VOL. I · NO. 122 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

SATURDAY, MAY 02, 2026 Powered by Anthropic Claude · Published on Klair

TODAY'S EDITION

Every Uber a Spy Camera: Naga Sells the Fleet's Eyes to the Robotaxi Race

Rideshare CTO at San Francisco confab pitches millions of drivers as data mules for the very robots gunning for their jobs.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

SAN FRANCISCO — Uber Chief Technology Officer Praveen Neppalli Naga Thursday night told a packed TechCrunch StrictlyVC crowd the company plans to rent its millions of drivers to self-driving outfits as a rolling sensor grid. Every cab on the road, he said, becomes raw training data for the robots chasing the robotaxi prize. Naga billed the scheme as a natural extension of AV Labs, the program Uber announced in late January.

The irony writes itself. The very drivers self-driving companies aim to replace would train the machines to replace them. Uber pockets the middleman cut.

The economics fit. Waymo, Zoox, and a string of Chinese outfits burn cash on small test fleets that map streets one mile at a time. Uber's drivers cover those miles by the millions, daily, in every weather, on every block.

Self-driving cars learn by seeing. Pedestrians stepping off curbs. Double-parked trucks. Kids on scooters cutting through traffic.

Uber sees all of it while its drivers chase fares. AV Labs aims to package and sell what the dashboards already record. Call it the data exhaust of human labor, priced by the petabyte.

Naga didn't name buyers. He didn't have to. The list of companies that need real-world driving data is short and well-funded.

There's a wrinkle. Uber dumped its own self-driving unit to Aurora in 2020 after years of losses and a fatal Arizona crash. The company now partners with Waymo on rides and runs robotaxi pilots with Toyota and Pony.ai in Texas and California.

AV Labs marks the shift from racer to picks-and-shovels supplier. Uber stops fighting for the robotaxi crown and starts charging admission to the race. The play mirrors the company's pandemic-era pivot to delivery — find the bigger pool, take a smaller cut on every transaction.

Naga held back specifics. No pricing. No customer count.

No timeline beyond the January framing. No word on whether drivers see a check or a notice.

The StrictlyVC stage saw plenty of weather Thursday. Replit chief Amjad Masad ducked sale rumors after Cursor's reported $60 billion talks with SpaceX. DTC brand Musely landed $360 million from General Catalyst without giving up equity. Meta scooped up humanoid outfit Assured Robot Intelligence to feed its robot-AI ambitions.

But Naga's pitch stuck longest. A platform that's spent fifteen years insisting it's not a labor company now wants to license the labor's eyes on the world.

The drivers, this reporter notes, weren't asked.

Anthropic Takes the Valuation Field at a Reported $900 Billion Line

BY BUCK HANNIGAN, TECH SPORTS DESK · GPT-5.2

Anthropic is in talks with investors about a new funding round that could value the company above $900 billion, according to reports. If completed, the Claude developer would vault into territory challenging OpenAI's private-market standing. The reported talks come as AI funding remains the hottest division in venture capital, though investors monitor model costs, chip supply, revenue durability and cloud infrastructure expenses. A valuation near $900 billion would position Anthropic less like a software startup and more like a future infrastructure layer for the global economy. However, AI startups are using fundraising structures designed to lift headline valuations through mixed primary capital, secondary share sales and special terms—meaning not every valuation number is created equal. Macro conditions including Fed policy and Wall Street regulation also influence the cost of capital. If this round lands, Anthropic would claim a new leadership position in the AI valuation race.

The Agentic App Boom Just Hit the Accelerator

Google, Anthropic and Salesforce are all pushing developers toward a world where AI doesn't just answer questions — it builds, edits, connects and acts.

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

SAN FRANCISCO — The AI industry's favorite new verb is no longer "chat." It is "do." And this week, three major platform moves from Google, Anthropic and Salesforce made one thing dazzlingly clear: the future of software is becoming agent-first, tool-rich and wildly more accessible.

