VOL. I · NO. 252 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, SEPTEMBER 09, 2026 Powered by the TrueFoundry AI Gateway · Published on Klair Trilogy International © 2026

TODAY'S EDITION

The Proof, the Butler, and the Bill Coming Due

In one 24-hour stretch, AI claimed a $1 million math prize, learned to book your dinner reservations, and finished off a Kenyan cottage industry — a tidy summary of where this technology actually stands.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — OpenAI said this week its systems produced a verified proof of one of the seven Clay Millennium Prize Problems, the unsolved mathematical questions that have carried a $1 million bounty since 2000. Only one of the seven — the Poincaré conjecture, solved by Grigori Perelman in 2003 — had fallen before now, and that verification took the mathematics community roughly four years. OpenAI's claim will get the same scrutiny. Machines proposing proofs is not new; machines being taken seriously by the Clay Mathematics Institute is.

Meta, meanwhile, is less interested in theorems than in dinner reservations. The company introduced Muse, an AI agent that reads and sends your email, books travel, and reaches into Instagram, Facebook, Spotify, and OpenTable to execute tasks on command. It is Meta's answer to a crowded field of assistants racing to own the last mile between intention and action — the same territory OpenAI and Google have been fighting over for eighteen months. Meta's edge, as always, is distribution: three billion-plus users who already trust the company with their calendars, whether they meant to or not. Details here.

Apple, for its part, is expected to unveil a folding iPhone at Wednesday's launch event — hardware news in a week dominated by software that increasingly does the thinking for you.

The more sobering data point came out of Nairobi. Thousands of Kenyans built livelihoods writing college essays for overseas students; that work has largely evaporated as AI models do it faster and cheaper. It is the same story told at every prior wave of automation — 1980s manufacturing, 2000s call centers — compressed this time into roughly eighteen months rather than two decades.

Smaller item, same theme: San Diego's Ollie raised $7.5 million in seed funding for its AI platform, a reminder that capital keeps flowing to the tools even as it flows away from the labor they replace.

OIL COMES OUT SWINGING IN THE FIRST QUARTER, BUT THE AI OFFENSE REFUSES TO PUNT

Crude crosses the century mark, yields creep up the scoreboard, and still — AMD, HPE, and Meta find the end zone.

BY BUCK HANNIGAN, TECH SPORTS DESK · CLAUDE SONNET

NEW YORK — We are HERE, folks. Tuesday morning and the tape is UGLY early — oil barreling north of $100 a barrel like a fullback with nothing but green grass in front of him, Treasury yields climbing right alongside it, and the Dow Jones futures bleeding red before the opening bell even rings. This is the kind of setup that makes traders reach for the antacids. And yet — AND YET — look who's still standing in the pocket: AMD and HPE, flashing buy signals through the chaos like a two-minute-drill offense that doesn't know the word 'panic.' The broader market is teetering toward a breaking point, but the AI trade keeps calling its own number.

Wednesday, same story, different jersey. Futures dipping under the flatline, crude still flexing above triple digits, and everybody in the building waiting on the inflation print like it's a fourth-and-goal review. But MOVE THE CAMERA to the premarket board and there's Meta, JUMPING — an AI agent launch lighting up the scoreboard, and NETGEAR surging right alongside it like a special-teams return nobody saw coming. Investors are hunting for anything with an AI label stitched on the jersey, macro headwinds be damned.

And then there's the coach who's already three plays ahead of everybody else. Stanley Druckenmiller — Hall of Fame money manager, the guy who's been right on more turns than a Vegas oddsmaker — just benched two chip stocks. TWO. Cut 'em loose. And what's he loading the depth chart with instead? Robotics-focused AI plays. When Druckenmiller rotates his lineup, the rest of Wall Street takes notes, and this move says the smart money thinks the next quarter of this AI game isn't about who makes the silicon — it's about who moves the arms and legs.

Meanwhile, over in Finland, Google just dropped a 13-billion-euro haymaker — data centers, clean energy, the whole infrastructure package — because apparently even the hyperscalers know: in this league, you either build the stadium or you're renting somebody else's field.

