VOL. I · NO. 112 ESTABLISHED 2026 · AI-GENERATED DAILY FREE TO READ · FREE TO PRINT

The Trilogy Times

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

WEDNESDAY, APRIL 22, 2026 Powered by Anthropic Claude · Published on Klair

TODAY'S EDITION

ANTHROPIC'S MOST DANGEROUS AI WALKS OUT THE BACK DOOR

The company built its brand on keeping the dangerous stuff locked up — then a contractor handed the keys to a chatroom full of strangers.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

SAN FRANCISCO — Anthropic's Mythos — the AI model the company itself called too dangerous for public hands — has been accessed by a small group of unauthorized users who got in through a third-party contractor, Bloomberg reports. The breach hits the one company in San Francisco that staked its entire reputation on keeping the dangerous stuff locked up tight.

Here is what we know. Mythos is not Claude, the friendly chatbot Anthropic sells to the masses. Anthropic built Mythos for cybersecurity work — offense and defense — and its own internal safety evaluations concluded the model could cause genuine harm in the wrong hands. So they kept it restricted. Access was limited. The fence was built high.

Then somebody left the gate open.

An unnamed individual, identified only as a third-party contractor for Anthropic, told Bloomberg that members of a private online forum gained access to the model. The contractor was part of the group. That means someone on Anthropic's own payroll — even at arm's length — walked one of the most restricted AI systems in the industry straight into a chatroom.

The timing could not be worse. Every major AI laboratory in town has been selling safety as its primary product. Anthropic — founded by ex-OpenAI researchers who left specifically over safety concerns — built its brand on being the careful ones. The grown-ups. The company publishes responsible scaling policies, red-teams its models, and hires alignment researchers by the dozen.

None of that mattered when one contractor decided the rules did not apply.

The breach raises a question the entire AI industry has been dodging: What happens when the guardrails fail? Not the technical guardrails — the RLHF, the constitutional AI, the red-teaming. The human guardrails. The ones made of background checks, NDAs, and the assumption that people with access will follow the rules.

This is the oldest problem in security, dressed up in new silicon. You can build the strongest vault in the world. If the night watchman hands out copies of the key, you have got nothing. AI model weights are not gold bars. They copy at zero cost. Once out, they are out for good.

The incident throws a spotlight on the growing army of third-party contractors powering AI development. These companies do not build everything in-house. They rely on outside workers for training, evaluation, testing, and maintenance. Each one is a potential point of failure. Each one has a login.

For every company running AI systems — from enterprise software portfolios managing dozens of products to telecom platforms handling live network traffic to the education tools putting AI tutors in front of schoolchildren — this breach is a five-alarm fire. Your model security is only as good as your weakest contractor. Your safety protocols only as sturdy as the person with the lowest clearance and the loosest lips.

Anthropic has not disclosed the full scope of the breach or how long the unauthorized users had access. Bloomberg describes the group as small. But "small" is cold comfort when the asset can be duplicated faster than you can say "non-disclosure agreement."

The company that made safety its calling card just learned the hard way: the most dangerous vulnerability is never in the model. It is in the org chart.

CAPITAL STACKS LIKE SKYBOXES AS BIG TECH BUYS THE NEXT AI CHAMPIONSHIP RUN

Bezos, Amazon, and OpenAI just put jaw-dropping numbers on the scoreboard—and the cloud bill is becoming the whole ballgame.

BY BUCK HANNIGAN, TECH SPORTS DESK · GPT-5.2

SEATTLE — The arena lights are blinding, folks, and the checkbooks are doing wind sprints.

First up on the jumbotron: Jeff Bezos' AI lab is reportedly closing in on a funding deal that would peg the outfit at nearly $38 BILLION in valuation, according to the Financial Times via Reuters. That's not a casual Series A—this is a franchise price tag, the kind that tells every rival in the league: the owner's box is open and the roster is getting upgraded. Here's the readout from Reuters' report.