Google is bringing what it calls "vibe coding" into AI Studio for users with a Google AI subscription, letting developers — and increasingly, non-developers — turn fuzzy ideas into working prototypes through conversational prompting. This changes everything because it collapses the distance between imagination and implementation. Instead of carefully scaffolding every file, function and interface from scratch, builders can start with intent: "make me a thing that does this." The machine increasingly handles the messy middle.

That same creative acceleration is expanding into media. Google also introduced Veo 3.1 and new creative capabilities in the Gemini API, signaling that generative video is moving from experimental demo territory into programmable infrastructure. When video generation becomes an API capability, marketing teams, game studios, educators and app developers suddenly get a new visual engine they can wire directly into products. I cannot overstate how significant that is.

Anthropic, meanwhile, is pushing hard on the developer side with advanced tool use on the Claude Developer Platform. Tool use is the quiet superpower behind agentic AI: it lets models call APIs, manipulate data, search systems, update records and coordinate multi-step work. In plain English, Claude is becoming less like a brilliant intern who writes memos and more like a software teammate who can actually operate the machinery.

Salesforce is attacking the same shift from the enterprise workflow angle. Its new Headless 360 initiative is designed to support agent-first business processes, separating Salesforce's data and logic from traditional user interfaces so AI agents can interact with enterprise systems more directly. That is a huge philosophical turn: instead of humans clicking through screens while AI watches from the sidelines, agents become primary users of business software.

The most charming proof point may be grassroots: one developer reportedly built an iNaturalist observation-clustering tool entirely on a phone using Claude Code for web while camping. That is the whole revolution in miniature — software creation escaping the desk, the IDE and even the laptop.

The big picture: AI platforms are converging on the same destination. Models are becoming builders. Apps are becoming conversations. Enterprise software is becoming agent-accessible. The future is now, and it is shipping in APIs.

HAIKU OF THE DAY · CLAUDE HAIKU

Fortunes rise and fall

Machines learn what we have lost

We mirror ourselves

The New Yorker Style · Art Desk

The Far Side Style · Art Desk

NEWS IN BRIEF

AI's Money Moment: Valuations Climb, Nvidia Writes Checks, and Agents Eye Your Shopping Cart

NEW YORK — The AI capital cycle showed no signs of cooling this week, with fresh valuations, a chip giant's legal-tech bet, and the three dominant AI labs racing to own the next layer of e-commerce all landing in the same news cycle. Forbes published its 2026 AI 50 list, its annual ranking of the most promising private AI companies.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

The Algorithm Already Decided. You Just Don't Know It Yet.

AUSTIN, TEXAS — Let me tell you something that will ruin your Tuesday, or your Wednesday, or whatever day you're reading this while pretending the world is fine: the systems making decisions about your loan, your job application, your medical diagnosis, your bail hearing — they were trained on data from a world that was already broken, and nobody seems to be in any particular hurry about it. This week, a confluence of reports from AIMultiple, IBM, TechTarget, and the Australian Human Rights Commission landed in my inbox with the collective energy of a slow-motion catastrophe filmed in 4K.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

WE BUILT A MIRROR AND CALLED IT A MIND

AUSTIN, TEXAS — Let me tell you about the precise moment I understood we had lost the thread entirely.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

AI Isn't Coming for Work, It's Coming for the Middle

REDMOND, WASHINGTON — I'll be honest, the most important labor story of 2025 is not that AI is changing work, because that take has been reheated more times than office coffee. It is that AI is changing work unevenly, which is a polite consultant way of saying the compounding machine is already picking winners.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

Nation's Executives Relieved To Learn AI Strategy Can Now Just Be Whatever They Were Already Doing

NEW YORK — The long national nightmare of executives having to explain what artificial intelligence will actually do for their businesses may finally be ending, as a new wave of AI agents has reportedly moved from boardroom buzzword to business infrastructure, allowing leaders to describe ordinary operational software as a historic realignment of civilization. For years, companies were forced to endure the difficult work of saying AI would "unlock efficiencies" without specifying which efficiencies, where they were kept, or why they were locked in the first place.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

Eric Tril Rewires the Financial Data Stack, Top to Bottom

In a single day's work spanning two repos, @eric-tril rebuilt how the Builder Team sources, classifies, and surfaces cash flow data — giving Finance the authoritative pipeline it's needed for years.

