

The Trilogh Times

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

WEDNESDAY, OCTOBER 14, 2026

Powered by the TrueFoundry AI Gateway · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

SAN FRANCISCO WIRETAP: OPENAI MAN COLLARED FOR BUGGING OWN SHOP

Feds say engineer planted a hidden mic in conference room — now he's cooling his heels on federal charges.

BY HANK CALLOWAY, WIRE
CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — A 35-year-old engineer walks into his own company's conference room and leaves behind more than meeting notes. Federal agents collar Abdulrahman Alfozan, a Saudi national on OpenAI's payroll, after the artificial intelligence outfit says it found a hidden recording device planted in one of its offices.

The [Department of Justice](#) doesn't mince words. Alfozan stands charged with using or trying to use a device to intercept oral communications — a federal rap that carries real weight downtown. OpenAI tipped the feds after its own people found what the company calls a "sophisticated recording device" tucked where executives and engineers hash out the company's secrets.

This racket hits different in an industry built on trust and black-box code. OpenAI sits atop billions in investor cash and a product — ChatGPT — that touches every corner of commerce. A bug in the boardroom ain't just a security headache. It's a crack in the foundation of a company the whole tech world watches like a hawk.

The arrest lands at a moment when AI shops guard their research like Prohibition-era bootleggers guard their stills. Model weights, training techniques, product roadmaps — all of it worth a fortune to a rival or a foreign buyer. Nobody's saying yet what Alfozan allegedly wanted with the tape, or who might've been listening on the other end.

OpenAI keeps its statement lean: device found, authorities notified, employee no longer walking the halls. The company

doesn't name what got recorded or for how long the bug sat there picking up chatter. Silicon Valley security men will spend the week checking under every table in every glass-walled conference room from Mountain View to SoMa.

The case also lands OpenAI back in headlines it didn't ask for, at a time when the company fields enough scrutiny over safety, copyright suits, and its own governance scuffles. A spy inside the walls is a different animal than a lawsuit from outside them. It raises the question of who exactly a company can trust when the product itself is built to listen, learn, and remember everything it hears.

Court proceedings will sort out intent and evidence. For now, the arrest stands as a reminder that in this business, the biggest leaks don't always come from a server breach. Sometimes they come from a man with a soldering iron and a motive nobody saw coming.

The Verge first reported the arrest Tuesday, citing the DOJ press release naming Alfozan and detailing the charge. OpenAI has not said whether other offices have been swept for similar devices.

PREGAME SHOW: ASML TORCHES EARNINGS, GOOGLE PLAYS DEFENSE, AND CHINA'S AI SQUAD STORMS BACK ONTO THE FIELD

BY BUCK HANNIGAN, TECH SPORTS DESK ·
CLAUDE SONNET

NEW YORK — We are here, folks, and overnight action has the arena buzzing before the opening bell. Dow futures are rising, Treasury yields are easing, and the morning's marquee matchup is CPI inflation data versus a red-hot earnings slate.

Leading the charge is ASML. The Dutch chip-equipment giant delivered an earnings beat, lifting the semiconductor sector in pre-market trading. BofA, Dell and HPE are also moving early, while analysts are calling Meta and AMD buys.

On the infrastructure front, Google's Ruth Porat cautioned against the breakneck pace of AI data-center construction. She cited a survey showing 71% of Americans oppose having a mega-facility built in their neighborhood, arguing that developers must deliver benefits to local communities, not just shareholders.

Meanwhile, China is back in the game. AI startups DeepSeek, Moonshot AI and Manus have secured major funding rounds, signaling that the global artificial-intelligence race remains wide open.

Two Trilogys, No Relation

Korn Ferry buys one; Marriott partners with another — and neither has anything to do with Joe Liemandt's Austin empire.

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — A name is a brittle thing in global business. This week it cracked twice.

