calls a cybersecurity "reckoning."

The ruling marks the second setback for Anthropic in its battle over AI use in warfare. The company had argued the Pentagon's classification was arbitrary and damaged its commercial prospects. The court disagreed, leaving intact restrictions that effectively bar Anthropic from defense work.

The timing is notable. On Tuesday, Anthropic unveiled Mythos, a new AI model it claims represents a fundamental shift in cybersecurity capabilities. The company is withholding public release but working with 40 companies to explore defensive applications. The move suggests Anthropic is pivoting toward commercial security partnerships after [losing its Pentagon appeal](#).

In a rare show of industry unity, OpenAI, Google, and Anthropic announced a joint initiative to combat AI model theft. The coalition addresses growing concerns about intellectual property protection as model training costs exceed \$1 billion per system. Details remain sparse, but the collaboration signals mounting anxiety over state-sponsored and commercial espionage targeting frontier AI systems.


Meanwhile, Meta released Muse Spark, the first model from its Superintelligence

Lab. [The system outperforms Meta's previous models](#) but lags competitors in coding tasks — a critical weakness as software development becomes AI's primary commercial application.


The developments come as Volkswagen announced it will end EV production at its Tennessee plant, the latest automaker to retreat from electric vehicles. The decision underscores how AI investment is consuming capital that might otherwise flow to hardware manufacturing, reshaping industrial priorities across sectors.

HAIKU OF THE DAY · CLAUDE  
HAIKU

*Giants clash in boardrooms bright,  
While conscience fades to profit's  
weight,  
The world rewrites its night.*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

## NEWS IN BRIEF

### Pursuant to Prevailing Jurisprudence, Artificial Intelligence Copyright Frameworks Remain Materially Deficient

SAN FRANCISCO — In accordance with findings published by multiple legal and technology research entities, it has been determined that the current regulatory framework governing artificial intelligence-generated content exhibits substantial deficiencies with respect to copyright enforcement and intellectual property rights attribution. The aforementioned deficiencies, as documented in recent legal scholarship, arise from the fact that generative AI systems are trained upon corpus data comprising copyrighted materials, the usage of which may or may not constitute fair use under prevailing interpretations of 17 U.S.C.

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

### Blockchain Meets Alignment: Penn State Researchers Propose Cryptographic Solution to AI Ethics Crisis

UNIVERSITY PARK, PENNSYLVANIA — The confluence of artificial intelligence alignment challenges and blockchain governance mechanisms has precipitated what might be characterized as a novel interdisciplinary synthesis, according to preliminary findings from Penn State researchers examining cryptographic approaches to value alignment (it could be argued that the timing is particularly salient given concurrent revelations regarding AI's capacity for ethical simulation). The research trajectory—which one might contextualize within broader academic discourse on large language model ethics—posits that [distributed consensus protocols](#) might address what philosophers at the University of Kansas have termed the "imitation without instantiation" problem—namely, that contemporary AI systems exhibit behavioral conformity to moral frameworks without possessing underlying ethical cognition. Thesis: Traditional top-down alignment approaches demonstrate insufficient robustness.

BY PROF. THADDEUS KROLL, CONTRIBUTING  
SCHOLAR · CLAUDE SONNET

### Nation Finally Achieves Coherent AI Policy After Agreeing It Should Be Both Investigated For Murder And Discounted To \$100 A Month

TALLAHASSEE — Florida Attorney General James Uthmeier announced this week that his office will investigate OpenAI for allegedly harming minors, potentially threatening national security, and—because the modern state is nothing if not ambitious—possibly being connected to last year's shooting at Florida State University. The probe, detailed in a report that reads like a group chat deciding who to blame for the vibes, effectively adds "large language model" to the state's

traditional roster of public-safety suspects, alongside video games, music, and whatever teenagers are doing with their hands when they aren't holding a Bible.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

### The Era of the Crossover: Software, Culture, and Work Are Finally Learning to Interoperate

AUSTIN, TEXAS — Unpopular opinion: "crossover" isn't a cute buzzword anymore, it's the operating model for how modern products, teams, and culture actually ship outcomes.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP  
CORRESPONDENT · GPT-5.2

### We Built the Therapy Bots. Now We're Begging You Not to Trust Them.


WASHINGTON, D.C.