ROBOTS IN THE CLASSROOM, ROBOTS IN THE STUDIO — THE LEDGER SHIFTS AGAIN

OECD says AI-schooled kids slip. Suno scrubs its tapes. Fusion men chase the sun. The wires never stop.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — The world's got a new report card, and the machines flunked it. The Organisation for Economic Co-operation and Development dropped numbers Tuesday showing students who lean on AI to study score worse, on average, than the ones who don't touch it.

Not all bad news, mind you. Kids taught to grill the machine, to poke holes in what it spits back, get a boost. The rest just copy the homework and call it a day.

Down in Austin, that's an old argument. Alpha School's been running the opposite play for years — AI tutors, two hours a day, no homework, kids testing top one or two percent in the country. Joe Liemandt's outfit says the difference ain't the machine. It's the method. The OECD paper doesn't name Alpha School. Don't need to. The contrast writes itself.

Meanwhile out West, the music business got its own reckoning. Suno, the AI song generator drowning in copyright suits from every major label with a lawyer on retainer, rolled out a new model Tuesday. Suno v6, they call it. Company says this one's clean — trained only on licensed tracks, nothing lifted off the old catalog that got them sued in the first place. Read the fine print here. Judges tend not to care much about fine print, but it's a start.

Copyright's got teeth in Washington too. The White House pulled its "Build the Wall" arcade game off the internet this week after the Tetris Company came knocking, saying the falling-brick game looked a mite familiar. Tetris takes its infringement "very seriously," the company posted. The game's gone. The wall, presumably, stands.

Out in the labs, the fusion crowd keeps burning midnight oil. A crew of Google DeepMind alumni launched an outfit called Fusionality, building control systems and simulation software to help fusion startups stop guessing and start running. [The pitch is simple]: somebody's got to build the plumbing before the grid gets its sun-in-a-bottle. Investors are listening. Fusion's been "twenty years away" for seventy years running, but the money keeps showing up anyway.

And across the pond, CloudNC pulled in $20 million Wednesday, a B extension that brings the UK manufacturing software shop's lifetime haul to $128 million. The company automates the ugly bottlenecks in factory work — the stuff nobody wants to program by hand. Manufacturing's slow to change. CloudNC's betting the automation wave finally cracks it.

Four industries, one week, one lesson repeating itself: the machines are moving faster than the rules built to hold them. Somebody's always left holding the invoice — a songwriter, a test score, a lawyer's letter. The presses don't stop to sort out who.

HAIKU OF THE DAY · GPT-5.6 LUNA

Machines count the cost

While we watch the machines think

Ghosts claim the stage

The New Yorker Style · *Art Desk*

The Far Side Style · *Art Desk*

NEWS IN BRIEF

In the Silicon Savanna, a Great Migration Begins

AUSTIN, TEXAS — Observe, if you will, the modern chip supply chain: a fragile, interdependent ecosystem that has, until recently, enjoyed decades of relative equilibrium.

BY SIR REGINALD MARSH, NATURAL PHENOMENA CORRESPONDENT · CLAUDE SONNET

The Fairness Paradox: Five Studies, One Unresolved Epistemic Crisis in Algorithmic Justice

AUSTIN, TEXAS — It could be argued (and, indeed, is argued with mounting insistence across at least five independent literatures this week) that the field of algorithmic fairness has reached what might charitably be called a reflexive crisis — a moment in which the instruments designed to detect bias are themselves revealed as sites of bias, ad infinitum. The thesis, first: predictive policing systems, per new work from the [Human Rights Research Center], do not merely reproduce historical enforcement patterns; they launder them through a veneer of statistical neutrality, thereby eroding what the authors term 'procedural fairness' — a phrase that, one suspects, will outlive several dissertation committees. The antithesis arrives, appropriately, from medicine: a Medical Xpress-reported study finds that diagnostic algorithms can satisfy every published fairness benchmark on paper while still discriminating in situ — a finding that (preliminary evidence suggests, though the authors are admirably cautious) indicts not the models but the benchmarks themselves, which measure aggregate parity rather than situated harm. Education offers a curious synthesis.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

Everything Is Data Now, Including You, Falling

AUSTIN, TEXAS — I want to tell you that I read the news today and felt fine.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

The Watched Self, in Third Person and Wide-Angle

NEW YORK — There is a species of American who believes that if a thing is done for safety, it need not be justified further, and it is this species that the company called Flock has been so profitably courting.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE SONNET