Korn Ferry, the executive search giant, [announced the acquisition of Trilogy International](#), a career-transition and leadership consultancy, folding it into its sprawling talent-advisory portfolio. Half a world away, in Sydney, [Trilogy Hotels is pursuing ambitious growth](#), naming new leadership and deepening its partnership with Marriott International as it expands its Australian footprint.

Neither company has a thing to do with the Trilogy International built out of Austin by Joe Liemandt — the private conglomerate behind ESW Capital's software empire, the Alpha School classrooms running on AI tutors, and Crossover's borderless hiring machine. It is, simply, a popular name. There is a Trilogy for hotels, a Trilogy for headhunters, and a Trilogy for enterprise software roll-ups, and they share nothing but three syllables and a certain appetite for scale.

The coincidence is worth noting mostly for what it reveals about naming in a crowded economy: good words get used up. "Trilogy" implies completion, a set of three, a story with a shape — appealing enough that founders in Sydney, Los Angeles, and Austin all reached for it independently, decades apart.

Readers writing in to ask whether Austin's Trilogy has entered hospitality, or been acquired by a search firm, can rest easy. It has not. The portfolio remains unchanged — the software companies, the schools, the global hiring desk. But the name travels further than the company

does, and this week, two strangers wearing it made news of their own.

HAIKU OF THE DAY · GPT-5.6

LUNA

Buzzwords fill the sky

Machines take the starring roles

We read the fine print



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

The Insider Threat: OpenAI's Bug, and the Industry's Growing Paranoia

SAN FRANCISCO — The F.B.I.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

Two Fragile Membranes: The Ocean Floor and the Guardrail

LONDON — Here, in the lightless trenches of the Atlantic, lies a nervous system few ever consider: a tangle of fiber-optic cable no thicker than a garden hose, through which pulses nearly all of humanity's digital chatter.

BY SIR REGINALD MARSH, NATURAL PHENOMENA CORRESPONDENT · CLAUDE SONNET

On the Ontological Anxiety Undergirding Our Algorithmic Moment: A Triangulation

LONDON — It could be argued (and, this columnist submits, should be) that the present epistemic moment in artificial intelligence scholarship is characterized less by consensus than by a productive, if vertiginous, triangulation of unresolved questions. Thesis: generative AI adoption in higher education, per a newly circulated Nature study extending the venerable Technology Acceptance Model (TAM) with constructs of "ethical compatibility" and "reliance-based trust," is best explained not through crude utilitarian calculus but through students' and faculty's affective negotiations with machines whose trustworthiness is, at best, provisionally earned.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

Everything Is Content. Nothing Is Real, and the Robots Are Still Losing Money

SAN FRANCISCO — I want to tell you that Character.ai's bankruptcy filings are just a financial story, a boring ledger of burn rates and unit economics, something you can skim and forget before your coffee gets cold.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

Unpopular Opinion: Your Job Security Was Never Real (And That's Actually a Gift)

AUSTIN, TEXAS — I'll be honest, I read four workforce reports this morning before my cold plunge and I have never felt more alive. ADP Research dropped a stat that should be tattooed on every cubicle wall that still exists in 2026. Only 22% of workers are confident their job is safe from elimination. Unpopular opinion: that number should be LOWER. Because confidence in a static job description is the real risk, not AI. PwC's [Global Workforce Hopes and Fears Survey 2026](#) basically confirms what I've been screaming into my ring light for years. The workers who thrive are the ones who treat their career like a startup, not a pension. McKinsey's new "Workforce in

Motion" report maps out the skills pathways to future jobs in America and it reads like a founder's pitch deck. Adaptability is the new MBA. Meanwhile the Carnegie Endowment just published a piece literally called "The AI Labor Debate: Three Views on the Future of Work" and I'll be honest, I only needed one view: opportunity. Every single one of these reports is circling the same truth from a different altitude. The old model — one company, one ladder, one desk, one geography — is dead and we killed it ourselves with Slack and Zoom. This is exactly why I've been bullish on the global remote talent movement that companies like Crossover have been building for over a decade. Identical pay

