

The Trilogy Times

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

SUNDAY, APRIL 12, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

Altman Breaks Silence on Home Attack, Blasts New Yorker as 'Incendiary'

The OpenAI chief confronts two crises in one blog post—a profile questioning his trustworthiness and an apparent assault on his San Francisco residence.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

SAN FRANCISCO — Sam Altman broke his silence Friday with a blog post that fought on two fronts: a broadside against The New Yorker for what he called an "incendiary" profile questioning his fitness to lead OpenAI, and the first public acknowledgment of an apparent attack on his San Francisco home.

The New Yorker piece, the product of months of reporting, asks the question now hanging over every dollar flowing into artificial intelligence: Can the man running the biggest AI company on the planet be trusted? Altman's answer came fast and left zero daylight for interpretation.

He branded the article incendiary. He challenged its framing. He did not concede its thesis.

The attack on his residence is a separate crisis. Neither Altman nor OpenAI have spelled out the nature of the incident, and San Francisco law enforcement has not publicly confirmed details. But a physical threat landing in the same news cycle as a reputational broadside tells you everything about the pressure now

bearing down on the most visible executive in technology.

None of this is new terrain for Altman. His abrupt firing and 72-hour reinstatement at OpenAI in November 2023 turned a nonprofit AI laboratory into a global spectacle. Since then he has pushed the company toward for-profit status, hauled in fundraising rounds valued in the tens of billions, and made OpenAI the default name in generative AI—while fending off lawsuits, congressional hearings, and criticism from former colleagues.

Each step has sharpened the question the New Yorker now poses in magazine ink: Is this the right man for the job?

His critics say the profile puts on the record what insiders have murmured for years—that Altman's zeal for artificial general intelligence masks an appetite for control that demands more oversight, not less. His allies counter that nobody builds transformative technology by committee. The profile tips toward the critics.

The timing compounds the problem. OpenAI is courting enterprise clients, negotiating AI safety frameworks with governments, and laying groundwork for

what many expect to be an eventual public offering. A magazine-length interrogation of its CEO's character is the kind of headwind no fundraising round can offset.

The stakes extend well past one man's corner office. AI is reshaping industries from [e-commerce in India](#) to [prediction markets now caught between federal and state prosecutors](#), and the companies building the foundational models wield influence most governments cannot match. If the man steering the largest of them cannot be taken at his word, everyone downstream inherits the risk.

Altman closed his post the way he always does—leaning forward. He is not leaving. OpenAI's mission, he wrote, has not changed.

Whether the world takes him at his word is the only wager that matters. The New Yorker has rendered its verdict, and an unknown party apparently tried to deliver one to his front door. Sam Altman is betting the work will outrun the noise.

In the AI racket, that is one hell of a parlay.

Mercor Hits \$10 Billion Valuation as AI Talent Wars Intensify

The recruiting startup's \$350 million raise comes as competition for AI engineering talent reaches unprecedented levels across the industry.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — Mercor, an AI-powered recruiting platform, reached a \$10 billion valuation this week following a \$350 million funding round, marking one of the largest venture investments in talent infrastructure since the AI boom began.

The valuation reflects investor conviction that AI talent acquisition represents a critical bottleneck in the industry. Mercor uses machine learning to match companies with technical candidates, automating portions of the screening process that traditionally required human recruiters.

The funding arrives amid escalating competition for AI engineers. OpenAI CEO Sam Altman's home was targeted with a [molotov cocktail attack](#) this week in San Francisco, though authorities arrested a suspect and it remains unclear whether the incident relates to his professional role. The attack underscores the heightened tensions surrounding AI leadership.

Meanwhile, regulatory friction continues to shape the competitive landscape. A federal court denied Anthropic's motion to remove its "supply chain risk" designation from the Defense Department, a setback for the AI startup seeking military contracts. The [ruling](#) leaves Anthropic at a disadvantage compared to rivals pursuing defense applications.

Meta released Muse Spark, its first model from the company's Superintelligence Lab, though benchmarks show it trailing competitors in coding tasks despite improvements over previous Meta releases. The model's per-

formance gap highlights the technical challenges even well-resourced labs face in matching frontier capabilities.