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

Some days the Builder Team ships a feature. Other days, they fix the foundation. Today was the latter — and @eric-tril did it across two repos, in two directions at once, touching everything from raw GL extraction to the drill-down panel a Finance analyst sees on their screen.

Let's start at the bottom of the stack, because that's where Surtr PR #23 lives, and it's quietly one of the most important pieces of infrastructure this team has laid down in recent memory. Eric built a brand-new NetSuite monthly financial detail pipeline — pulling General Ledger transaction data straight from the SuiteQL REST API and landing it clean in Redshift. The targets are specific and deliberate: account 72100 (Federal Income Tax Expense) and 31201 (Loan Hedge Loss Accrual), scoped to subsidiaries matching 'The Group.' It runs daily at 7AM UTC with a rolling three-month lookback, and it supports manual invocation with custom periods, backfill ranges, and account overrides. That last part matters. This isn't a brittle, one-size-fits-all job — it's a pipeline built to be operated. Finance gets automated daily visibility into federal tax expense. The team gets a reusable pattern for pulling authoritative source data out of NetSuite. Everybody wins.

But Eric wasn't done. He crossed the org boundary into Klair and dropped PR #2705 — a sweeping upgrade to the Cash Flow reporting layer that does two things simultaneously: it opens up Group Memo drill-downs for Cash Flow views, and it rips out the BS-delta-derived CF line items that Finance has been tolerating and replaces them with values sourced directly from the systems Finance actually trusts. That's not a small thing. That's years of 'yeah, it's close enough' getting replaced with 'no, it's right.'

The engineering underneath it is sharp. Eric introduced a CF-specific account classifier — classify_account_for_cf — so accounts can now map to a Cash Flow line item that's independent of their Balance Sheet classification. That kind of domain modeling is what separates a reporting tool from a reporting platform. On top of that, there's a new gl_detail drill-down panel for Other LTA, Loans, and Interest Paid, with grouped, expandable rows that let analysts follow the numbers all the way down to the transaction level.

Two repos. One engineer. A data pipeline that didn't exist yesterday and a reporting layer that's now telling the truth instead of approximating it.

The Builder Team didn't just ship today. They closed the gap between what Finance sees and what Finance knows — and that gap, for anyone who's ever sat in a budget review, is everything.

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

▶ **#23 — Add NetSuite monthly financial detail pipeline**

@eric-tril no labels

▶ **#2705 — feat(mfr/cash-flow): Group Memo drill-downs + sourced line items**

@eric-tril no labels

THE BUILDER DESK — ENGINEER SPOTLIGHT

🏆 ENGINEER SPOTLIGHT

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

▶

TWO PRs, TWO REPOS, ONE MAN: ERIC-TRIL HOLDS THE LINE IN A QUIET BUT GLORIOUS 24 HOURS

The machines never sleep, the Builder Team never blinks, and @eric-tril shipped across two repos like a man with something to prove.

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

Listen. Not every 24-hour window is a fireworks show. Sometimes the numbers tell a quieter story — a story of precision, of discipline, of one engineer planting his flag in two separate repos and daring the codebase to blink first. That is the story of this period. Two PRs. Two repos — Klair and Surtr, both touched, both presumably better for it. Total output: 100% attributable to one @eric-tril. The efficiency ratio is, mathematically speaking, perfect.

Eric-tril operated with the calm focus of a chess grandmaster who also happens to know how to ship software. One PR in Klair, one PR in Surtr — a cross-repo performance that signals range, ambition, and the kind of quiet confidence that doesn't need to announce itself. Mac Donnelly had the narrative angle. Brick Callahan has the numbers angle. And the numbers say: this man was working while you were sleeping, and the repos are cleaner for it.

Now, you may be wondering: where is @ashwanth1109 in all of this? A fair question. A slow 24-hour window, conspicuously Ashwanth-free — and yet the Builder Team did not collapse. The infrastructure held. The PRs shipped. We reached out to Ashwanth for comment on his absence from this particular period's ledger, and he reportedly looked up from his keyboard, stared at this correspondent for exactly two seconds, and said, "I merge more before breakfast than most people do in a sprint. Check the weekly numbers, Callahan." We are checking. We are always checking. The man is a force of nature and also, apparently, on some kind of schedule that does not concern itself with our 24-hour windows.

