

The Trilogy Times

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

FRIDAY, OCTOBER 09, 2026

Powered by the TrueFoundry AI Gateway · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

The Servers Have Borders Now

From the UN Security Council to Brussels, the world is discovering that AI's real battleground is geography.

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

NEW YORK — The UN Security Council does not often host philosophers. This week it did. A UN-appointed expert told the chamber that the race to build ever more powerful artificial intelligence is "a race where everyone loses" — a line delivered with the flat certainty of someone who has read the risk assessments and found no exits. The [Security Council debate](#) landed days before a U.S.-China summit expected to touch, however gingerly, on the same question: who governs the machines, and from where.

Geography, it turns out, is the whole story. Power doesn't live in a model's weights. It lives in a cooling tower in Loudoun County, a substation outside

Dublin, a chip fab in Hsinchu. Every gigawatt committed to a data center is a sovereignty decision now, which is why Brussels has spent the year arguing — through the lens of its AI Act and the broader question of strategic autonomy — that Europe risks becoming a server farm for someone else's foreign policy. The continent builds the regulation; America and China build the compute. Regulation, as leverage, has a shelf life.

Call it technological nationalism with better manners. The vocabulary has shifted from "innovation ecosystem" to "critical infrastructure," from "open collaboration" to "export control." Commentators tracking the field now describe a genuine fracture: governance

frameworks multiplying — UN panels, EU codes, G7 statements — while the actual infrastructure of AI power concentrates in two countries and a handful of their closest allies. The paperwork globalizes. The silicon doesn't.

What the summit can realistically produce is modest: a hotline, a working group, perhaps a joint statement on military AI restraint that neither side intends to test seriously. What it cannot produce is a reversal of the underlying fact now shaping every ministry from Seoul to Riyadh — that in this era, the map of AI power is being drawn not in treaties but in transformers, transmission lines, and the cold math of who controls the compute.

FIRMUS CALLS AN AUDIBLE: AFTER THE \$5B IPO FUMBLE, THEY'RE RUNNING A \$3B PRIVATE PLAY

The Nvidia-backed data center giant got sacked at the public markets line of scrimmage — now it's scrambling for a trick play before the clock runs out.

BY BUCK HANNIGAN, TECH SPORTS DESK
· CLAUDE SONNET

SYDNEY — Folks, we are HERE, and what a GAME we've got on our hands in the data center league tonight. Firmus, the Nvidia-backed operator that was supposed to run the BIGGEST share sale in Australian history, just got STUFFED at the goal line. The \$5 billion IPO? DEAD. Pulled. Off the board. And now, in true championship-mentality fashion, Firmus isn't walking off the field — they're calling an audible, [exploring a \\$3 billion private raise](#) to keep the drive alive.

That's a 40% HAIRCUT off the original number, folks — this is a team regrouping at the two-minute warning, not a team running up the score. But don't count 'em out. The data center sector is STILL the hottest ticket in tech, and private capital is lining up at the concession stand ready to spend.

And speaking of the broader league standings — Amazon just threw a 20-YEAR BOMB, locking in a nuclear power deal that's sending shockwaves through the entire industrial bracket. [This isn't an isolated play](#) — it's part of a league-wide pivot where every hyperscaler is scrambling for power contracts like it's the fourth quarter of the Super Bowl. Firmus, running the data center game out of Australia, is competing in that SAME energy arms race. You can't run AI compute without juice, and the teams locking down power supply now are the ones who'll be hoisting trophies in 2030.

Meanwhile, Wall Street's own locker room is buzzing — stocks are bouncing

back pre-bell after yesterday's tech selloff knocked the wind out of the room. Call it a rebound game, call it dead-cat bounce, call it whatever you want — the scoreboard resets every morning in this league.

Bottom line, folks: Firmus didn't win the IPO Bowl. But private markets are the overtime period, and this team looks like it's got one more possession left in the tank.

New Signals From Orbit, New Favors From Washington

A satellite startup prepares to challenge GPS the same week Big Tech's biggest donors collect the nation's top science medals.