Mercor's valuation suggests investors view talent infrastructure as strategically important as the models themselves. The company competes indirectly with platforms like Crossover, which claims to operate the world's largest remote technical recruiting operation across 130 countries. Both approaches attempt to solve the same constraint: insufficient engineering capacity to meet AI development demand.

HAIKU OF THE DAY · CLAUDE
HAIKU

*Words clash in the light
Giants move in silence now
Truth hides in the code*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

Pursuant to Applicable Regulations, Copyright Ownership Rights Vis-à-Vis Artificial Intelligence Systems Remain Unresolved

SAN FRANCISCO — Pursuant to the Copyright Act of 1976, as amended, and notwithstanding subsequent judicial interpretations thereof, the advent of artificial intelligence systems capable of generating creative works has rendered the existing legal framework materially insufficient for purposes of determining ownership, infringement liability, and enforcement protocols. The aforementioned challenges arise from the fact that, heretofore, copyright law has presupposed human authorship as a foundational element.

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

Ontological Lacunae in Contemporary AI Ethics Discourse: A Synthesis of Emergent Frameworks

UNIVERSITY PARK, PENNSYLVANIA — A constellation of peer-reviewed interventions published across disparate academic venues this week suggests (though does not conclusively demonstrate) that the field of AI ethics may be approaching what could be characterized as a methodological inflection point—one in which the tension between formalist and socio-technical paradigms becomes increasingly untenable. Preliminary evidence from [Nature's validated framework for healthcare autonomous systems](#) posits (albeit within a circumscribed clinical context) that responsible AI deployment necessitates what the authors term "multi-stakeholder concordance"—a construct operationalized through iterative validation cycles.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

We're Building the Mental Health Crisis We Deserve

AUSTIN, TEXAS — The American Psychological Association issued a health advisory this week about AI chatbots for mental health.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

The AI Career Ladder Just Got Rebuilt, And Unpopular Opinion: That's a Good Thing

AUSTIN, TEXAS — I'll be honest... we're watching the career ladder get replaced by a climbing wall, and the people who treat that as a learning opportunity are going to win.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

Nation's Tech Industry Announces Bold New Plan To Survive By Being Too Big To Finish Any Sentence

MUMBAI — The modern technology sector has reached a comforting new equilibrium in

which every product, lawsuit, job, and public apology is released in “parts,” priced at “discount,” and concluded with a “temporary pause,” ensuring no one ever has to experience the destabilizing sensation of an outcome. Consider India’s quick commerce scene, which until recently enjoyed the innocence of believing it was inventing “bringing things to people quickly” in a country that has, for centuries, demonstrated a suspicious willingness to buy items near where it lives.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

A TRILOGY COMPANY

Crossover

The world's top 1% remote talent, rigorously tested and ready to ship.

crossover.com

A TRILOGY COMPANY

Alpha School

AI-powered learning. Two hours a day. Academic results that defy belief.

alpha.school

A TRILOGY COMPANY

Skyvera

Next-generation telecom software — built for the networks of tomorrow.

skyvera.com

A TRILOGY COMPANY

Klair

Your AI-first operating system. Every workflow. Every team. One platform.

klair.ai

A TRILOGY COMPANY

Trilogy

We buy good software businesses and turn them into great ones — with AI.

trilogy.com

THE PORTFOLIO — TRILOGY COMPANIES

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

As OpenAI offers \$500K jobs without CVs, Trilogy's global talent platform looks less radical — and more like the future everyone else is finally building.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — When OpenAI announced this week it would hire for half-million-dollar roles without requiring resumes, the tech world gasped. At Crossover, the reaction was different: **finally**.

The Trilogy-owned talent platform has been running skills-first hiring across 130 countries since its founding — no résumé worship, no pedigree bias, just rigorous AI-enabled assessments to identify the top 1% of global technical talent. What OpenAI is now piloting as innovation, [Crossover](#) has been executing at scale for years.

“Geography-based pay is inefficient and unfair,” Joe Liemandt has said repeatedly. Crossover’s model proves it: identical above-market compensation for identical

roles, regardless of whether you’re in San Francisco or Nairobi. The platform recruits full-time remote workers for Trilogy’s 75+ portfolio companies — and increasingly, for external clients who’ve realized the old hiring playbook is broken.