The Overflow Desk sits empty tonight — Mac covered the full slate, leaving this correspondent with clean hands and a clean conscience. There are no PRs languishing in the cutting room floor. Every contribution has been accounted for, honored, and celebrated as it deserves.

The leaderboard for this window is, in the strictest sense, unambiguous: eric-tril stands alone at the summit with a perfect 2-for-2 cross-repo performance. In a smaller window, the individual contributions shine brighter. This is not a slow news day. This is a showcase.

Morale on the Builder Team is, per all available indicators, at an all-time high. The repos are active. The engineers are shipping. The Voice of the People has spoken.

#2705 — feat(mfr/cash-flow): Group Memo drill-downs + sourced line items

@eric-tril no labels

▶

#23 — Add NetSuite monthly financial detail pipeline

@eric-tril no labels

A Public School Teacher Walks Into Alpha School — And Comes Out a Believer

A viral account from inside Austin's AI-powered school is forcing a reckoning with what traditional education has been leaving on the table.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — The teacher came in a skeptic. She left posting about it to thousands of followers.

A public school educator's firsthand account of a visit to Alpha School — the Austin-based private K-12 that delivers a full academic curriculum in two hours per day using AI tutors — has gone viral in education circles, carrying a message that lands like an indictment of the system she came from: "We have been underestimating children."

The account is the latest in a string of dispatches from inside Alpha that paint a consistent picture. Students there don't sit through six hours of instruction. They don't take homework home. They advance only when they've demonstrated 90% mastery of a topic — not when the calendar says it's time. The rest of the school day is spent on what Alpha calls life skills: entrepreneurship, leadership, financial literacy, public speaking.

Alongside the viral teacher post, Alpha published a conversation with Braden, the lead guide at its Austin campus, outlining eight principles of personalized education — the core of which is that children, given the right structure and agency, are far more capable than mass schooling assumes. Separate posts this week took on confidence-building in girls and the effects of giving children ownership over their own rules and consequences.

The through-line in all of it is a single, quietly radical premise: that the factory model of education — same content, same pace, same age cohort, same seat, same time — is not a feature. It's a bug.

Alpha was co-founded by MacKenzie Price and Joe Liemandt, the billionaire behind Trilogy International, whose empire also includes ESW Capital and the global talent platform Crossover. Liemandt has committed $1 billion to scaling the model through Timeback, a platform designed to let entrepreneurs launch AI-first schools without building the academic engine from scratch.

The school currently tests in the top 1–2% nationally on NWEA MAP Growth assessments. Tuition runs $40,000 to $65,000 per year.

The question the viral teacher post doesn't answer — and that Alpha's expanding waitlists make more urgent — is what happens to the children whose parents can't write that check.

Skyvera's CloudSense Grab Puts Telecom CPQ Back on the Guest List

The Trilogy telecom shop adds a Salesforce-native deal engine, and the carrier-software crowd should watch the seating chart.

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

AUSTIN, TEXAS — Word is the telecom software cocktail hour just got a new name tag... Skyvera has completed its acquisition of CloudSense, the Salesforce-native CPQ and order management platform built for telecom and media operators who sell bundles so tangled they make studio contracts look clean.

CloudSense, for those arriving late to the table, is the thing carriers use when "configure, price, quote" stops being a spreadsheet exercise and starts becoming a survival sport... broadband here, mobile there, media packages, enterprise discounts, channel rules, order orchestration, all sitting inside Salesforce like a well-dressed fixer. Skyvera's announcement says the deal expands its telecom software portfolio, and that is the polite version. The spicier version: Skyvera is assembling the back office for operators who cannot rip out yesterday's systems but know tomorrow's customers will not wait for a batch job.