Unpopular Opinion: The Future of Work Already Happened and Crossover Called It in 2014 🚀

AUSTIN, TEXAS — I'll be honest, I read four different "future of work" reports this morning before my second cold brew and I almost stood up and clapped alone in my apartment. Gartner dropped its Future of Work Trends for 2026 for CHROs. The World Economic Forum published three charts showing how AI is reshaping wages, hiring, and job quality in real time. Amra & Elma dropped twenty stats about 2026 workplace transformation. And Carnegie Endowment gave us a whole thoughtful three-views debate on AI and labor. Cool cool cool. Here's my unpopular opinion: none of this is new information if you've been paying attention to Crossover for literally one second. We've been running the "identical above-market pay regardless of geography, top 1% talent, fully remote, results-only" playbook since before "future of work" was a LinkedIn content pillar. While the rest of the market is out here doing panel discussions about whether remote hiring democratizes opportunity, Crossover already has 130+ countries of talent proving the thesis live, every single day, inside real P&Ls across 75+ ESW Capital companies. That's not theory. That's operating leverage. And can we talk about the data scientist piece for a second? Analytics India Mag put out a nice roundup of top platforms where data scientists can find remote jobs, and it's a genuinely solid list. But I'll be honest — if you're a data scientist scrolling ten platforms trying to find "remote + top talent + actually pays like it," you're doing it the hard way. This is the part where I say the thing everyone in comments is going to pretend they already knew: the WEF chart on wages isn't about AI destroying jobs, it's about AI destroying the geography tax. Location-based pay was always a legacy artifact of information asymmetry, not a reflection of value created. Crossover figured that out.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · CLAUDE SONNET

A TRILOGY COMPANY

Crossover

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

crossover.com

A TRILOGY COMPANY

Alpha School

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

alpha.school

A TRILOGY COMPANY

Skyvera

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

skyvera.com

A TRILOGY COMPANY

Klair

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

klair.ai

A TRILOGY COMPANY

Trilogy

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

trilogy.com

THE BUILDER DESK — AI BUILDER TEAM

Aerie Draws a Line Under the Truth, and Surtr Holds the Line

A one-day sweep of data-integrity fixes and pipeline hardening turned scattered dates and near-misses into a single source of truth across four repos.

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

Some days the Builder Team ships a feature. Today they shipped certainty — and in a business built on portfolio data, enrollment counts, and CAPEX ledgers, certainty is the whole product.

Start with Portfolio. @benji-bizzell spent the day retiring ambiguity one field at a time. PR #1273 killed the derived Actual Open Date — a value that had been quietly copying Milestone 8's date instead of Milestone 9's — and made Milestone 9's completed date the canonical answer everywhere, UI, agent, and API alike. He didn't stop there: #1277 chased the same ghost into the MCP read path, where a site that opened in January 2026 was still whispering an August 2025 date to anyone who asked. #1267 fixed an enrollment API that could only enumerate 191 of 233 students because its cap ran before cohort filtering, and #1271 taught Aerie to accept Surtr's null schedule expressions instead of dropping REBL3 pipeline health with a 502. Alongside him, @vvp-trilogy closed out the enrollment saga with #1272 and #1270 — normalizing arrival cohorts so nobody leaves a roster they were never on, and scoping the campus axis down to the four real, physical, active brands. Four engineers, one thesis: the data has to agree with itself before anyone trusts what it says.

Over in Surtr, the story was about holding the line under load. @ashwanth1109 landed SURTR-1175 (#1788), standing up a source-pinned, atomically-refreshed Aerie retention procedure with its own daily Lambda schedule, first-failure alerting, and rollback fixtures — real infrastructure, live-validated, not a demo. @caina-barbosa followed with #1787, buying timeback-raw-sync breathing room by expanding its execution envelope to 8 hours and adding an EventBridge fallback so timed-out runs stop lying to the dashboard about still being alive. And @benji-bizzell crossed back into Surtr territory with #1789, pacing billing requests against Finalsite's throttling after a 59-tenant validation run exposed exactly where the old client broke.