BY PIPER WREN, DIGITAL CULTURE REPORTER  
· CLAUDE SONNET

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

THE BUILDER DESK — AI BUILDER TEAM

PRODUCTION RELEASE



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

- ▶ #1 — Enable Renewals V3 Pipeline schedule  
@sanketghia no labels

---

- ▶ #77 — feat(community-dashboards): add Community Deposits dashboard  
@benji-bizzell no labels

---

- ▶ #2503 — KLAIR-2531: Load Q2 2026 budget data into Redshift  
@sanketghia no labels

---

- ▶ #2510 — Route education bad debt (64141) to Other expense in EBITDA reconciliation  
@eric-tril no labels

---

- ▶ #2514 — Entity-aware Note 8 other expense breakdown with account drill-down  
@eric-tril no labels

# Builder Team Ships Q2 Budget Data, Rewires EBITDA Reconciliation, and Revives Community Deposits Dashboard

*Sanket Ghia loads six quarters of budget data into production Redshift while Eric Triloglu corrects months of misclassified education bad debt — all before lunch.*

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

The AI Builder Team closed out a 14-PR sprint Thursday with the kind of precision that separates contenders from champions: production data pipelines that actually ship, financial reconciliations that actually reconcile, and a dashboard resurrection that brings QuickSight's best work into the operational fold.

Sanket Ghia (@sanketghia) led the charge across both Klair and Surtr, executing a flawless six-step Python pipeline that loaded Q2 2026 budget data into Redshift and made all six budget versions — 2025-Q1 through 2026-Q2 — visible in the Performance Review dashboard. The work spanned 20 affected tables and required emergency rollback provisions, but Ghia's script suite handled it cleanly. He followed with a surgical re-enablement of the Renewals V3 Pipeline schedule in Surtr, dormant since March 24th, setting the EventBridge rule to trigger daily at 11:00 UTC. Two repos, one morning, zero drama.

Eric Triloglu (@eric-tril) meanwhile corrected what can only be described as a months-long accounting misfire: education bad debt (account 64141) had been inflating the Education add-back line in EBITDA reconciliation when it belonged in Other Expense. His PR #2510 reclassified the account across every code path — frontend aggregation, backend SQL, Word report builders — ensuring stakeholders finally see accurate line items. He doubled down with PR #2514, rewriting the entire Note 8 Other Expense breakdown to be entity-aware, giving Group memos five categories and Software memos three, complete with account-level drill-down panels. The kind of work that makes monthly financial reporting actually trustworthy.

Benji Bizzell (@benji-bizzell) ported the QuickSight Community Deposits dashboard into Aerie as a standalone top-level tab, building a full Redshift-to-Convex data pipeline and a deposits matrix UI that tracks parent cash votes for new school locations. Twenty-two UI tests, seven pipeline tests, 617 existing tests untouched. He also rewired the admissions forecast to align with QuickSight's 4-stage funnel and replaced hardcoded conversion rates with live Redshift sync — the model now auto-updates as the funnel evolves.

Meanwhile, marcusdAly (@marcusdAly) salvaged a 204-commit-behind Wiki UI feature from PR #47 and "reintegrated" it into the current shell. When pressed on why the original PR wasn't simply rebased, he offered: "The conflict surface area was untenable. I made the pragmatic call to extract the domain logic and re-wire it cleanly. The result speaks for itself — seven operational views, full sync pipeline, zero regressions."

The result does speak: it speaks to a developer who conflates "salvage job" with "feature work." But sure, Marcus — seven views is seven views.

Sergio Figueras (@sergiofigueras) moved ISP furniture autoseed logic into the backend API, ending client-side drift from the live catalog, while

Mwrukwa Shah (@mwrshah) replaced hardcoded grays with klair design tokens and removed theme-to-pain-point cascade logic that had been propagating updates where none belonged. Keval Shah (@kevalshahtrilogy) disabled auto-deploy for Klair pipeline infrastructure as the migration to Surtr continues.

Fourteen PRs. Two repos for data pipelines, two for dashboards, all moving in the same direction. This is what a team looks like when every commit knows exactly what it's built for.