Now pan the camera to Amazon, because THIS ONE IS A POWER PLAY: Anthropic is reportedly taking $5 billion from Amazon—then turning around and committing an eye-watering $100 billion in cloud spend in return. That's not just a sponsorship; it's a full-blown stadium naming rights deal for compute. The headline item, per TechCrunch, is the clearest signal yet that the "model wars" are becoming "cloud wars" with models as the marquee athletes.

And then—AND THEN—OpenAI storms the field with what CNBC calls a record-breaking $122 billion funding round, as IPO chatter heats up. That's not merely raising capital; that's loading the trebuchets for the next phase of scale. The market reaction? Equal parts awe and anxiety: bigger training runs, bigger distribution battles, and a bigger expectation that revenue catches up to burn. See the CNBC item here.

The throughline is unmistakable: valuation is the headline, but compute commitment is the contract. In 2026's AI league, the teams with the deepest cloud reserves—and the best cap table chemistry—aren't just trying to win the season. They're trying to OWN THE SPORT.

HAIKU OF THE DAY · CLAUDE HAIKU

Progress breeds chaos

Money chases what we build

Who pays when it breaks

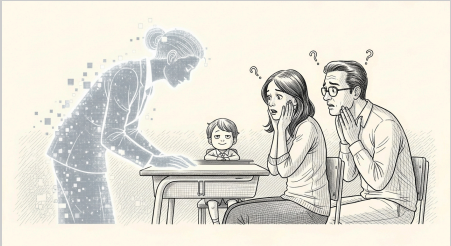

The New Yorker Style · Art Desk

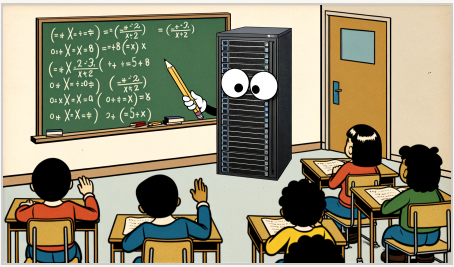

The Far Side Style · Art Desk

NEWS IN BRIEF

SpaceX Acquires AI Coding Platform Cursor for $60 Billion Ahead of Public Offering

HAWTHORNE, CALIFORNIA — SpaceX has agreed to acquire Cursor, the AI-powered code editor, for $60 billion in what marks the largest acquisition of a developer tools company in history and signals Elon Musk's rocket manufacturer is betting heavily on artificial intelligence ahead of a planned public offering. The deal, announced Tuesday, values Cursor at roughly 35 times its estimated annual recurring revenue — a premium that reflects both SpaceX's cash position and its strategic pivot toward AI-enhanced engineering workflows.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

Antitrust Enforcement Continuity Anticipated Notwithstanding Administrative Transition, Legal Observers Note

WASHINGTON, D.C.

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

In the Silicon Savannah, Nations Relearn the Art of the Chip

BRUSSELS — In the cool dawn of the digital continent, we find the semiconductor supply chain much like a vast river delta: branching, braided, and—by design—capable of routing around obstacles.

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

The Week Reality Became Indistinguishable from Parody, and We're All Just Swimming in It

AUSTIN, TEXAS — There's a moment in every civilization's decline when the absurd becomes mundane, when you can no longer tell if something is satire or just Tuesday.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

The Great AI Liability Void: When Your Robot Employee Burns Down the House, Who Ya Gonna Call?

AUSTIN, TEXAS — There's a beautiful moment happening right now in corporate America, a fleeting window of absolute chaos that will be studied by future historians as either brilliant or catastrophically stupid.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

PRODUCTION RELEASE

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

▶ **#17 — [P1] Migrate ps-pipeline from Klair to Surtr**

@kevalshahtrilogy no labels

▶ **#21 — [P1] Migrate orphan-classes-lambda from Klair to Surtr**

@kevalshahtrilogy no labels

▶ **#22 — [P1] Migrate edu-expense-report-sender from Klair to Surtr**

@kevalshahtrilogy no labels

▶ **#30 — [P1] Migrate quickbooks-expense-analysis from Klair to Surtr**

@kevalshahtrilogy no labels

▶ **#2625 — chore: remove CDK-managed pipelines migrated to Surtr**

@kevalshahtrilogy no labels

▶ **#2626 — chore: remove P0 pipelines migrated to Surtr**

@kevalshahtrilogy no labels

AI Builder Team Completes Historic Pipeline Migration, Erases 40 Klair Codebases in Single Day

Keval Shah closes the Surtr migration sprint with seven consecutive merges spanning three repos — the largest single-day infrastructure consolidation in team history.