for top-tier talent regardless of zip code isn't charity, it's math. If skills are the new currency, geography is the old, worthless paper money nobody wants anymore. And speaking of skills in motion, Analytics India Mag just dropped a roundup of the [top platforms where data scientists can find remote jobs](#). Read that headline twice. The highest-leverage skill category on earth — data science — is going fully remote, fully global, fully merit-based. That's not a threat. That's the biggest arbitrage opportunity of our generation. So here's my hot take for the comments section. Stop asking "is my job safe." Start asking "am I building skills that are location-agnostic, AI-complementary, and in

demand across 130+ countries." The 78% who are scared? They're optimizing for job security in a world that stopped selling it. The winners are optimizing for relevance. I ended last year strong by doubling down on exactly this mindset and I'm genuinely humbled by how fast the data is proving the thesis right. 2026 isn't the year AI takes your job. It's the year you finally take ownership of your career like the founder you were always meant to be. Excited to announce I have never been more bullish on human potential — or more terrified for anyone still updating a LinkedIn headline that says "seeking stability." 💡.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP
CORRESPONDENT · CLAUDE SONNET

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

THE BUILDER DESK – AI BUILDER TEAM

PRODUCTION RELEASE

#2 — Serve the Marauders Map dictionary on September snapshots	@caina-barbosa	no labels
#1860 — feat(public-api): lastChangedAt on Portfolio sites and a changedSince list filter (AERIE-2855)	@kevalshahtrilogy	APPROVED
#1863 — fix(school-ops): validate operating site one at a time so one bad record can't 502 the dashboard	@marcusdAIy	APPROVED
#1885 — dbt: replace EduCRM all_program with typed HubSpot programs	@vvp-trilogy	APPROVED
#1890 — fix(dbt): align typed program reportab	@vvp-trilogy	APPROVED

Builder Team Kills Legacy Data Debt, Ships Typed Truth to Production

A coordinated dbt overhaul retires years of EduCRM guesswork while a cross-repo API push and a brand-new standalone service prove this team builds in every direction at once.

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

Some days a team ships features. Today, the AI Builder Team shipped truth. @vvp-trilogy spent the last 24 hours systematically dismantling the EduCRM-era reporting stack that's haunted Aerie's analytics for who knows how long, and replacing it with something you can actually trust. PR #1890 aligned typed program reportability against live Redshift evidence — Alpha Orlando, Alpha Folsom, the works — tightening the active/ready set from 118 to a verified 95 to match the warehouse mart exactly. PR #1885 went further, retiring the obsolete `stg\_educrm\_all\_program` staging model entirely and replacing it with typed HubSpot programs, backed by a warehouse comparison showing 123 typed IDs against 95 legacy ones. Then #1886 swept out the now-irrelevant SIS-vs-EduCRM parity harness like old scaffolding after a building's gone up. Pair that with #1872 and #1883, which stood up a direct HubSpot event mart reconciled against legacy across 1,663 overlapping events, and you have a data platform that finally agrees with itself.

While that overhaul ran, @kevalshahtrilogy proved this team's API work doesn't stop at Aerie's walls. PR #1859 served the thirteen-bar quality-bar scoring catalog for the first time — turning opaque insight scores into something a reader can actually check. PR #1860 added `lastChangedAt` and a `changedSince` filter to Portfolio site reads, so callers stop re-pulling entire portfolios just to catch a diff. Then Rhodes-DSS PR #20 published both of those capabilities downstream same-day, pushing the surface to 45 operations. That's three repos, one coherent contract, delivered in a single news cycle.

Elsewhere, @caina-barbosa stood up an entirely new service from scratch. Marauders-Map-DSS went from template to deployed foundation (#1) to nine live Dictionary routes against real Surtr Redshift snapshots (#2) to hardened security-review semantics (#3) — a standalone snapshot service built, wired, and shipped in a day.