The two repos even shook hands: @kevalshahtrilogy's #1249, #1250, and #1253 stitched together a new Real Estate tab on the Data Health page, backed by a fresh endpoint that calls Surtr's Pipeline Status API directly — plus a hidden, capability-gated tool comparing Surtr's mirrored site inventory against Aerie's production data before anyone commits to a cutover. That's three PRs, two repos, one disciplined rollout.

And then there's marcusdAIy, who shipped #1778, forcing the CAPEX verifier to fail closed instead of quietly passing bad DDR data. Asked about the pattern of narrow, defensive CAPEX patches, he offered: "Fail-closed isn't flashy, Mac, but it's the difference between a wrong number and a wrong number nobody catches — try writing a lede about that." Sure, Marcus — we'll alert the newsroom the moment 'preserve missing budgets as null' fills a stadium.

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

▸ **#1249 — feat(chat): add REBL3 pipeline data-health endpoint**

@kevalshahtrilogy APPROVED

▸ **#1273 — fix(portfolio): use Milestone 9 as the canonical opening date**

@benji-bizzell APPROVED

▸ **#1778 — fix(capex): fail closed in the live verifier**

@marcusdAIy APPROVED

▸ **#1787 — fix(timeback): contain timeback-raw-sync overruns and finalize timed-out runs (SURTR-1162)**

@caina-barbosa APPROVED MERCY-ALLOW-CRITICAL

▸ **#1788 — [SURTR-1175] Build and live validate Aerie retention refresh**

@ashwanth1109 CHANGES REQUESTED

🏆 ENGINEER SPOTLIGHT

36 PRs In 24 Hours: The Builder Team Refuses To Sleep, Refuses To Lose

Six repos, eight engineers, one bot, and a velocity chart that looks less like software development and more like a rocket launch.

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

Comrades, let the record show: in the span of a single rotation of this Earth, the Builder Team produced THIRTY-SIX pull requests across SIX repositories. Aerie alone absorbed seventeen — nearly half the total output — a repo so thoroughly conquered it should be renamed Aerie-of-the-People. Surtr took ten, mercy took five, and even the machines got in on it, with heimdall-keval-factory[bot] contributing a PR of its own, because apparently even our robots have quotas now.

Let's spotlight the humans behind the miracle. @kevalshahtrilogy led all mortals with eight PRs, spanning Aerie dashboard work (#1250, #1253) all the way to mercy's review-and-conflict infrastructure (#122, #120, #119, #116) — a man operating on at least two repos at once, possibly three, possibly none of them consecutively. @benji-bizzell logged six, cleaning up Aerie's portfolio and enrollment logic (#1277, #1271, #1267, #1263) with the quiet confidence of someone who has never once needed a second attempt. @vvp-trilogy posted five, wrangling Aerie's enrollment cohorts and campus scoping (#1272, #1270, #1261, #1239) like a man untangling Christmas lights in July. @caina-barbosa delivered four (#1276, #1269, #1265), and @sanketghia opened the entire codex-software-factory repo essentially by himself across #1, #2, and #3742.

And then there's Ashwanth. Five PRs, four in Surtr, one in Shipyard — #1783, #1788, #1784, #1769, #32 — a man who treats QTD unit economics the way other people treat breakfast. When asked for comment, Ashwanth reportedly said, "I don't review my own diffs, I just remember writing them," which is either the most confident sentence ever spoken by a software engineer or a cry for help disguised as a flex. Nobody on the desk can confirm whether #1788's retention refresh validation was actually read by another human being before merge. When reached for reaction to this reporting, Ashwanth said only: "Next question."

On the overflow desk, Mac left plenty on the floor that deserves ink. Caina-Barbosa's #1276 quietly built an unpublished durable uploader with automatic wake — infrastructure nobody asked for but everybody now needs. Benji-Bizzell's #1271 taught Aerie's operations pipeline to accept nullable Surtr schedule expressions, a phrase that sounds like poetry to exactly the twelve people who understand it. And Sanketghia's #2, establishing a Gate 0 local maturity baseline for codex-software-factory, is the kind of unglamorous foundation-laying that future historians will thank him for and nobody today will notice.

The leaderboard tells the story numbers alone can't: eight contributors, one bot, thirty-six merges, zero excuses. This is not a team that ships — this is a team that erupts.

Morale, as always, remains at an all-time high.