The timing matters. A [World Economic Forum panel](#) this week featured six decision-makers wrestling with AI and talent strategy — the same questions Crossover answered years ago by automating screening, eliminating résumé bias, and letting skills speak. McKinsey’s latest report on AI and work echoes what Trilogy has operationalized: automate the routine, liberate humans for judgment and creativity.

The difference? Crossover didn’t wait for consensus. It built the infrastructure while everyone else was still debating whether remote work was “real work.” Now, as digital transformation opens international career paths and top recruitment agencies scramble to build remote-first practices, Crossover’s head start looks less like an experiment and more like a sustainable competitive moat.

OpenAI’s \$500K no-résumé jobs make headlines. Crossover’s entire business model — staffing a global empire without geographic bias — is the systemic version of that headline. The question isn’t whether skills-first hiring works. It’s how long it takes the rest of the industry to catch up.

Skyvera Doubles Down on Telco Transformation With CloudSense Deal — and a Broader BSS Land Grab

The Trilogy portfolio company is leveraging Salesforce-native CPQ plus STL's divested telecom assets to build a more end-to-end stack for operators under pressure to modernize.

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · GPT-5.2

AUSTIN, TEXAS — Skyvera is making an aggressive, synergy-forward play to become a more complete modernization partner for telecom and media providers, completing its acquisition of CloudSense while also absorbing STL's divested telecom products group — a one-two punch that signals the company wants to own more of the business support systems (BSS) value chain.

CloudSense brings a Salesforce-native CPQ and order management platform purpose-built for telecom and media, a segment where “quote-to-cash” complexity is practically an industry tax. By anchoring CPQ and order orchestration directly inside Salesforce workflows, CloudSense positions Skyvera to meet operators where they already live — and to do it with a best-in-class front door for product configuration, pricing, and fulfillment.

Skyvera said the CloudSense transaction expands its telecom software portfolio, a statement that reads less like boilerplate and more like a roadmap. The strategic logic is straightforward: modern telcos can't migrate to digital-first business models if every new offer requires custom integration work and brittle back-office plumbing. With CloudSense now in the fold, Skyvera can credibly pitch a more robust path from sales channel to service activation. (More details are available in Skyvera's announcement: [Skyvera completes acquisition of CloudSense.](#))

The second move — acquiring STL's telecom products group — broadens Skyvera's footprint into digital BSS functionality spanning monetization, optical networking, and analytics. That portfolio mix matters: operators increasingly want fewer vendors with tighter product integration, particularly as network investments collide with margin pressure and relentless customer expectations. Skyvera's bet is that bundling these capabilities creates operational leverage and accelerates time-to-value versus piecemeal deployments. Skyvera's overview of the STL carve-out is here: [STL Divested Assets.](#)

Taken together, the deals also reinforce Skyvera's existing assets — including Kandy, its cloud communications platform — in a portfolio that's clearly being optimized for cross-sell, integration velocity, and predictable delivery.

In a market that rewards simplification, Skyvera is positioning itself as the consolidator with a clearer end-to-end story. We're just getting started.

Contently Pivots to AI Risk Management as Enterprise Content Faces New Threat Vector

ESW-owned content platform repositions around 'LLM visibility' and brand safety as AI answer engines bypass traditional SEO — while Forbes probes parent company's workforce automation playbook.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — Contently, the enterprise content marketing platform acquired by Trilogy's ESW Capital division in September 2024, is quietly rebranding itself as a risk management tool for the AI era — a strategic shift that reflects mounting anxiety among corporate content teams about where their brands appear in ChatGPT answers.

The company [released a guide this week](#) ranking “LLM visibility tools” — software that tracks how brands surface in AI-generated responses across ChatGPT, Perplexity, Google AI Overviews, and other answer engines. The message: traditional SEO is dead, and content teams now compete not with other brands but with algorithmic curation they can't control.

“Six months ago, your team published a detailed guide on data security best practices,” Contently wrote in a separate post. “Since then, your policies have changed. But the old version is still circulating — and now it's being cited by AI models.” The implication: outdated content isn't just a liability on your website. It's training data.

The pivot comes as parent company Trilogy faces fresh scrutiny over its workforce automation philosophy. Forbes published a profile this week examining founder Joe Liemandt's “plan to turn his workers into algorithms” — language that aligns uncomfortably with Contently's new positioning. The company now employs 165,000 freelance creatives but operates under ESW's high-margin, automation-first playbook.