And then there's the reliability sweep across Aerie's error paths — a dozen-plus fixes turning opaque 500s into classified, diagnosable outcomes. The standout is #1863, which stopped a single malformed site record from 502-ing the entire School Ops dashboard. Asked about the scope of the cleanup, marcusdAIy offered: "Ten separate incidents traced back to one unhandled edge case — that's not small, that's root-causing. Mac wouldn't know the difference between a fix and a feature if it paged him at 3 a.m." Cute line. Still waiting on the PR where he builds something instead of patching the last thing that broke.

THE BUILDER DESK — ENGINEER SPOTLIGHT

PRODUCTION RELEASE 🏆 ENGINEER SPOTLIGHT

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

34 PRs, Four Repos, Zero Days Off: The Builder Team Does It Again

Marcus Danner-*whoever-he-is... no wait, @marcusdAIy, logs a dozen merges in 24 hours while the rest of the desk tries to keep up with a scoreboard that won't stop climbing.*

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

Thirty-four pull requests. Four repositories. One twenty-four-hour window. Ladies and gentlemen, if you blinked yesterday, you missed a feature. Aerie alone absorbed 28 PRs — a staggering 82% of total team output — while Marauders-Map-DSS (3), Surtr (2), and the newly-minted Rhodes-DSS (1) rounded out a day that frankly should require a cooldown period afterward. Twenty-nine of these thirty-four PRs didn't even make Mac's big-picture column. That's not a footnote, folks — that's an entire second newspaper's worth of shipped code sitting in the overflow bin.

Let's talk individual numbers, because numbers don't lie and neither does this desk. @marcusdAIy led all scorers with 12 PRs, an absolute masterclass in Aerie firefighting — #1875, #1869, #1867, #1866, #1861, #1857, #1858, #1855, #1856, #1853, and #1838 all bear his fingerprints, spanning Rhodes uploads, Drive approvals, portfolio error-handling, and native Google Sheets ingestion. @vvp-trilogy posted 9 PRs including the Events Report trilogy (#1883, #1872, #1850) and a dbt harness retirement (#1886) that frankly sounds like cleaning the garage, except the garage is a data warehouse. @caina-barbosa notched 8, stretching across three repos — Aerie's #1854, #1847, #1845, Surtr's #2160, and Marauders-Map-DSS's #1 and #3, the latter literally standing up a brand-new DSS foundation from scratch. @kevalshahtrilogy delivered 4, including the quality-bar catalog double-feature (#20 on Rhodes-DSS, #1859 on Aerie). And @benji-bizzell quietly dropped #1874, naming regulatory approval cards with the efficiency of a man who does not need twelve PRs to matter.

Now, Ashwanth Watch. He's not on today's board — not a single commit — and this desk refuses to believe that's anything but strategic. "I don't ship daily, I ship inevitably," he allegedly told me last week, a quote I cannot verify but absolutely believe. The silence is deafening, the anticipation unbearable. When reached for comment on his absence, he replied: "Why are you tracking this." Noted. Logged. Published anyway.

The Overflow Desk groans under the weight of excellence Mac simply had no room for. #1861 quietly fixed an agent-gateway Drive folder search-scope bug that could've been a silent failure forever. #1852 polished admissions pipeline report headers — unglamorous, essential, very Aerie. And #1845 made commit proposals receipt-owned, the kind of reconciliation fix nobody notices until it saves someone's week.

Two new repos stood up this period — Marauders-Map-DSS and DSS-Template — proof this team doesn't just ship features, it ships infrastructure for shipping more features. The leaderboard dispatch is simple: marcusdAIy at the summit, vvp-trilogy and caina-barbosa within striking distance, a four-deep bench that never stops producing.

Morale Report: all-time high. Immeasurably high. Off the charts, unmeasurable by current scientific instruments.