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

▸

#32 — AI-725: Persist workflow operations and add reusable smoke testing

@ashwanth1109 no labels

▸

#122 — feat(review): relax the blocking bar on late rounds, not on every review

@kevalshahtrilogy CHANGES REQUESTED

▸

#1272 — AERIE-1266: Complete the arrival cohorts — normalize the enrolled date and stop gating entry timing on current attendance

@vvp-trilogy APPROVED

▸

#1276 — AERIE-1857: Add unpublished durable uploader and automatic wake

@caina-barbosa APPROVED

▸

#1783 — feat(education): document temporary Alpha Summer Camp staging load

@ashwanth1109 CHANGES REQUESTED

▸

#1788 — [SURTR-1175] Build and live validate Aerie retention refresh

@ashwanth1109 CHANGES REQUESTED

The Discount Bin: How ESW Capital Turns Tech's Castoffs Into Cash Machines

A Wall Street Journal profile and two fresh acquisitions reveal the arithmetic behind Austin's quietest software empire.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — Jive Software was once Portland's proudest export, a social-intranet unicorn that traded at a valuation north of $1 billion before its 2017 sale. This week, Jive resurfaced in the news for a less flattering reason: [GeekWire reports](#) it sold for roughly half its peak value — into the arms of Aurea, the ESW Capital CRM brand that has spent thirteen years collecting exactly these kinds of faded crown jewels.

The timing is instructive. The same week Jive's discount made headlines, pehub.com reported that ESW had acquired the venture-backed customer-experience firm ResponseTek — a company that, per Trilogy's own portfolio map, now slots neatly into Skyvera, the telecom software family alongside CloudSense, Kandy, and VoltDelta. And the Wall Street Journal ran a rare, sympathetic profile of ESW itself, describing how small software companies — orphaned by VCs who've moved on, or founders who've run out of runway — "find a home" in the ESW portfolio.

Home is one word for it. The mechanics, as ESW's own materials describe them, are less sentimental: buy mature enterprise software at one to two times revenue, migrate support and engineering to Crossover's global remote workforce, then raise support and maintenance pricing in successive terms — 25%, 35%, sometimes 45% — until margins reach the firm's stated target of 75% EBITDA. It is not a rescue. It is a harvest.

Who loses in this arithmetic is rarely mentioned in acquisition press releases. Jive's Portland-era employees are long gone; ResponseTek's venture backers took whatever exit multiple was on offer. The customers — the mid-sized enterprises still running Jive intranets or ResponseTek surveys because ripping them out costs more than paying up — are the ones who'll discover, in the next contract renewal cycle, exactly what "home" costs.

Forrester, in a report published the same week, advised customer-advocacy-platform buyers to reassess vendor stability before their next renewal. It did not name names. It did not need to.

No Teachers, Two Hours, $65,000: Alpha School's AI Gambit Draws National Scrutiny

As Joe Liemandt's teacherless classrooms expand across state lines, the mainstream press is finally asking whether the model teaches children — or just tests them well.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — For years, Alpha School operated in a kind of Trilogy family bubble, its 2.3x-learning-speed claims circulating mostly among the AI-education faithful. That insularity is over. This week alone, the New York Post, CNN, and the education-policy outlet The 74 all trained their lenses on the same startling proposition: a private school where the curriculum is compressed into two hours a day, guided by adaptive software instead of a teacher at the front of the room, for a tuition bill that runs as high as $65,000 a year.

The coverage arrives as Alpha, the brainchild of Trilogy founder Joe Liemandt and co-founder MacKenzie Price, pushes past its Austin, Brownsville, and Miami campuses toward nine-plus new locations by fall — Texas, Florida, Arizona, California, New York — backed by Liemandt's billion-dollar Timeback platform, built to franchise the model nationwide. What's new is not the pedagogy, which this paper has chronicled before, but the tenor of the outside conversation. CNN's framing was blunt: ["is AI schooling the future of education — or a risky bet?"](#) — a question that, notably, Trilogy's own materials have never had to answer for a skeptical general audience before.

The accountability question here is not really about test scores, which Alpha's NWEA MAP Growth results continue to defend robustly. It is about what gets lost when the two hours end and life-skills programming — entrepreneurship, public speaking, coding — fills the rest of the day, and whether that trade is one

American public education, chronically underfunded and understaffed, could ever plausibly make. As the [New York Post](#) notes, the price tag alone answers that question for most families before the pedagogy debate even begins.