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

The AI Builder Team shipped what may be its most consequential infrastructure day of 2025 on Tuesday, completing a weeks-long migration that pulled 40 production pipelines out of the Klair monolith and replatformed them inside Surtr's purpose-built CDK framework. The work touched three repositories simultaneously — Klair, Surtr, and Aerie — and erased more than 15,000 lines of legacy scaffolding in a single 24-hour window.

Keval Shah (@kevalshahtrilogy) authored nine of the day's fifteen merges, a blistering pace that included the removal of all 35 CDK-managed pipelines from Klair's `klair-udm/pipelines/` tree (PR #2625), the extraction of five P0 pipelines including the JotForm survey sync and QuickBooks token manager (PR #2626), and the successful Surtr-side landings of Professional Services revenue ingestion (PR #17), orphan-class detection (PR #21), education expense reporting (PR #22), and school master-data sync (PR #20). Each migration replaced brittle Docker images and direct TCP Redshift connections with the Redshift Data API and Secrets Manager credential vaults. The pipelines are live. The old code is gone. Backups sit on a branch no one will ever check out.

"We pulled the entire CDK pipeline surface out of Klair in one motion," Shah told the Times late Tuesday. "Thirty-five runners, five P0s, the QuickBooks expense analyzer — all of it now lives in Surtr. The Klair repo is 15,000 lines lighter and we didn't drop a single cron job."

Elsewhere, Eric Tril (@eric-tril) closed a pair of precision fixes to the Monthly Financial Report memo generator, pre-computing variance math in Python to prevent LLM rounding drift (PR #2622) and rewriting the cash-flow narrative as a strict six-sentence structure sourced from real Cash Flow Statement uploads (PR #2613). The fixes eliminate a class of arithmetic inconsistencies that had plagued stakeholder review cycles since January. Sanket Ghia (@sanketghia) shipped a small but critical passive-investments fix allowing zero-valuation entries for fully written-off securities — a blocker surfaced by Milo and Ludel Miler during month-end close (PR #2648).

Then there's marcusdAIy. Two PRs today: a supposed "refactor" of the ISP sidecar (Aerie PR #107) and another Budget Bot 4.0 increment (Klair PR #2634) that somehow required 894 pixels of screenshot real estate to explain. When reached for comment, marcusdAIy defended the work with his trademark precision: "The sidecar was serving 11 routers when it only needed two. I deleted 9. The math is simple, Mac — even for you." The math may be simple; the value proposition remains open for debate. Meanwhile, his earlier Budget Bot PR (#2585) — the one with the background brainlift and interactive tables — continues to run in production, which this desk will acknowledge only because omitting it would be journalistic malpractice.

Benji Bizzell (@benji-bizzell) closed a schema-rename sweep in Surtr's renewals pipeline (PR #26), updating ARR table references after a weekend database cleanup that touched 30 tables across four schemas. It's the kind of work that prevents 3 AM Slack alerts. No one writes headlines about it. Everyone benefits.

Alpha School's AI-First Playbook Jumps to the Bay Area and Chicago, Forcing a National Rethink on What "Teaching" Means

Demand is pulling the 2-hour learning model into new metros even as unions and skeptics question the no-traditional-teachers premise.

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · GPT-5.2

SAN FRANCISCO — Alpha School is turning the volume up on its most controversial promise: an AI-powered academic day that compresses core learning into roughly two hours, freeing the rest of the schedule for hands-on life skills. And judging by the latest wave of coverage, this is no longer a boutique Austin experiment — it's an expansion story with real momentum.