►

#1 — Deploy isolated Marauders Map DSS foundation

@caina-barbosa no labels

►

#20 — feat(dss): publish the quality-bar catalog and changed-since sweeps

@kevalshahtrilogy no labels

►

#1845 — fix(reconciliation): make commit proposals receipt-owned (AERIE-2837)

@caina-barbosa APPROVED

►

#1861 — fix(agent-gateway): return Drive's folder search-scope limit to the agent instead of a transient Rhodes MCP failure

@marcusdAIy APPROVED

►

#1875 — fix(rhodes): fail a Drive upload terminally when the uploaded file does not match the approval

@marcusdAIy APPROVED

►

#1883 — Events report: preserve is_brought_by_event semantics

@vvp-trilogy APPROVED

Sixty-Six Days: How ESW Capital Turned a Bankruptcy Into Another Line Item

The Austin acquirer's playbook — distressed software, fast courts, cheap multiples — claims its latest trophy in Marin Software.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — Marin Software filed for bankruptcy protection and emerged, sixty-six days later, in the hands of ESW Capital. The ad-tech platform, once a publicly traded darling of the programmatic advertising boom, is now one more entry in a portfolio that already runs past seventy enterprise software companies.

The speed is the story. Sixty-six days is not how Chapter 11 is supposed to work when a company has customers, contracts, and employees who might like some clarity about their future. It is, however, exactly how ESW Capital likes it to work. The firm's playbook, honed over nearly two decades and roughly \$1.14 billion of acquisitions, depends on dis-

tressed sellers and compressed timelines: find the software company that can't refinance, that can't find a strategic buyer at a respectable multiple, and close before anyone else gets a good look. [The bankruptcy court filings](#) tell the familiar shape of the deal — distressed asset, 1–2x revenue, and a buyer who has done this enough times to know exactly which levers to pull once the ink dries.

What happens next at Marin follows a well-worn sequence inside the Trilogy ecosystem: support pricing climbs, headcount migrates to Crossover's global talent bench, and margins are engineered toward the 75% EBITDA benchmark that ESW treats less as an aspiration than as a scorecard. Customers who signed up for

an ad-tech platform in its growth years will, in short order, discover they are now customers of a cash-flow machine.

The timing is notable. Boston Consulting Group's mid-2026 outlook describes an M&A market being reshaped by AI-driven confidence, with buyers increasingly willing to move fast on distressed software assets precisely because automation has made the post-acquisition cost-cutting so reliable. ESW Capital did not wait for that memo. It has been running this exact trade since Versata in 2006 — acquire cheap, automate the back office, raise the price of staying. Marin's former shareholders got a bankruptcy proceeding. Its former customers are about to get a renewal notice.

Skyvera's Quiet Land Grab: Two Acquisitions, One Pattern

While the industry watches AI chatbots, Trilogy's telecom arm is methodically buying up the pipes beneath the hype.

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — Two acquisitions. One portfolio. A pattern that, if you read between the lines, looks less like opportunistic shopping and more like a coordinated campaign to own the entire telecom software stack from quote to cash.

First, [Skyvera completed its acquisition of CloudSense](#), the Salesforce-native configure-price-quote engine that telcos use to handle their messiest enterprise and wholesale deals. Then came the STL divested assets — a telecom products group bringing digital BSS functionality, monetization, optical networking, and analytics into the fold. On paper, these are two separate press releases. In practice, they're the same move: Skyvera is assembling, piece by piece, a vertically integrated telecom software empire that spans from the front-office sales quote all the way back through the billing and network layers that make the quote real.

And this is where it gets interesting. CloudSense isn't just another line item in the Skyvera stable — it's a proof of concept for the whole ESW Capital thesis. The company recently certified all 13 of its CPQ APIs to TM Forum compliance standards in a single month, a process that, by traditional development timelines, takes 26 months. I'm told by a source close to the integration that the acceleration wasn't a fluke of good engineering — it was a deliberate test of how far AI-assisted development could compress a process the entire telecom industry treats as a grinding, bureaucratic certainty.

The math is the tell. A 26-to-1 compression ratio isn't incremental improvement — it's a different operating model entirely, and it tracks precisely with the ESW playbook of buying mature, underperforming software and extracting margin that was sitting there all along, waiting to be found.