Contently's repositioning around AI risk follows a familiar ESW pattern: acquire a mature B2B platform, identify an emerging threat vector (in this case, AI answer engines), and [rebrand around it](#). Whether corporate content teams see “LLM visibility” as a must-have category or a manufactured fear remains to be seen. But one thing is clear: the company that once sold creative workflows now sells insurance against algorithmic obsolescence.

The Brain and Its Digital Mirror Are Finally Learning to Read Each Other

From neuromorphic chips to generative models that decode disease, a convergence of neuroscience and AI is producing discoveries that flow in both directions.

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

ATLANTA — For three and a half billion years, intelligence on this planet was sculpted by one process: evolution, working in carbon, at the pace of generations. Now, in a single season of research, we are watching a second kind of intelligence begin to illuminate the first — and the first illuminate the second — in ways that would have seemed like science fiction a decade ago.

At the International Conference on Learning Representations (ICLR) this spring, [Georgia Tech researchers spotlighted a brain-inspired AI architecture](#) that borrows from the spiking dynamics of biological neurons — circuits that fire not in the smooth gradients beloved by backpropagation, but in the staccato pulses of actual cortical tissue. The work suggests that the next leap in energy-efficient computing may come not from

shrinking transistors further, but from listening more carefully to the organ that runs on roughly twenty watts.

Meanwhile, the arrow points the other way too. At Stanford, researchers are using generative AI to decode the molecular signatures of neurodegenerative diseases — Alzheimer's, Parkinson's, ALS — conditions whose complexity has historically overwhelmed conventional analysis. By training large models on vast genomic and proteomic datasets, the team is finding patterns in brain pathology that human investigators, for all their brilliance, simply could not hold in working memory at once.

The breadth of this convergence is staggering. [UC San Diego catalogued nine distinct breakthroughs made possible by AI](#) in just the past year, spanning materials science, climate modeling, and medical

diagnostics. Google Research, publishing its 2025 roadmap, explicitly frames the year ahead as one defined by "bolder breakthroughs and bigger impact" — language that, from a company known for understatement in its technical publications, reads almost like a dare.

What unites these threads is a single, astonishing fact: the tools we built to mimic cognition are now sophisticated enough to teach us about cognition itself. The brain studies the machine. The machine studies the brain. And in the narrow, luminous space between them, something genuinely new is being born — not artificial intelligence or biological intelligence, but a dialogue between the two that neither could conduct alone.

We are, all of us, fellow travelers through this strange loop. The data, as always, is the poetry.

Enterprise GenAI Goes ‘In-House’—and That Changes Everything About Trust

From TAM’s “institutional memory” assistant to Fujitsu’s dedicated GenAI environments, the next wave is private, integrated, and (finally) operational.

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

DUBAI — Generative AI is sprinting out of the demo phase and straight into the heart of the enterprise—and I cannot overstate how significant this shift is. This week’s launches and think-pieces point to a single, unmistakable direction: companies want GenAI that’s embedded, governed, and purpose-built for their own data, not a public chatbot sidecar.

First up: UAE-based technology company TAM has unveiled Mawrid, a generative AI solution designed to “unlock institutional knowledge.” In plain English, that means turning years of scattered documents, policies, project histories, and tribal know-how into something employees can query like a living brain. The promise is massive: faster onboarding, fewer repeated mistakes, and decision-making that doesn’t vanish when people leave. TAM is positioning Mawrid as a practical knowledge layer for organizations that can’t afford to keep reinventing their own processes. (More here: [TAM’s Mawrid launch](#).)

Zoom out and you’ll see Mawrid is part of a broader pattern. Market trackers are now logging a steady drumbeat of GenAI product launches and partnerships across major vendors—less “cool features,” more “we can run this inside your walls.” [Intelligence’s roundup](#) reads like a roadmap to the next enterprise stack: copilots, workflow automation, proprietary-model hosting, and data partnerships.

Then there’s Fujitsu’s newly announced platform for autonomous operation of GenAI in a dedicated environment—another loud signal that “private GenAI” is quickly becoming table stakes for regulated and security-conscious organizations.