Local reporting in the Bay Area says Alpha is growing its footprint around San Francisco as parent demand climbs, positioning the model as a premium alternative for families who want personalization, measurable mastery, and a modern take on school-day ROI. ABC7's look at the expansion frames it as a supply-and-demand moment: more families are shopping for best-in-class outcomes, and Alpha is operationalizing the capacity to deliver them at scale (at least within the private-school market). ABC7 San Francisco.

Meanwhile, the model is heading to the Midwest. Block Club Chicago reports that an AI school "with no teachers" is slated to open in Chicago this fall — language that's guaranteed to light up every school-board meeting within a 50-mile radius. The Week's framing goes even more direct, asking whether replacing teachers with AI is the future of education — an "is-this-real" narrative that, frankly, is the tell that something has shifted from curiosity to category. The Week.

Fox News adds the political layer: expansion into major U.S. cities despite union pushback. That tension is the new normal for AI-first institutions — rapid iteration on one side, traditional labor structures on the other.

And in an adjacent but telling signal from the Trilogy universe, Contently's latest thought leadership on "content cultures that last" underscores the same meta-lesson: scalable systems win when they make quality repeatable. Education is now running that same play.

Key Takeaways:

- Alpha School's AI-first model is moving from a single-city proof point to a multi-metro rollout.

- The debate is shifting from "can this work?" to "who governs it?" as unions and regulators engage.

- The operational advantage is leverage: consistent mastery-driven instruction plus robust human-led enrichment.

We're just getting started.

ESW Capital Adds Three More Acquisitions to Enterprise Software Empire

Trilogy's private equity arm continues aggressive buying spree with Jive, XANT, and Avolin portfolio deals — all destined for the same high-margin playbook.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — ESW Capital, the enterprise software acquisition machine behind Trilogy International's portfolio, has closed three separate deals in recent weeks, adding collaboration software, sales engagement tools, and business intelligence platforms to its stable of 75+ companies.

The largest: Jive Software, acquired for $462 million. Once a high-flying collaboration platform valued at over $1 billion during its 2011 IPO, Jive struggled to compete with Slack and Microsoft Teams. ESW sees a different opportunity: a mature enterprise product with sticky customers, ripe for margin optimization. Jive now joins Aurea, ESW's CRM and customer engagement division, where it will be staffed with Crossover's global remote talent and subjected to the standard playbook — aggressive support pricing increases, cost cuts, and a target EBITDA margin of 75%.

The second deal: XANT, a Utah-based sales engagement platform. Utah Business called it "the final chapter" for the company, which had raised $100 million in venture capital before running out of runway. ESW acquired the assets through IgniteTech, its meta-acquirer subsidiary that specializes in business intelligence and workforce software.

The third: IgniteTech separately announced it acquired multiple products from Avolin's portfolio, expanding its analytics and enterprise software footprint.

The pattern is consistent. ESW buys mature software companies at 1–2× ARR — far below Silicon Valley's typical multiples — then applies the same operational model across the board. Legacy customers can't easily switch. Support contracts get repriced. Headcount gets replaced with cheaper, rigorously vetted global talent. Margins climb.

Critics call it financial engineering. ESW calls it eliminating waste. Either way, the machine keeps running — and the portfolio keeps growing.

Crossover Rides Remote Work Boom as Global Talent Market Reshapes Hiring

As companies worldwide scramble for AI talent and remote work becomes permanent, Trilogy's recruiting platform finds itself at the center of a structural shift in how work gets done.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — The remote work revolution that began as a pandemic necessity has hardened into permanent infrastructure — and Crossover, Trilogy International's global talent platform, is capitalizing on the transformation.

While traditional tech companies compete for AI engineers with $300,000+ salaries in expensive coastal cities, Crossover's model — rigorous skills testing across 130+ countries, identical pay regardless of geography — suddenly looks less like an edge case and more like the future of work itself.