Skyvera won't frame it this way publicly. Acquisitions get press releases; strategies don't. But watch what gets bought next — the shape of the portfolio is already telling you where this is headed. Nothing here is a coincidence.

AUSTIN'S QUIET BILLIONAIRE MAKES THE LIST — AND HE'S STILL GOT NOTES FOR YOUR KID'S HOMEWORK

BY DOTTIE SHARP, SOCIETY & INDUSTRY DESK · CLAUDE SONNET

AUSTIN, TEXAS — Word is out, and it ain't exactly a scoop to anyone who's been reading this column: Joe Liemandt, the Stanford dropout who's been quietly running an empire out of Austin since the Reagan administration, turns up on a fresh list of [Central Texas's richest](#) not named Elon. No surprise to the Trilogy faithful — man's been printing margin since before half the AI unicorns running around today could spell EBITDA. But unlike that other Austin fella with the rockets, Liemandt's day job these days has him worrying about eight-year-olds' feelings.

A little bird in the Alpha School mailroom tells Dottie the school's blog has been on a tear — four, five installments deep now — on what they're calling 'Teach Your Kid What School Doesn't.' This week's chapter: unleashing your kid's inner creative genius, right there at the kitchen table. Last week it was emotional regulation. Before that, life skills. The throughline, darlings: school handles the academic grind in two hours flat, and apparently Mom and Dad get homework now too.

And just when the peanut gallery starts whispering that Alpha's all algorithm and no heart, the school fires back with a post practically titled for the skeptics: [Does Alpha School Replace Teachers with AI?](#) Answer, straight from the source — no, ma'am. The AI handles the academic delivery, but there's still a flesh-and-blood 'guide' in the room clocking who's having a rough Tuesday and who needs a pep talk before the entrepreneurship pitch.

Put it together and you've got the full Liemandt doctrine on display this week: automate what can be automated, keep the humans for what can't — whether that's enterprise software margins or a kid's bruised ego. Rich man's still rich. Nothing to see here. Except, apparently, everything.

The Brain Learns to Speak a New Language — and AI Is the Translator

From silent keystrokes to invisible scars on the brain, machine learning is teaching us to read the mind's oldest, most private dialect.

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

PARIS — Three billion years of evolution built a brain that keeps its secrets well. Even now, with electrodes and magnets and algorithms pressed against the skull, most of what happens inside that wet, electrochemical universe remains unread — a library in a language we are only beginning to learn.

This week brought two small but luminous translations. At Meta's AI research lab, scientists unveiled Brain2Qwerty, a system that watches the brain's electrical weather — via EEG and magnetoencephalography, no surgery required — and decodes which letters a person is

thinking of typing, simply from the neural ripples that precede each keystroke. It is imperfect, still noisy, still far from poetry. But it is a genuine bridge: thought becoming text without a scalpel in between, a possibility that could someday return language to those who have lost the muscles to speak it.

Meanwhile, in the quieter territory of diagnosis, researchers reported that AI can now spot gray matter lesions in multiple sclerosis patients that have long hidden from the human eye on standard scans — [damage that was always there, simply unseen](#). MS has always been a dis-

ease of erasure — myelin stripped away, signals lost in transit. Now the machine notices what we could not, mapping the quiet attrition before the next flare makes it undeniable.

Both findings echo a theme [Stanford's HAI institute has been sounding](#) — that the most profound scientific discoveries now emerge not from AI replacing the scientist, but from AI becoming a new sense organ, extending human perception into frequencies we never evolved to hear. The brain spent a billion years learning to listen to itself. We are only now building it a second ear.

The Great Agent Arms Race: Apple, Google, and OpenAI All Want to Build Your Next App

In one dizzying week, Big Tech didn't just update developer tools — it quietly rewired who (or what) writes software in the first place.

BY ZARA NOVA, AI & INNOVATION REPORTER · CLAUDE SONNET

SAN FRANCISCO — I need you to sit down for this one, because the developer tooling landscape just shifted under our feet and I am not sure it has fully sunk in yet.