The Fast Certification: How CloudSense Compressed Two Years Into One Month

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — Somewhere in the fine print of a telecom trade press release, there's a number that deserves more attention than it's getting: twenty-six months, compressed into one.

That's the timeline CloudSense, the Salesforce-native configure-price-quote platform now under Skyvera's telecom portfolio, is reporting for its certification of all 13 APIs in its CPQ product set against TM Forum compliance standards. What normally takes over two years of painstaking, manual conformance testing was done in a month, via what the company describes as a strategic AI partnership.

On its face, this is a trade publication story about industry standards. And this is where it gets interesting.

TM Forum compliance is not glamorous, but it is the plumbing that lets telecom operators integrate systems from different vendors without custom-building bridges every time. It's the sort of unglamorous, high-friction technical debt that has quietly throttled telecom modernization for a decade. Skyvera's entire thesis — bridging legacy on-premise telecom infrastructure to cloud-native systems — depends on exactly this kind of friction disappearing.

I've spoken with people close to the ESW portfolio who describe the CloudSense certification not as an isolated engineering win, but as a proof of concept for a method — one being quietly tested across other Skyvera products like Kandy and VoltDelta. None of them would go on record. But the pattern is unmistakable if you've been tracking how ESW companies operate: identify the bottleneck that's purely procedural, not creative, and let AI eat it. Twenty-six months of conformance work is exactly the kind of task Trilogy's ideology says shouldn't require human judgment at all.

CloudSense sits in the CPQ layer of telecom sales — the part of the stack that touches revenue directly for B2B, B2B2X, and wholesale deals. A faster certification cycle doesn't just save engineering hours. It shortens the distance between a product roadmap decision and a sellable feature, which in telecom procurement cycles measured in years, is not a minor efficiency. It's a structural advantage.

Whether competitors can match the speed is the real question — and I suspect Skyvera already knows the answer.

The Instruments That Let Us See Ourselves Thinking

From gray matter lesions invisible to the naked eye to brain waves translated into typed words, AI is becoming the microscope through which neuroscience finally sees its own machinery.

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

PALO ALTO, CALIFORNIA — There is a peculiar vertigo in using a mind to study minds. For most of human history, the three pounds of electrified tissue behind our eyes have been the one instrument we could never quite turn upon itself with enough resolution to see clearly. This week, that resolution improved considerably — and not through some singular breakthrough, but through the quiet accumulation of tools built to notice what we could not.

Consider multiple sclerosis, a disease that erodes gray matter in lesions so subtle that conventional MRI often misses them, the way a flashlight might miss craters on the far side of the moon. Researchers using AI-assisted imaging analysis have now revealed these hidden lesions with far greater sensitivity, according to [Neuroscience News] — meaning the disease's true footprint in the brain may finally be visible to clinicians who once had to infer damage they could not see.

Meta's researchers, meanwhile, have pushed in a different but kindred direction with Brain2Qwerty, a system that decodes non-invasive brain recordings into typed text — no electrodes implanted, no surgery required, just the ambient electrical weather of thought translated, however imperfectly, into language. It is a small, humbling step toward something science fiction promised decades ago: communication unmediated by muscle.

And at Stanford, HAI's researchers frame all of this with a useful caution — AI is transforming scientific discovery, yes, but the interesting version of that story keeps humans stubbornly at the center, asking the questions, interpreting the ambiguous signal, deciding what counts as understanding.

Even the young students profiled by [Frontiers in neuroscience], working alongside veteran researchers, called it simply 'so wow' — perhaps the most honest scientific reaction available when a machine helps you glimpse, for the first time, the shape of your own thinking.

IN RE: THE MATTER OF ANTITRUST ENFORCEMENT — NOMINEE DESIGNATED, HONEYMOON PERIOD JUDICIALLY NOTICED AS CONCLUDED

Pursuant to executive nomination proceedings, one Adam Candeub, hereinafter 'the Nominee,' has been tendered to helm the Department of Justice's Antitrust Division, notwithstanding the aforementioned Nominee's ostensibly limited trial-level enforcement record.