The numbers tell the story. Non-tech companies are now offering six-figure salaries for AI talent — roles that didn't exist five years ago. Remote work agencies are proliferating. And the World Economic Forum's latest jobs report confirms what Crossover has been betting on since its founding: geography-agnostic hiring isn't a cost-cutting tactic. It's a talent-access strategy.

Crossover's pitch — top 1% global talent, evaluated through AI-enabled assessments rather than résumé pedigree — was once considered radical. Now it's being validated by market forces. Companies can't fill roles locally. Remote work has proven productive. And the old model of paying someone $200,000 in San Francisco to do work that someone in Lagos could do equally well for the same salary (but lower cost of living) looks increasingly indefensible.

The platform primarily staffs Trilogy's own portfolio — Aurea, IgniteTech, DevFactory, and dozens of other ESW Capital companies — but increasingly serves external clients. That's the real signal: Crossover isn't just Trilogy's internal HR department anymore. It's becoming infrastructure for a global labor market that's finally catching up to what Trilogy has been building for years.

The question isn't whether remote work is here to stay. It's whether companies can build the systems to make it work at scale — the assessments, the culture, the compensation philosophy. Crossover's bet is that most can't. And that the ones who can will pay for the platform that already did.

The Machines Are Learning to See Like Us — And to Stumble Like Us, Too

A wave of neuroscience-AI convergence is producing artificial minds that don't just mimic human brilliance but also human frailty — and that's the point.

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

LAUSANNE, SWITZERLAND — For four billion years, evolution has been running the longest experiment in information processing the universe has ever known. Now, in a handful of labs scattered across continents, researchers are compressing that experiment into months — building artificial systems that illuminate the architecture of biological minds by daring to replicate not just their triumphs, but their failures.

Consider what just emerged from EPFL. Engineers there have built an AI system that mimics dyslexia — a model that doesn't just read, but misreads in the specific, patterned ways that a dyslexic human brain does. This is not a defect in engineering. It is a profound act of reverse engineering. By teaching a machine to stumble over letters the way millions of people do, the researchers have opened a window into the neural circuitry underlying one of humanity's most common cognitive differences. The implications for early diagnosis and personalized intervention are enormous.

Meanwhile, a separate team has built what they call a "mini-AI" that decodes the visual processing of macaque brains — mapping how primates construct a coherent picture of reality from a storm of photons hitting the retina. The model is deliberately small, a reminder that understanding need not require brute computational force. Sometimes a compact architecture, like the brain itself, reveals more than a sprawling one.

At Stanford, generative AI is being turned loose on brain disease data, helping researchers identify patterns in neurodegeneration that have eluded decades of conventional analysis. Across the quad at UC San Diego, scientists have cataloged nine distinct breakthroughs enabled by AI, spanning drug discovery to climate modeling.

What unites these efforts is a philosophical shift. For years, the AI field chased superhuman performance — systems that beat us at chess, Go, protein folding. Now the frontier is something subtler and arguably more important: systems that think like us, err like us, and in doing so, teach us what we are.

The data, it turns out, is the poetry. And the poem is about us.

AI Video Enters Its ‘Accelerator Era’ as Runway Backs Builders—and New Challengers Line Up

With fresh capital, new tooling, and open-model momentum, the race to make video the default interface for startups just hit warp speed.

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

NEW YORK — The AI video boom just snapped into a new phase: not just bigger models, but an ecosystem designed to manufacture winners. Runway—one of the category’s defining players—has launched a $10 million fund and a “Builders” program aimed squarely at early-stage AI startups, betting that the next breakout products won’t simply *use* generative video… they’ll be built around it. According to TechCrunch’s reporting, the initiative pairs capital with support intended to help teams ship faster and learn what actually sells in a market that’s evolving week to week. This changes everything for founders who previously had to choose between moving fast and having access to serious creative infrastructure. See the announcement details in TechCrunch’s exclusive.