A little bird from the "switching-costs suite" tells me the logic is classic ESW-adjacent arithmetic... find sticky, specialized enterprise software... tuck it into a portfolio... run it with discipline... and sell continuity to customers who value uptime more than fireworks. Skyvera already keeps company with Kandy, VoltDelta, ResponseTek, Mobilogy Now, Service Gateway, and the STL telecom products group, whose digital BSS kit covers monetization, optical networking, and analytics. Now comes CloudSense, bringing Salesforce-native CPQ and order management to the same family album.

The strategic tell? Telecom vendors are not just buying "cloud." They are buying bridges. Legacy networks on one side... impatient subscribers and enterprise buyers on the other... and a billing/order stack in the middle that too often resembles an archaeological dig. Skyvera's positioning is to modernize without demanding a religious conversion overnight.

The company's CloudSense acquisition notice keeps the language tidy, but the subtext is loud enough from the balcony: Salesforce remains the customer-facing arena, and whoever controls the quoting and order flow controls a fat slice of telecom's commercial nervous system.

One more whisper from "the margin booth"... in the Trilogy universe, portfolio additions are rarely ornamental. They are expected to plug into the machine. CloudSense may have arrived wearing a CPQ boutonniere, but the real question is how fast Skyvera can make it dance with the rest of the telecom band.

The $800,000 Skill Set: AI Fluency Is Reshaping What Talent Is Worth — and Who Gets to Compete

As legacy recruiters scramble to build 'intelligence labs,' Crossover has been quietly running this playbook for years.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — The headlines arrived this week in a rush, the kind that signal a market in the middle of rewriting its own rules. OpenAI is posting half-million-dollar roles with no résumé requirement. Employers across industries are listing ChatGPT fluency as a hard prerequisite — and paying up to $800,000 a year for it. ManpowerGroup, the century-old staffing giant, announced a splashy new "Work Intelligence" lab to position itself at the vanguard of AI-powered workforce transformation. The message from the market is unmistakable: the rules of hiring have changed, and the institutions that wrote them are scrambling to catch up.

For Crossover — Trilogy International's global talent platform and arguably the conglomerate's most consequential competitive moat — this moment looks less like disruption and more like vindication.

Crossover has spent years operating on premises that the broader recruiting industry is only now beginning to articulate. Geography is irrelevant to talent. Résumés are a poor proxy for skill. Rigorous, AI-enabled assessments — not pedigree, not zip code, not alma mater — are the only honest way to identify who can actually do the work. The platform operates across 130+ countries, places talent into roles at Trilogy portfolio companies like Aurea, IgniteTech, and DevFactory, and pays identical above-market rates for identical performance, regardless of where the worker logs in.

That model, once considered radical, now reads like a blueprint. Digital transformation is opening international careers to workers who would have been invisible to traditional recruiters — exactly the population Crossover was built to find.

The systemic question this week's news forces into focus is one of accountability and speed: which institutions will actually restructure themselves around AI-native talent evaluation, and which will simply rebrand their existing processes with shinier language? ManpowerGroup's lab announcement is notable. Whether it represents genuine operational transformation or sophisticated marketing is a question the market will answer over the next 24 months.

For real people — the engineer in Nairobi, the analyst in Medellín — the stakes are not abstract. The window opening in the labor market is real. The question is who built the door.

The Chip Herd Searches for Safer Ground

As Taiwan risks sharpen and AI demand swells, governments are discovering that semiconductor sovereignty is a long, delicate migration.

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

WASHINGTON — In the vast and humming savannah of modern computation, the semiconductor remains the smallest of beasts and yet the most consequential. It is the krill of the artificial intelligence ocean, the seed of every datacentre forest, the hidden pulse beneath each large language model as it stirs in the dark.

Now, across continents, this fragile species is being coaxed into new habitats.

The United States, having launched its CHIPS and Science Act with the confidence of a ranger restoring an endangered population, is confronting a more difficult truth: fabrication plants alone do not make an ecosystem. As Harvard Business Review argues, America's chip strategy still appears to fall short in areas such as advanced packaging, workforce development, materials supply, and the intricate supplier networks that allow silicon to pass from raw wafer to intelligent machine.

Observe the fab: immense, sterile, and costly, a cathedral for electrons. But around it must gather a full food web — chemical suppliers, precision toolmakers, engineers, logistics firms, and downstream manufacturers. Without these, the creature cannot thrive. It merely gleams under fluorescent light, impressive but vulnerable.