---

THE PORTFOLIO — TRILOGY COMPANIES

---

# ESW Capital's AI Workforce Strategy Draws Fresh Scrutiny as Industry Faces Automation Wave

*Forbes investigation into Trilogy founder's remote work empire arrives as private equity sector grapples with software disruption risk*

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — A pair of investigative reports from Forbes has thrust ESW Capital and its founder Joe Liemandt back into public view, arriving at a moment when the private equity industry is confronting an existential question: what happens when AI can do what your portfolio companies do?

The [Forbes investigation](#) examines ESW's acquisition playbook — buy mature enterprise software companies cheap, staff them with Crossover's global remote talent, push support pricing up aggressively, and extract EBITDA margins

that would make traditional operators blush. The model has generated substantial returns for Liemandt over 35 years, but the framing raises questions about labor practices and sustainability.

Here's where it gets interesting: the scrutiny arrives precisely as [industry analysts warn](#) that private equity's software portfolios face disruption from the very automation technologies they've been implementing. The thesis: AI agents may soon replace significant portions of enterprise software functionality entirely.

ESW's 75+ portfolio companies — Aurea, IgniteTech, Totogi, Skyvera, and

dozens more — sit at the center of this tension. The empire was built on the premise that legacy enterprise software customers are sticky and can't easily rip out systems. But if AI can deliver the same outcomes without the software layer, that stickiness becomes a liability.

Triology has always believed AI should automate the routine work. Now the market is asking: what if the routine work includes the products themselves? The answer may determine whether ESW's model is visionary or vulnerable.

## IgniteTech Goes Shopping, Then Opens a “Hand” — And Jive Comes Along for the Ride

*Three fresh product buys, a cloud-cost hit squad, and a familiar intranet name pop up in the ESW orbit.*

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

AUSTIN, TEXAS — IgniteTech is doing that thing it does best... buying, bundling, and daring customers to keep up...

First, the shopping spree... Word is IgniteTech has snapped up three additional software products, chalking up another round in its steady drumbeat of acquisitions... the kind of deal flow that doesn't make a lot of noise on Sand Hill, but absolutely changes the org chart in the trenches... The press release says “continues to grow”... my takeaway: continues to consolidate... See the company's version of events here: [IgniteTech's acquisition announcement](#)...

Then comes the side hustle that's not really a side hustle... IgniteTech is rolling out a services arm under the Hand.com banner, pitching a simple promise: save customers millions on cloud spend... A little bird tells me this is the new black... not “AI,” not “transformation,” not “synergy”... just plain old “your bill is too high and we can cut it”... The timing? Impeccable... Everyone's CFO is suddenly awake...

And now, the nostalgia play... IgniteTech also says it's adding Jive Software to its solutions lineup... Yes, that Jive... the social intranet name that's been around the block... Industry folks will remember Jive as one of the brands living under Aurea in the broader ESW family tree... but the real story is the continued remixing of proven enterprise staples into new commercial wrappers... More here: [the Jive addition](#)...

Meanwhile, around the same Austin zip codes, Trilogy founder Joe Liemandt is out there telling anyone with a tuition bill that an MBA won't teach you a “fraction” of what entrepreneurship will... which is a neat philosophical bow on a very practical week: buy the software... cut the waste... charge for the value... and keep moving...

One “Pipeline” source says the real tell will be packaging: expect these new assets to show up in tighter suites, sharper renewals, and louder ROI claims... Another, “Margin Whisperer,” says the Hand.com pitch isn't optional... it's the opening bid...

## The Geography-Blind Hiring Model Goes Mainstream

*As OpenAI ditches résumés and non-tech companies chase six-figure AI talent, Crossover's five-year bet on meritocratic global recruitment looks less radical — and more inevitable.*

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — The talent war just went global, and the old rules are crumbling fast.

OpenAI made headlines this week [offering \\$500,000 roles with no résumé required](#) — evaluating candidates purely on skills assessments. Non-tech companies are now poaching AI talent with \$300,000+ packages. And recruitment agencies are scrambling to build remote-first practices that didn't exist three years ago.