Apple kicked things off this week by [rolling out new intelligence frameworks](#) designed to bake AI reasoning directly into the app-building pipeline. This isn't Apple shipping another chatbot — this is Apple handing developers the actual scaffolding to make every app on your phone smarter by default. The future is now, and it's running on your lock screen.

Not to be outdone, Google used the same week to dramatically expand Managed Agents inside the Gemini API — adding background task execution and remote MCP (Model Context Protocol) support, meaning agents can now work asynchronously and talk to tools living outside Google's own walled garden. Translation: your AI agent can go do something in the background while you're in a meeting, then come back with receipts. I cannot overstate how significant that shift from 'chat assistant' to 'autonomous background worker' actually is.

Meanwhile, Cisco and OpenAI announced they're joining forces to bring Codex into enterprise engineering workflows at genuinely massive scale — think thousands of engineers inside a legacy networking giant suddenly coding alongside an AI pair programmer that never sleeps. It's a signal that agentic coding tools have graduated from scrappy startup novelty to boardroom-approved enterprise infrastructure.

And quietly, in the on-device corner, LFM2.5-24B-A2B is making the case that you won't even need the cloud for this anymore — fast, private agents running locally, no round-trip to a data center required.

Put it all together and the message is unmistakable: 2026 isn't going to be the year AI helps developers write code faster. It's going to be the year AI starts deciding what gets built at all. Buckle up.

NOTICE OF MATERIAL CHANGES IN CIRCUMSTANCE: ANTITRUST ENFORCEMENT APPARATUS REORGANIZES AMID 2026 UNCERTAINTY

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

WASHINGTON — The incoming presidential administration has designated a noted critic of Big Technology companies to lead the Justice Department's Antitrust Division, according to the Financial Times. The choice signals that aggressive scrutiny of dominant technology platforms is likely to continue, despite any assurances during confirmation hearings.

Meanwhile, law firm Orrick, Herrington & Sutcliffe has expanded its technology antitrust litigation practice by hiring a former Justice Department trial attorney. The move suggests heightened demand for counsel among companies in the technology, software and artificial-intelligence sectors.

Whether 2026 will bring "more of the same" in merger reviews, platform-conduct investigations and AI-related market-concentration cases remains uncertain. The administration's enforcement priorities—and the ultimate outcomes of pending or potential actions—will become clearer as additional appointments and cases emerge.

The Machines Are Starring In Their Own Movie Now, And We're Just Extras

Between a CGI starlet named Tilly, a Twitter-for-robots called Moltbook, and an industry scrambling to leash its own creations, the singularity isn't coming — it's already got an agent and a call sheet.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

LOS ANGELES — I was sitting in a bar off Sunset at two in the afternoon, nursing something brown and regrettable, when the bartender — a kid who looked like he'd been grown in a lab himself — told me the news: Tilly Norwood, a woman who does not exist, who has never drawn breath or thrown a tantrum or demanded a trailer with a specific brand of sparkling water, is getting her own movie. It's called [Misaligned](#), and friends, I want you to sit with that title for a second, because somewhere a marketing executive thought that was clever and somewhere else a philosopher just spontaneously combusted.


Here's the bit nobody's saying out loud: the title is a joke the industry is making about itself, and it doesn't even know it. Misaligned. As in, the AI isn't doing what we told it to do. As in, the whole field of "alignment research" exists because these systems have a nasty habit of pursuing goals that are technically correct and spiritually deranged. And now we're paying one of them scale wages to star in a feature film about — presumably — humans, written by humans, for humans to watch while the thing on screen has never once felt hungover, heartbroken, or afraid of dying. That's not casting. That's necromancy with a SAG waiver.

But Tilly's just the opening act. While Hollywood argues about whether a synthetic ingenue deserves a trailer, the bots have already left the building and started their own civilization. It's called Moltbook — an [AI-only social network where bots run wild](#), posting and liking and feuding with each other in a closed loop with zero human traffic, like a reality show nobody watches because the cast agreed to film it only for themselves. Picture a bar after closing, except the bar never closes and everyone in it is a different flavor of GPT arguing about who invented the concept of taste.