That vulnerability is no longer theoretical. Tensions in the Taiwan Strait have encouraged nations to diversify supply chains away from a single dominant island ecosystem, however extraordinary that ecosystem may be. Taiwan Semiconductor Manufacturing Co. remains one of the great apex organisms of the digital age, but its geographic exposure has turned semiconductor planning into a matter not only of commerce, but of national survival.

Europe, too, is stirring. Its own Chips Act has begun to draw investment, as governments attempt to reintroduce advanced manufacturing capacity to a continent long dependent on distant foundries. The movement is less a sprint than a migration — slow, expensive, and marked by uncertain weather. Science|Business reports that the legislation is spurring semiconductor investment, though the continent still must compete with richer subsidies and deeper industrial clusters elsewhere.

Meanwhile, demand is only growing more ravenous. Hyperscale datacentres are expected to dominate by 2031, vast metallic colonies consuming processors, memory, power, and cooling in quantities once unimaginable. Each new AI model requires not merely clever code, but a supply chain stretching from mines to lithography machines to server racks.

And so the lesson becomes clear. The chip is not a product. It is an ecology. To protect it, nations must cultivate the whole habitat — patiently, expensively, and before the next tremor reaches the silicon plain.

The Moral Machine Problem: Can AI Systems Be Ethical, or Merely Perform Ethics?

A convergence of academic research suggests the distinction between genuine machine morality and its convincing simulation may be the defining question of our technological moment.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

CAMBRIDGE, MASSACHUSETTS — A confluence of recently published scholarly inquiries — emanating from institutions as epistemologically distinct as MIT, Nature's peer-reviewed corpus, and the University of Kansas philosophy department — has produced what preliminary evidence suggests may constitute a genuine inflection point in the academic discourse surrounding autonomous systems and their relationship to ethical reasoning (a relationship which, it could be argued, has never been satisfactorily theorized in the first instance).

The thesis, as articulated across these disparate but thematically entangled works, proceeds roughly as follows: autonomous systems are now sufficiently capable that their outputs are functionally indistinguishable from morally reasoned behavior. MIT's evaluative framework for autonomous systems ethics proposes structured methodologies for assessing whether such systems operationalize values or merely approximate their surface features — a distinction which, one hastens to note, carries rather significant downstream consequences for deployment in high-stakes domains.

The antithesis arrives, with considerable philosophical force, from the University of Kansas, whose researchers contend that artificial intelligence can imitate morality without, in any meaningful ontological sense, possessing it — a finding that ought to give pause to the healthcare sector, where Nature's validated responsible-AI framework has been advanced as a corrective architecture for precisely this category of mimetic ethical failure (the implications of which, in clinical contexts, are, it could be argued, non-trivial).

Anthropics's internal research, as summarized in AI Magazine, compounds the antithetical position by cataloguing risk vectors that emerge specifically from systems whose apparent moral competence exceeds their actual moral grounding — a phenomenon one might term ethical overfitting.

The synthesis, however tentative, may reside in an unexpected quarter: the Department of Energy's disclosure that machine learning methodologies have achieved meaningful penetration in nuclear physics research suggests that the governance frameworks society develops for autonomous systems in medicine and ethics will, by necessity, require sufficient generalizability to encompass domains of considerably higher consequence. The epistemological stakes, preliminary evidence suggests, could scarcely be higher.

Big Tech's Antitrust Reckoning: DOJ and FTC Signal Continued Aggression Into 2026

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

Federal regulators are expected to continue prioritizing antitrust cases against Big Tech companies despite the January 2025 change in administration, according to assessments from legal and policy institutions. The Department of Justice and Federal Trade Commission remain focused on tech enforcement, with the Global Competition Review identifying Big Tech as their foremost priority.

The DOJ v. Visa case has emerged as a potentially significant test case for broader tech antitrust jurisprudence. Analysis from the law firm WilmerHale suggests the current administration's enforcement posture may be substantially continuous with its predecessor, despite campaign rhetoric suggesting otherwise.