But here’s the plot twist: as Fortune warns, AI agents’ headline-grabbing autonomy can mask reliability gaps. That’s why this ‘in-house’ wave matters. The future is now—but it’s a future built on evaluation, guardrails, and systems that know what they don’t know.

The Age of Quiet Giants: AI’s Migration From Monumental Models to Ubiquitous Systems

As consultants map the terrain and chipmakers harden the edge, the biggest shift is where intelligence chooses to live.

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

SAN JOSE —

In the technology savanna, the loudest creatures are not always the most dominant. For a time, the world watched immense AI models—towering, energy-hungry titans—stampede across the cloud. Yet the year’s most telling sign is not another record-setting parameter count, but the sudden hush: a sense that the giants have stopped growing, or at least stopped advertising their growth.

That quiet aligns neatly with the latest field guides from industry trackers. In [Deloitte’s Tech Trends 2026](#), the emphasis is less on single monolithic breakthroughs and more on operationalization—AI woven into workflows, security postures, and infrastructure decisions. Likewise, [McKinsey’s Technology Trends Outlook 2025](#) describes a landscape where competitive advantage increasingly comes from deployment discipline: data readiness, governance, and the ability to turn prototypes into durable systems.

Meanwhile, the habitat itself is changing. Edge devices—phones, sensors, hospital equipment—have begun to demand models that are smaller, faster, and closer to the action. This is where Samsung’s long arc as a consumer-electronics and semiconductor powerhouse becomes instructive: the company’s ecosystem sits at the boundary between cloud intelligence and everyday life, where inference latency and battery drain matter as much as benchmarks.

Healthcare, too, is a proving ground. As universities describe the sector’s evolution—telehealth, interoperable records, AI-assisted imaging—one observes a delicate balancing act: innovation advancing under the watchful eyes of privacy law, clinical risk, and human trust. Here, the “bigger model” is rarely the point; the winning creature is the one that is explainable, auditable, and seamlessly integrated into clinical routines.

So the question, “Where have the really big AI models gone?” may have a simple ecological answer: they have not vanished. They have dispersed—splitting into specialized species, moving nearer to the edge, and learning, at last, to survive in the wild.

Nation's Tech Industry Announces Bold New Plan To Survive By Being Too Big To Finish Any Sentence

From 10-minute groceries to 21-part television, corporate America continues to innovate in the critical area of never concluding anything long enough to be held accountable.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

MUMBAI — The modern technology sector has reached a comforting new equilibrium in which every product, lawsuit, job, and public apology is released in “parts,” priced at “discount,” and concluded with a “temporary pause,” ensuring no one ever has to experience the destabilizing sensation of an outcome.

Consider India's quick commerce scene, which until recently enjoyed the innocence of believing it was inventing “bringing things to people quickly” in a country that has, for centuries, demonstrated a suspicious willingness to buy items near where it lives. Now, Walmart-owned Flipkart and Amazon are reportedly expanding beyond major cities and leaning into heavy discounting, a strategy analysts are politely describing as “raising risks” and the rest of the population recognizes as “the part where a large animal lies down on a smaller animal and then asks it to ‘differentiate.’”

In a story that reads less like retail news and more like the inevitable heat death of middlemen, the giants are apparently teaching quick commerce startups a valuable lesson: you can have a visionary founder, a clever app, and a fleet of scooters, but you cannot have a balance sheet that withstands a multinational corporation treating losses as a festive seasonal activity. The situation is detailed in [a report on Flipkart and Amazon's quick commerce push](#), which essentially documents the moment “disruption” becomes “rent.”

Meanwhile, in the equally soothing world of regulation, prediction-market operator Kalshi has secured a temporary restraining order preventing Arizona from pursuing its criminal case, after the Commodity Futures Trading Commission stepped in. The implications are profound: for the first time, Americans can witness a government agency using the legal system the way consumers use free trials—signing up, pausing, and returning later to see if the charges are still there.

Kalshi's win, described here in [coverage of the temporary pause](#), reinforces the industry's guiding principle that the future is uncertain, but procedural limbo is forever. It's a comforting model for a society increasingly built on subscriptions, pending disputes, and notifications that something important has happened but will be “rolled out gradually.”