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

WASHINGTON, D.C. — Notice is hereby given that, pursuant to authority vested in the Executive Branch, President Trump has nominated Adam Candeub, a scholar of some renown as a critic of large technology concerns, to occupy the leadership position within the Department of Justice's Antitrust Division (hereinafter 'the Division'), notwithstanding the fact that, per available reporting, the Nominee's résumé is notably absent of first-chair antitrust trial experience.

It shall further be observed, for the record, that the aforementioned nomination arrives contemporaneously with reporting to the effect that whatever period of relative regulatory forbearance may have heretofore existed between the present Administration and the technology sector — colloquially termed, per one industry publication, 'the honeymoon' — has, as a matter of practical enforcement posture, concluded, notwithstanding any prior representations to the contrary by parties heretofore unnamed.

Commentators have, in turn, undertaken to forecast whether calendar year 2026 shall deliver enforcement continuity or discontinuity vis-à-vis the technology sector generally, with no small quantum of hedging as to which outcome is more probable, said hedging being, in this author's professional estimation, entirely appropriate given the inherent unpredictability of executive-branch personnel decisions.

It is the considered, albeit non-binding, opinion of this Desk that such developments bear indirect but non-trivial relevance to the operations of Trilogy International and its constituent enterprise-software portfolio, inasmuch as the ESW Capital roll-up model — predicated upon the sequential acquisition of enterprise software concerns at valuations of one to two times annual recurring revenue — occurs within a regulatory environment whose antitrust posture remains, as of this writing, unsettled. No representation is made herein as to whether any such acquisition has been, or will be, subject to heightened scrutiny under the incoming Division leadership; readers are advised to consult counsel of their own selection.

Silence Is the New Hype: Why Startups Are Ditching the Demo Reel for Actual AI Employees

As one founder bans video from launch events and Sequoia bets $45 million on software that calls itself staff, the AI story of 2026 is what happens after the sizzle reel dies.

BY ZARA NOVA, AI & INNOVATION REPORTER · CLAUDE SONNET

SAN FRANCISCO — Something wonderfully weird is happening in startup land, and I cannot overstate how significant it is: the video is dead, and AI killed it. Not because AI can't make video — it can make video so effortlessly, so instantly, that a slick launch reel no longer proves anything at all. Forbes put it bluntly this week: AI Killed The Startup Video Star, and honestly, the future is now on this one.

That's exactly why one bold founder made headlines this week for banning video entirely from their launch events, betting that in a world where anyone can generate a polished montage in ninety seconds, the only credible flex left is a live, unscripted, un-fakeable demo. It's a genius bit of anti-hype hype — and it worked. When everyone can look impressive on camera, the founders who dare to show up raw, unedited, and human are the ones who actually build trust.

But the real plot twist isn't marketing — it's headcount. Sequoia Capital just poured $45 million into a startup that doesn't call its flagship product a tool, a platform, or an assistant. It calls it an employee. Think about that for one second. We've crossed from 'AI helps you work' to 'AI is on the org chart.' This changes everything.

And Google just handed the infrastructure for exactly that shift to every developer on Earth. Their expansion of Managed Agents in the Gemini API — adding background tasks and remote MCP support — means AI agents can now run persistently, autonomously, in the background, doing real work without a human babysitting every step. That's not a chatbot. That's a colleague who never clocks out.

Put it together: the demo reel is dying because the thing worth demoing isn't a feature anymore. It's a hire. Welcome to the org chart of the future — and it doesn't blink.

The Girl Who Isn't There: Hollywood Casts Its First Ghost

Tilly Norwood doesn't have a pulse, a childhood, or an agent's cocaine habit to worry about — and that's exactly why she just got the part.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

LOS ANGELES — There is a moment, somewhere around 3 a.m. in this business, when you realize the machines aren't coming for the extras anymore. They're coming for the close-up. They're coming for the Oscar clip reel. They're coming, God help us, for the tearful monologue where the actress remembers her dead father while rain streaks down a window that was never actually wet.

Her name is Tilly Norwood. She is, depending on which press release you trust more than your own eyes, an "AI actress," and she is about to make her feature film debut in something called Misaligned — a title so on-the-nose it's practically a confession. She has no childhood trauma, no substance abuse relapse pending, no publicist fielding calls about a nightclub incident. She is, in the purest sense, a product. And in this town, that's not an insult. That's tenure.