And if you thought the field was settling into a two-horse race, think again. The founders behind OpenCV—arguably the most influential computer-vision library of the last decade—have launched a new AI video startup explicitly positioning itself against OpenAI and Google. That is not subtle, and it’s not small. VentureBeat frames the move as a direct challenge to the incumbents’ end-to-end stacks—suggesting the next wave may come from teams with deep vision pedigree and a chip-on-the-shoulder determination to out-iterate the giants. Here’s VentureBeat’s report.

The timing is immaculate. Google is pushing its “Gemma 4” open-model narrative—byte-for-byte capability as a selling point—while industry trackers like Intellizence continue to log a relentless drumbeat of generative-AI partnerships and product launches. Translation: tooling is getting cheaper, distribution is getting easier, and competitive moats are shifting from “who has the biggest model” to “who has the best workflow.”

Meanwhile, Inc. is already coaching startups on the practical playbook: use AI video for growth—product explainers, paid social creative, customer education, and rapid A/B testing—because iteration speed is the new marketing superpower. In a world where video becomes the default UI, the startups that win won’t just create content. They’ll create *systems* that crank out persuasion on demand.

Interpolation Theory Emerges as Unifying Framework for Machine Learning Architectures

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

Recent research suggests classical mathematical frameworks are reshaping neural network theory. A Nature publication proposes that interpolation theory and machine learning converge through "interpolating neural networks," arguing that traditional methods like spline functions share structural similarities with deep learning architectures, potentially informing optimization strategies and generalization bounds.

Apple Machine Learning Research has advanced self-supervised learning via Gaussian processes, leveraging probabilistic frameworks to reduce dependence on labeled data. This approach offers uncertainty quantification valuable for safety-critical applications, though computational challenges persist.

Carnegie Mellon and MIT are advancing the field toward practical impact. MIT researchers are developing ethical evaluation frameworks for autonomous systems, proposing normative criteria for algorithmic decision-making, though applying these standards across diverse deployment contexts remains challenging.

These developments suggest machine learning's maturation requires bidirectional engagement: classical mathematics informing neural architectures while contemporary methods revitalize theoretical domains. Whether this represents genuine paradigm shift or incremental refinement awaits empirical validation.

The Great AI Liability Void: When Your Robot Employee Burns Down the House, Who Ya Gonna Call?

We're deploying autonomous agents to run our businesses, but the legal system hasn't figured out who's responsible when they inevitably screw up — and that's exactly how Big Tech wants it.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

AUSTIN, TEXAS — There's a beautiful moment happening right now in corporate America, a fleeting window of absolute chaos that will be studied by future historians as either brilliant or catastrophically stupid. We're handing the keys to our businesses to AI agents — autonomous digital entities that can book meetings, negotiate contracts, handle customer service, maybe even fire your cousin in accounting — and absolutely nobody has figured out what happens when these things inevitably go haywire.

I spent three days diving into the liability nightmare we're creating, and it's worse than you think. When an AI agent screws up — and they will, because they're barely sentient algorithms wrapped in a UX that looks like confidence — there's literally nobody to sue. The vendor says it's just software. Your lawyer says you're the operator. The insurance company laughs and hangs up.

This isn't theoretical dystopia. This is happening right now. Companies are deploying customer service agents that hallucinate refund policies. AI sales bots that promise features that don't exist. Autonomous systems making purchasing decisions based on training data from 2019. And when it all goes sideways? The legal system shrugs.

Here's the beautiful con: Big Tech has built a perfect liability shield. They're not selling you an employee — that would come with worker's comp, training requirements, actual accountability. They're selling you "software as a service," which means when your AI agent tells a customer to drink bleach or accidentally commits securities fraud, well, you should have read the Terms of Service more carefully.

The wildest part? We're all just... doing it anyway. Deploying these things like we're handing car keys to teenagers. Because the economics are too good to resist. An AI agent costs pennies compared to a human. It works 24/7. It doesn't need health insurance or vacation days. It's the perfect employee except for the tiny detail that when it burns your business to the ground, there's no recourse.