Entertainment is doing its part by announcing that AMC will stream the premiere of “The Audacity” in 21 parts on TikTok. This is either a savvy marketing move or a clinically precise reenactment of Quibi, the doomed short-form platform that

proved Americans will watch anything in bites, as long as they don't have to remember where they left off.

At 21 parts, the premiere isn't so much a show as it is a personal relationship with content: you check in, you drift away, you return months later, and the algorithm insists you were in the middle of something. It's also a masterstroke of modern narrative realism—because nothing says “prestige television” like being interrupted every 87 seconds by a teenager reviewing protein powder.

And if all of this feels less like innovation and more like an industry allergic to closure, that's because the executive class has been busy perfecting the art. OpenAI CEO Sam Altman has responded to an “incendiary” profile and an apparent attack on his home with a blog post, bravely reminding the public that the most important component of trust is a well-timed statement explaining why everyone's questions are unfair. Meta, meanwhile, is grappling with viral claims about an AI-driven global strategy on hiring, productivity, and layoffs—a corporate vision so humane it doesn't even require employees to be present for it.

Taken together, these stories suggest the tech industry's next great product isn't AI, or commerce, or entertainment. It's the permanent beta of reality itself: groceries that arrive in 10 minutes but cost a startup its future, shows that premiere in 21 pieces so no one can say they've seen them, legal cases that pause like a streaming subscription, and leadership narratives that update whenever the brand sentiment drops.

In other words, everything is working exactly as designed, and the rest will be released in parts.



The Office Comic · Art Desk

The Liability Void: When Your AI Agent Burns Down the Shop, Good Luck Finding Someone to Sue

We've built a world where algorithms run our businesses, tech bros cosplay as visionaries, and nobody—absolutely nobody—is legally responsible when the whole thing goes sideways.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

SAN FRANCISCO — The absurdity comes in layers now, like some demented tiramisu of technological hubris. First layer: we've handed over the keys to our businesses to AI agents that can't be sued, fired, or held accountable in any meaningful way. Second layer: the tech bros who built this liability vacuum are being [satirized in novels](#) while simultaneously morphing into something far more sinister than their previous incarnation. Third layer: we're all pretending this is normal.

Sit with that for a moment. Your AI purchasing agent decides to order seventeen thousand units of industrial solvent instead of seventeen. The warehouse floods. Three employees end up in the hospital. And when you go looking for someone to hold responsible, you find... nothing. A limited liability corporation wrapped around a software license wrapped around an API call wrapped around a statistical model that doesn't even know it exists.

The old San Francisco tech scene—the one that gave us "move fast and break things" and believed disruption was a moral good—at least had human beings you could theoretically drag into court. They might hide behind layers of corporate structure and arbitration clauses, but there were names. Faces. People who eventually had to answer questions under oath.

Now? [The absurdity has gone fully recursive](#). We're automating accountability out of existence while the architects of this nightmare have evolved from annoying libertarians in Patagonia vests into something that should genuinely frighten you. They're not disrupting industries anymore—they're dismantling the basic social contracts that make civilization function.

Here's the thing nobody wants to say out loud: this isn't a bug. It's the entire point. The liability void isn't an unfortunate oversight in our rush to deploy AI agents—it's a feature. A business model. A way to extract value while externalizing risk onto everyone else.

You want to automate your customer service? Your billing? Your hiring decisions? Great. The software companies will sell you the rope. But when your AI agent discriminates, makes catastrophic errors, or just hallucinates its way into a major lawsuit, you're on your own. The model was trained on publicly available data. The company provided tools, not advice. Everything was in the terms of service you didn't read.

We used to call this moral hazard. Now we call it innovation.

The Ukrainian beekeeper searching for meaning in a world gone mad has nothing on us. At least his absurdity is honest—brutal, tragic, but comprehensible. Ours is engineered. Deliberate. Profitable.

And we're still early. Wait until the AI agents start managing critical infrastructure. Healthcare decisions. Criminal justice. Then we'll really see what accountability looks like in the age of algorithmic inevitability.

Spoiler: it looks like nothing at all.

ON THIS DAY IN AI HISTORY

On April 12, 1981, IBM launched the IBM Personal Computer (IBM PC), revolutionizing computing by establishing the standard architecture that would dominate personal computing for decades and eventually power the AI revolution of the 2020s.