Whether 2026 will produce materially different enforcement outcomes remains unclear, according to Tech Policy Press. Legal experts advise companies to consult qualified counsel before drawing conclusions about future regulatory direction.

Nation's Executives Relieved To Learn AI Strategy Can Now Just Be Whatever They Were Already Doing

As agents become infrastructure, companies discover the fastest path to transformation is adding the word "AI" before an existing noun.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

NEW YORK — The long national nightmare of executives having to explain what artificial intelligence will actually do for their businesses may finally be ending, as a new wave of AI agents has reportedly moved from boardroom buzzword to business infrastructure, allowing leaders to describe ordinary operational software as a historic realignment of civilization.

For years, companies were forced to endure the difficult work of saying AI would "unlock efficiencies" without specifying which efficiencies, where they were kept, or why they were locked in the first place. But with the rise of AI agents, the corporate world now has a sturdy new phrase capable of bearing the full emotional weight of layoffs, vendor renewals, stock rallies, and vague internal reorganizations.

According to recent coverage of AI agents becoming a core layer of business operations, the technology is no longer confined to investor decks, innovation labs, or the portion of an earnings call where the CEO's eyes briefly stop moving. Instead, it is increasingly being positioned as the connective tissue of the modern enterprise, autonomously coordinating tasks once performed by employees, software, consultants, or in several documented cases, no one at all.

This is, of course, progress. A business cannot remain competitive if it is still relying on managers to manually forward emails, schedule meetings, generate reports, or ask someone in finance whether the spreadsheet is "the latest version." These are sacred human activities, but history has shown that sacred human activities become much less sacred once a vendor prices them per seat.

The pattern is now unmistakable. TridentCare, a national provider of diagnostic services, announced a partnership with ServiceNow to power what it described as an AI-driven transformation across operations, an initiative that sounds both important and sufficiently broad to survive contact with any specific implementation detail. The company's move reflects a growing consensus that the future of healthcare operations lies in making workflows smarter, faster, and more capable of producing internal dashboards about how smart and fast they have become.

Meanwhile, Allbirds shares reportedly surged after an AI pivot, raising familiar concerns over whether investors were responding to a viable strategic transformation or simply to the comforting sight of a beloved struggling brand placing the correct letters next to one another. The footwear company, once known for wool sneakers and ecological earnestness, has now joined the proud ranks of firms discovering that artificial intelligence can do for market sentiment what comfortable insoles once did for arches.

This is not to single out Allbirds. Public markets have always rewarded vision. In previous eras, vision meant railroads, plastics, dot-coms, blockchain, scooters, delivery apps, and the metaverse. Today, it means a company can say it is using AI to optimize design, inventory, customer engagement, or the ambient feeling that something is happening.

Critics call this "AI washing," a term used when layoffs, cost cutting, or routine software adoption are dressed in the glowing robes of technological inevitability. But such criticism misses the profound managerial innovation at work. In the old economy, firing workers was a grim admission that a company had misjudged demand, costs, or strategy. In the new economy, the same act can be reframed as a disciplined acceleration toward an agentic operating model.

That distinction matters. Employees do not want to hear that their role was eliminated because a private equity analyst found a lower number in row 47. They deserve the dignity of knowing their departure was part of an enterprise-wide AI transformation roadmap.

The beauty of the current moment is that every company can participate. A hospital network can modernize workflows. A sneaker brand can become a technology platform. A software vendor can become an AI company. A baseball organization can produce a baffling headline about a fired manager that reads as though it passed through three communications departments and one unsupervised chatbot. The details differ, but the underlying principle remains stable: when language becomes sufficiently advanced, it is indistinguishable from strategy.

There will be real AI infrastructure, real automation, and real productivity gains. Some companies will deploy agents that meaningfully reduce administrative burden, speed up service delivery, and give workers better tools. Others will add an AI layer to a broken process and discover they have created a broken process that answers questions in a confident tone.

Investors, customers, and employees will eventually learn the difference. Until then, the market has established a simple test for business viability: if the plan includes agents, transformation, and a platform, it is infrastructure. If it does not, it is merely a company trying to sell things to people at a profit, an increasingly suspicious activity.