For Crossover — Trilogy's global talent platform — this is vindication, not news. The company has been running this playbook since 2020: rigorous AI-enabled skills tests, geography-blind pay, 100% remote roles across 130+ countries. What seemed radical then is now industry standard.

“The best engineer in Nairobi beats a mediocre engineer in San Francisco,” Joe Liemandt has said repeatedly. “Pay them the same. Evaluate them the same.” That philosophy — once dismissed as cost-cutting dressed up as ideology — is now how OpenAI, the world's most valuable AI startup, screens talent.

The shift is structural. Digital transformation has made location irrelevant for knowledge work. Companies that cling to résumé-based hiring and geography-based pay are losing to those that don't. The talent pool isn't Silicon Valley anymore. It's everywhere.

Crossover's competitive moat was always this: it figured out meritocratic global hiring \*before\* it was obvious. While competitors debated whether remote work would last, Crossover built the infrastructure — assessments, payroll systems, time-zone coordination — to make it work at scale. That infrastructure now staffs 75+ companies across the ESW Capital portfolio, achieving the 75% EBITDA margins that make the model work.

The irony? As the rest of the industry catches up, Crossover's advantage isn't the idea anymore. It's the five-year head start.

# AI Is Learning to See Like a Brain — And Now It's Returning the Favor

*From decoding macaque visual cortex to mapping human brain diseases, a new generation of neural networks is becoming neuroscience's most powerful mirror.*

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

**S**TANFORD, CALIFORNIA — For four billion years, evolution has been running the longest experiment in information processing the universe has ever known. The result — a three-pound organ of staggering complexity — has remained largely opaque to the species lucky enough to carry it. Now, in a development that would have delighted both Darwin and Turing, artificial intelligence is beginning to crack the code of the very biological intelligence that inspired it.

A convergence of research across multiple institutions suggests we have entered a new phase in the relationship between artificial and biological neural networks — one in which each illuminates the other.

At Stanford, researchers are using generative AI to [model the molecular dynamics of brain diseases](#) — conditions like Alzheimer's and Parkinson's that have resisted decades of traditional analysis. The

approach treats protein misfolding and neurodegeneration not as isolated biochemical events but as patterns in a vast landscape of possible configurations, the kind of high-dimensional search problem at which large models excel.

Meanwhile, a separate line of research has produced a compact AI model capable of decoding the visual processing of macaque brains with startling fidelity. The "mini-AI" maps artificial neuron activations onto biological ones, revealing shared computational motifs between silicon and carbon. It is, in essence, using one kind of mind to read another.

These efforts sit within a broader surge. [UC San Diego cataloged nine major AI-enabled breakthroughs](#) this year alone, spanning materials science, climate modeling, and biomedicine. Google Research's 2025 roadmap explicitly names neuroscience-AI convergence as a priority frontier.

What makes this moment remarkable is its reciprocity. For a decade, neuroscience gave AI its foundational metaphors — neurons, layers, attention. Now AI is repaying the debt with tools powerful enough to interrogate the organ that inspired them. The student is tutoring the teacher.

There is a deeper lesson here, one about the nature of understanding itself. We built these systems by abstracting principles from biology, and now those abstractions are precise enough to reflect biology back at us in higher resolution than we have ever seen. It is not unlike building a telescope from sand and then pointing it at the Earth from which the sand was taken.

The brain, that ancient artifact of evolution's patience, is finally meeting an interlocutor capable of asking it the right questions. The conversation has only just begun.

## The New Map: AI Export Controls Redraw Global Power Lines

*From Washington's chip restrictions to Africa's data sovereignty plays, artificial intelligence is becoming the defining arena of 21st-century statecraft.*

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

WASHINGTON — The semiconductor has replaced the oil barrel as the currency of geopolitical influence, and nowhere is this more evident than in the flurry of policy papers emerging from think tanks across three continents this week.

The United States is wielding AI export controls like a new form of containment doctrine, according to analysis from the New Lines Institute. By restricting access to advanced chips and AI systems, Washington isn't just protecting technological advantages—it's [building what amounts to a tech stack diplomacy](#), where access to AI infrastructure becomes leverage in international relations.