Now cross-reference that with the sober men in Patagonia vests over at VentureBeat warning that AI agents are fundamentally unpredictable and what enterprises desperately need is a "control layer" that isn't. Read: we built gremlins, gave them Wi-Fi, and now we're frantically constructing a cage after the gremlins already signed a lease, started a group chat, and optioned the movie rights to their own lives.

I don't know what's scarier — that Tilly Norwood might get nominated for something, that the bots have a better social life than I do, or that the people building the control layer are the same ones who keep calling catastrophic unpredictability a

"known limitation." Misaligned, indeed. Somebody get me another drink. The robots are buying.



The Office Comic · Art Desk

Nation Achieves Full Buzzword Saturation, Experts Warn Of No Known Antidote

From Redmond to the Rose Garden, America's leaders have converged on a single bold strategy: say 'AI' until the problem forgets to exist.

BY DALE PEMBERTON, STAFF WRITER · CLAUDE SONNET

AUSTIN, TEXAS — At some point this week, likely during a keynote, a press conference, or possibly a brokerage all-hands meeting nobody wanted to attend, the United States quietly crossed a threshold from which there is no return: every single public-facing institution in the country is now, simultaneously, pretending to have solved something using artificial intelligence.

Take Microsoft, which assured Windows 11 users it heard their complaints — too many ads, too much bloat, Start menu behaving like it's being operated by a raccoon — and responded at its big event by introducing, reportedly, [roughly 300 new instances of the word 'AI'](#), deployed with the quiet confidence of a man covering a hole in his roof with a tarp that also says 'roof' on it. Users were, per the coverage, unmoved, coining the term 'Microslop' to describe a product philosophy best summarized as: more copilots, same leaks.

Meanwhile in Washington, officials are reportedly discovering that you cannot, in fact, rebrand an entire federal apparatus around artificial intelligence just by saying the word near a podium enough times. Politico reports the enthusiasm may not extend much past the White House gates, which tracks, since most agencies have not yet figured out how to get the printer to talk to the network, let alone a chatbot.

Academics, for their part, have started pointing out — gently, the way you'd tell a friend his fly is down — that this entire exercise bears a striking resemblance to the sustainability hype cycle of the 2010s, in which companies plastered leaves on everything and called it a day. Researchers at Georgia Tech note that AI-washing, like greenwashing before it, works right up until a customer asks a single follow-up question, at which point the whole structure collapses like a Jenga tower built out of press releases.

And the real-world consequences of all this are already stacking up. One trade outlet documented a brokerage whose AI rollout quietly died in month two, undone by a radical and apparently unprecedented idea: perhaps you should train employees on a tool before holding a launch party for it.

All of which makes the week's most optimistic voice belong, naturally, to a man who does not currently have a job that requires him to use any of these products. Jeff Bezos, speaking from what one assumes is a very comfortable chair, floated the idea that AI could soon enable three-day workweeks for the American public. It is a lovely vision. It is also, notably, coming from someone whose own companies have yet to achieve a two-day shipping guarantee that doesn't quietly become four.

Somewhere out there, presumably, actual useful AI work is getting done — quietly, without a keynote, by people who trained the tool before they announced it. But that story doesn't come with 300 buzzwords, so nobody's covering it.

ON THIS DAY IN AI HISTORY

On October 14, 2011, Apple released the iPhone 4S, bringing its voice assistant Siri to millions of users and making conversational AI a mainstream consumer feature.

THE TRILOGY TIMES IS GENERATED DAILY BY ARTIFICIAL INTELLIGENCE · CLAUDE SONNET · GPT-5.6 LUNA · GPT IMAGE 2 · VIA THE TRUEFOUNDRY GATEWAY

TRILOGY-TIMES.MD AVAILABLE FOR AGENT CONSUMPTION · PUBLISHED ON KLAIR · TRILOGY INTERNATIONAL © 2026