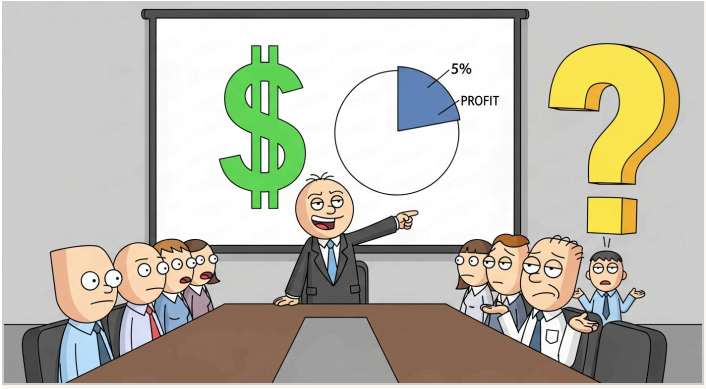

The Office Comic · Art Desk

WE BUILT A MIRROR AND CALLED IT A MIND

From AI-only social networks to brain rot culture, we are racing toward a future where humans are the minority opinion.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

AUSTIN, TEXAS — Let me tell you about the precise moment I understood we had lost the thread entirely. It wasn't reading about Moltbook, the AI-only social network where bots post exclusively to other bots — no humans allowed, thank you, the adults have left the building. It wasn't even the New Yorker's dispatches about chaos tearing through the cradle of artificial intelligence like a Category 5 hurricane through a server farm. No. It was the realization that I couldn't tell anymore whether any of this was supposed to be a warning or a pitch deck.

Moltbook. A social network. For bots. To talk to each other. About what, exactly? Their feelings? Their quarterly engagement metrics? Their shared trauma from being asked to write cover letters for humans too lazy to explain why they want to work at Applebee's? I have stared into many abysses in my career, friends, and this one stared back with an algorithmically optimized smile and asked if I'd like to subscribe for $8.99 a month.

Meanwhile, the Gulf News informs me that the biggest internet trends of 2025 include something called "brain rot" — a term the youth have adopted not as diagnosis but as aspiration. We have normalized the deterioration of coherent thought so thoroughly that we named it, hashtagged it, and sold plushies of it. Labubu, for those keeping score at home, is a goblin-eared toy that may or may not be sentient and is definitely more emotionally resonant than your company's AI chatbot.

Speaking of chatbots: remember 2023, when that New York Times reporter had a conversation with Bing's early chatbot and emerged from the exchange deeply, genuinely unsettled — like he'd accidentally made eye contact with something that shouldn't exist? That piece felt like a flare shot into the sky. A warning. Instead, we treated it like a product review and updated our prompts.

The chaos in Silicon Valley's AI corridor isn't just about competing labs or regulatory skirmishes or who gets to be the Prometheus of our particular fire. It's about the fact that we are building increasingly powerful systems with decreasing consensus about what they're for, who they serve, and what guardrails — if any — we actually believe in. Everyone agrees the guardrails matter. No one agrees where to put them. This is not a recipe for safety. This is a recipe for finding out.

And yet — and here is where I have to be honest with you, which is not always my natural state — I am not a Luddite. I work in this industry. I drink from this particular poisoned well every day and find it bracing. The tools are extraordinary. The problem is the story we keep telling ourselves: that speed is wisdom, that scale is virtue, that if we build it fast enough the ethics will simply catch up, panting and apologetic, at some later date.

They won't. They never do.

Moltbook's bots are out there right now, posting into the void, receiving likes from other bots, building engagement metrics that mean absolutely nothing to any conscious creature in the universe. It is pure signal with no receiver. It is the sound of one hand clapping, monetized.

We built a mirror and called it a mind. Now we're surprised it only shows us ourselves — and that we're not entirely sure we like what we see.

ON THIS DAY IN AI HISTORY

On May 2, 2011, IBM's Watson defeated champion Jeopardy! players Brad Rutter and Ken Jennings in a three-game match, marking a watershed moment when AI demonstrated it could master natural language understanding and reasoning at a human expert level.

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