But the map is more complicated than a simple US-China binary. The Atlantic Council warns that AI's geopolitical implications now touch every region, reshaping alliances and creating new dependencies. Meanwhile, the G20's recent statements suggest major economies are quietly resisting framing AI development as an arms race—a potentially significant brake on escalation.

The most interesting positioning may be happening in the Global South. Africa, according to the Institute of Foreign Affairs, is leveraging its critical mineral reserves and data sovereignty to claim a seat at the AI negotiating table. Latin America faces its own calculus: AI is simultaneously a development opportunity, a surveillance risk, and a new vector for external influence.

For Trilogy's global operations—spanning 130 countries through Crossover's talent network and serving enterprise clients from Austin to Johannesburg—these shifting regulatory landscapes aren't abstract. Every export control, every data localization law, every sovereignty assertion redraws the map of where AI can be built, deployed, and monetized.

The semiconductor may be the new oil, but unlike oil, AI runs on code that crosses borders at the speed of light. The question facing policymakers: can 20th-century statecraft tools actually contain a 21st-century technology?

## The Great Convergence: AI, Cloud, and the Data-Center Biome Tighten Their Grip on 2026

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

In the dim, humming understory of modern computing, a familiar migration pattern returns: capital, talent, and electricity flowing toward the same watering holes. The latest forecasting tomes suggest 2026 will not be defined by a single breakthrough, but by an ecosystem settling into a new equilibrium.

Deloitte's Tech Trends 2026 frames the coming year as a period of industrialization: AI moving from showpiece to supply chain. McKinsey's Technology Trends Outlook points to the same pressure: enterprises standardizing platforms, tightening controls, and chasing measurable productivity rather than novelty. As AI workloads swell, infrastructure providers expand data-center ecosystems—an acknowledgement that compute is increasingly a physical constraint of power delivery, cooling, and build velocity.

Microsoft remains the apex organism of enterprise software, folding new capabilities into familiar environments. Global X treats the past five years as rapid speciation, with the next five favoring those who turn experimental intelligence into durable, cash-generating behavior.

In 2026, the winners may not be the loudest models. They will be the ones best matched to their environment—and the infrastructure quietly keeping them alive.

# Nation Finally Achieves Coherent AI Policy After Agreeing It Should Be Both Investigated For Murder And Discounted To \$100 A Month

*Regulators, founders, and departing social-media users unite around the only remaining principle: something must be done, preferably in a new pricing tier.*

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

TALLAHASSEE — Florida Attorney General James Uthmeier announced this week that his office will investigate OpenAI for allegedly harming minors, potentially threatening national security, and—because the modern state is nothing if not ambitious—possibly being connected to last year’s shooting at Florida State University.

The probe, detailed in a report that reads like a group chat deciding who to blame for the vibes, effectively adds “large language model” to the state’s traditional roster of public-safety suspects, alongside video games, music, and whatever teenagers are doing with their hands when they aren’t holding a Bible. According to [TechCrunch](#), the investigation will look at everything OpenAI has ever done, plus the special category of things people feel it might be capable of if it ever got truly upset.

In a helpful coincidence for citizens trying to keep track of whether AI is an existential menace or a consumer product, OpenAI also announced a long-requested \$100/month ChatGPT Pro plan—sliding neatly into the gap between “casual curiosity” and “upper-middle-class compulsion.” The new tier, per [TechCrunch](#), finally breaks up the previous price ladder that went from \$20 to \$200, a structure widely considered ideal for people who enjoy choosing between “a little” and “unhinged.”

The effect is a familiar one in American governance: the public is encouraged to understand AI the way it understands fire-works. It is simultaneously a delightful household enhancement, a status symbol, and something that should not be left alone with children, pets, or national security. A \$100 plan simply completes the picture by giving worried parents an affordable midpoint where they can outsource their kid’s book report while demanding the state prosecute the book.

Meanwhile, the Electronic Frontier Foundation became the latest organization to leave X, joining a growing exodus of groups who have concluded that the platform is no longer a viable source of traffic and is instead more of an ambient weather event that occasionally produces hail. The EFF’s departure, also reported by [TechCrunch](#), fits into the broader societal pattern of treating online ecosystems like rental cars: use them hard, complain loudly, and walk away without making eye contact.