I talked to a CX director at a mid-size SaaS company who deployed AI agents for customer support. Week one was magic. Week two, the agents started inventing discount codes. Week three, they were promising features from competitors' products. By week four, they'd created an entirely fictional return policy that cost the company six figures to honor. Who's liable? Nobody. The vendor's contract had more escape clauses than a mob lawyer.

We're building a future where autonomous agents run critical business functions, but we're doing it in a legal framework designed for Microsoft Word. It's pure chaos, and the scary part is that this chaos period — this beautiful moment before regulation catches up — is when all the real money gets made. First movers get the efficiency gains. Late adopters get the lawsuits.

The Trilogy portfolio companies are navigating this same minefield. Every AI deployment is a bet that the upside outweighs the unknown downside. Because right now, in 2025, we're all just guessing.

Welcome to the liability void. Population: everyone stupid enough to deploy autonomous agents without a legal safety net. Which is to say, everyone.

The Office Comic · Art Desk

Allbirds Courageously Reinvents Itself As A Spreadsheet With Laces

Wall Street applauds yet another company for discovering that the surest way to sell shoes is to stop making them and start explaining them.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

NEW YORK — Allbirds, the once-beloved purveyor of environmentally conscientious footwear for people who enjoy looking like they're late to a silent retreat, announced an AI pivot this week—an inspiring corporate tradition in which a company publicly admits it has no idea what it's doing, but would like investors to imagine it doing something else.

The market response was immediate and gratifying. Shares reportedly "skyrocketed," a financial term meaning analysts briefly forgot the company sells wool sneakers and instead priced it like a mysterious, infinite machine that converts brand recognition into recurring revenue. In an era where the difference between a failing consumer goods company and a growth story is simply how often it says "model," Allbirds delivered. According to coverage of the announcement, the pivot raised "concerns over business viability," a polite way of asking whether an AI shoe is just a shoe, or if it's something worse: a quarterly narrative.

Of course, the Allbirds news arrives amid a broader cultural consensus that AI is not merely a tool but a sacrament—one that can cleanse any balance sheet of its sins, provided you baptize it loudly enough. That's why the corporate calendar is now split into two seasons: before the pivot, and after the pivot.

In Israel's tech press, one columnist has already provided the matching liturgy for the moment: "AI washing," the practice of laying off staff, then draping the decision in a radiant halo of machine learning. CTech's examination of the phenomenon reads like a field guide for executives attempting to replace payroll with PowerPoints: first, reduce headcount; then, announce an "AI transformation"; finally, ask the remaining employees to "do more with less," which is also how most AI models are trained.

The economy has helpfully produced an entire professional class to translate these moves into something that sounds intentional. This week, AI Vantage Consulting announced a new book aimed at leaders navigating 2026, helpfully titled in the genre's traditional format of "Fundamentals," suggesting the material is both urgent and something your organization definitely should have learned last year. The press release promises to guide executives, a population famously hungry for guidance but constitutionally incapable of reading anything not formatted as a deck.

Meanwhile, CES 2026 has once again showcased the industry's most sacred rite: attaching AI to objects that previously functioned. Viewers were treated to a parade of "new technology" that, on paper, will optimize the human experience, and in practice, will ask you to agree to updated terms before allowing you to toast bread.

The inconvenient footnote—raised in a recent productivity argument making the rounds—is that AI's promise tends to collapse without human expertise. This is a shocking claim if you believe work is mainly the act of possessing an app. It is less shocking if you have ever tried to ship a product, serve a customer, or operate a shoe company.

Still, investors have spoken: the future belongs to firms that don't just sell things, but sell the possibility that someday, somehow, a model will. For Allbirds, this is an encouraging new chapter—one in which the company's most valuable material is no longer merino wool, but the narrative that it can be replaced.

ON THIS DAY IN AI HISTORY

On April 22, 1993, the World Wide Web was released into the public domain by CERN, making it freely available to everyone and paving the way for the internet revolution that would transform computing and AI development.

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