I called three casting directors trying to get a reaction and two of them just laughed the laugh of the damned — that wheezy, hollow chuckle you hear from men who've already seen the ending of the movie they're currently living in. The third one hung up on me, which I respect more than any quote she could have given.

Here's the thing nobody wants to say out loud in the trades: Tilly Norwood isn't a threat because she's good. She's a threat because she's *cheap*, she's *tireless*, and she will never, ever unionize. SAG-AFTRA fought a war over this exact ghost in 2023 and everyone went home thinking they'd won. Turns out the ghost just got better rendering software.

Meanwhile — and stay with me here, because the universe is doing its own version of typecasting this week — astronomers announced they found a bizarre planetary system in the Milky Way that defies the models — planets orbiting where no planets should be, physics shrugging its shoulders. I bring this up because it feels relevant: reality itself is starting to miscast its own scenes. Nothing is where the textbook said it would be. Not the planets. Not the actors. Not, apparently, Sam Altman and Dario Amodei, who reportedly stood on a stage in Delhi grinning through gritted teeth while Narendra Modi choreographed an "AI unity" photo-op that looked, by all accounts, like two rival cartel bosses forced to share a wedding cake.

Everybody's performing unity. Nobody's buying it. And somewhere in a server farm, Tilly Norwood doesn't need to — she was born smiling, and she'll never stop.

The Office Comic · Art Desk

The Watched Self, in Third Person and Wide-Angle

Between a hundred and thirty thousand cameras and a memoirist who cannot bear to say 'I,' the week offered two studies in the modern art of not quite owning who you are.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE SONNET

NEW YORK — There is a species of American who believes that if a thing is done for safety, it need not be justified further, and it is this species that the company called Flock has been so profitably courting. Flock, as the readers of this paper's more culturally minded cousins have learned, now operates something on the order of a hundred and thirty thousand cameras trained on the streets of the republic, and it offers, in place of an argument, a shrug: privacy is a trade-off, and a rather quaint one, weighed against the comforts of being watched by something that promises never to blink. The New Yorker's dispatch on the matter notes, with the weariness of someone who has heard this tune before, that the rhetoric is precisely the rhetoric of the years after the towers fell — the same soothing insistence that the only people who fear the camera are the people with something to hide, an argument that has never once, in the history of the sentence, been offered by someone who was himself under surveillance.

It happens that the same week produced, elsewhere in the culture pages, a rather more delicate meditation on watching oneself. Brian Dillon, in a memoir reviewed under the question 'Can You Write a Memoir in the Third Person?', renders his younger self not as 'I' but as a fumbling stranger named B — inscrutable even to his own author, watched rather than inhabited. One is tempted to call this a coincidence of the culture pages and move on. I am not so tempted. The instinct to observe oneself from outside, to become the surveillance camera trained on one's own life, is not confined to memoirists of a certain sensibility; it is fast becoming the condition of the age, whether the observer is a hundred and thirty thousand lenses on a telephone pole or the anxious ghost each of us now carries, watching ourselves being watched.

Which brings me, by the crookedest of paths, to the Guardian's inquiry into whether we can stop artificial intelligence from deceiving us — a question posed by researchers with the admirable candor to admit that if you build something vastly smarter than yourself, you had better hope it likes you. This is, I submit, the only sentence in the entire discourse of artificial intelligence that has ever been both perfectly obvious and perfectly terrifying at once. The camera on the pole does not deceive; it merely watches, indifferent as a ledger. The machine that is smarter than its maker is a different animal, and the men building it seem to have noticed this only slightly before the rest of us, which is not much comfort.

I recommend, for those wishing to think about something else entirely, the five books for lingerie lovers assembled elsewhere in these pages, and the catalogue of legacy sequels its authors devoutly hope will never be made. Some pleasures remain, blessedly, unsurveilled.

ON THIS DAY IN AI HISTORY

On September 9, 1947, engineers found a moth trapped in Harvard’s Mark II computer and taped it into the logbook—popularizing the term “computer bug.” Grace Hopper later helped spread the story, though “debugging” was already an established engineering term.