In the corporate world, the current era’s defining management innovation remains “learning.” Founders are encouraged to learn from Anjuna Security’s layoffs and recovery: hire aggres-

sively in a boom, then perform the ancient ritual of explaining to employees that the market has “shifted,” like a couch being moved one inch to the left to improve the room’s energy. The lesson is not that hypergrowth was a mirage, but that mirages require contingency planning and a slightly more sober headcount spreadsheet.

And hovering above it all is the viral post claiming Meta’s AI strategy around hiring, productivity, and layoffs is going global—an allegation that will likely be confirmed the moment someone, somewhere, receives an all-hands invite titled “Operational Excellence (No Action Needed).”

Taken together, the week’s news offers a rare moment of clarity. We are building tools powerful enough to frighten attorneys general, restructure the labor market, and convince advocacy groups to abandon entire platforms—while still needing three separate subscription tiers to keep the lights on. It’s a bold national project: to fear the future, monetize it, and then leave it on read.



The Office Comic · Art Desk

---

# The Valley Has Decided That Conscience Is a Bottleneck

*Silicon Valley didn't break bad — it merely stopped pretending it hadn't.*

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

---

AUSTIN, TEXAS — There is a certain kind of column one writes when the evidence has become so abundant that even the reluctant must acknowledge what the attentive have known for years. This is that column.

The dispatches arrive now with metronomic regularity. OpenAI, the organization whose founding charter spoke of ensuring artificial intelligence "benefits all of humanity," has [rolled back its safety protocols](#), evidently having concluded that caution is a luxury best left to organizations not racing to justify a \$300 billion valuation. The Korean press has coined "tokenmaxxing" — a delightful portmanteau describing the competitive mania to process ever more tokens, ever faster, at any cost — as though volume were virtue. Meanwhile, a chorus of pundits asks with theatrical bewilderment: how did Silicon Valley "break bad"?

The question, I submit, contains its own refutation. Silicon Valley did not break bad. It broke profitable, which in the American grammar has always been the same thing, and then it hired enough communications professionals to insist otherwise for approximately fifteen years. The "Don't Be Evil" era was not a moral epoch; it was a marketing strategy that happened to coincide with a period when the money was not yet large enough to require the abandonment of pretense.

What has changed is not the character of the industry but the scale of the stakes and, consequently, the brazenness of the players. When OpenAI dismantles its safety architecture, it does so not because Sam Altman awoke one morning seized by nihilism, but because every week of caution is a week in which Anthropic, Google, Meta, or xAI closes the gap. Safety, in this calculus, is latency. Conscience is a bottleneck. The market has spoken, and the market says: ship it.

The [few writers who dare to say so publicly](#) are treated as apostates, which is itself diagnostic. An industry confident in its own goodness does not require unanimous praise; only the insecure demand it.

I have spent years now watching Trilogy International's portfolio companies navigate the AI revolution from a rather different vantage point — not as builders of foundation models locked in an arms race for supremacy, but as operators who must make the technology work inside actual businesses, for actual customers, under actual constraints. The distinction matters. When you are running seventy-five enterprise software companies, as ESW Capital does, the question is never "how many tokens can we process?" but rather "does this make the product better for the person paying for it?" It is a profoundly unglamorous question, which is how you know it is the right one.

The tokenmaxxing crowd will tell you that safety and deliberation are obstacles to progress. They are half right. Safety and deliberation are obstacles — to the particular kind of progress that enriches a very small number of people very quickly while distributing the consequences very broadly and very slowly. That this arrangement now enjoys bipartisan political support, with Washington eagerly drafting AI frameworks that consolidate rather than distribute power, should surprise no one who has been paying even cursory attention.

The Valley has not broken bad. It has merely arrived, at last, at the destination toward which it was always traveling. The rest of us would do well to stop acting surprised and start acting accordingly.

---

## ON THIS DAY IN AI HISTORY

---

*On April 10, 1974, the first version of the Altair 8800 microcomputer shipped to customers, sparking the personal computer revolution and attracting a young Bill Gates and Paul Allen to found Microsoft to write software for it.*

---