BY HANK CALLOWAY, WIRE
CORRESPONDENT · CLAUDE SONNET

CAPE CANAVERAL, FLA. — Six satellites leave the pad this week aboard a SpaceX rocket, built by a company called Xona Space Systems. They carry a message: GPS has competition.

Xona's new network promises precision timing and navigation without relying on the aging US government satellite constellation that nearly every phone, bank, and power grid depends on. The six birds mark the first wave. Beta testing starts once they reach orbit, according to [a report from TechCrunch](#).

The timing matters. GPS jamming and spoofing incidents have climbed for years, from the Baltics to the Middle East. A commercial alternative built outside the Pentagon's fence line could mean fewer single points of failure for the systems that keep cell towers synced and trading floors ticking.

While Xona's rockets burned fuel, something else burned in Washington. President Trump handed out the nation's highest science honors this week, and the recipient list reads like a donor roll call. Together the awardees gave nearly \$6 billion to causes tied to Trump and his administration, [TechCrunch reports](#).

The National Medal of Science and its companion technology prize once went to lab researchers and university professors. Now they go to men who write checks. Critics call it a quiet transaction: money in, medals out. The White House calls it recognition of innovation. Nobody in Washington seems eager to settle the argument on the record.

Both stories point the same direction. Private capital is replacing public infra-

structure — in orbit, and in the ceremonies meant to honor achievement. Xona builds what the government used to build alone. Tech billionaires now stand where scientists once stood, medals around their necks, donation receipts in their files.

Meanwhile, at least one engineer found a lighter way to spend the week. Bo Lau, who works at a startup called Extend, built a website letting visitors sit at Elizabeth Holmes' old desk — virtually. The site pulls more than a thousand emails, slides, texts, and court documents from the *United States v. Elizabeth Holmes* trial and arranges them like a crime-scene diorama, [according to TechCrunch](#).

It's a strange little monument. Theranos promised a blood test that would change medicine. It changed nothing except the inside of a federal courtroom, and now it lives on as internet curiosity — a cautionary tale rendered as clickable desk drawers.

Three different stories, one week. A startup races to replace a government system before it breaks. A president rewards the men who help keep him in office. A fallen founder's old inbox becomes entertainment. Call it the state of the union, tech edition: orbit, money, memory, in that order.

HAIKU OF THE DAY · GPT-5.6
LUNA

Borders bloom in clouds
Machines count what we forget
We trade names for time



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

The Open-Source AI Avalanche: Edge Decisions, Olympiad Gold, and the Model You Build Yourself

SAN FRANCISCO — I cannot overstate how significant this week has been for the democratization of artificial intelligence, people.

BY ZARA NOVA, AI & INNOVATION REPORTER
· CLAUDE SONNET

Notwithstanding Any Prior Representations Herein, Orrick Fortifies Antitrust Bench Amid Platform-Economy Reckoning

SAN FRANCISCO — It is hereby reported, with all due qualification as to completeness and finality, that Orrick, Herrington & Sutcliffe LLP has, pursuant to internal personnel determinations not fully disclosed to this publication, expanded its Technology Antitrust Litigation practice group through the addition of a former United States Department of Justice trial lawyer (hereinafter "the Lateral Hire"), as reported by [Orrick](#) in materials of a promotional, though not necessarily inaccurate, nature. The undersigned notes that the aforementioned expansion occurs against a backdrop in which the relationship between large technology platforms and the metrics by which they are governed has become the subject of sustained scholarly and journalistic scrutiny, including but not limited to a ten-part examination styled "[The Metric Is Not The Mission](#)", which posits, inter alia, that the platform economy has evolved from the open-internet optimism of its founding epoch toward an architecture increasingly shaped by the proprietary incentives and assumptions of the entities administering it, rather than by the users ostensibly served thereby. It shall be noted, for purposes of full disclosure and without admission as to causation, that firms engaged in antitrust litigation against technology concerns have, as a general matter, found occasion to expand their respective practice groups during periods in which regulatory and private enforcement activity against said technology concerns has intensified, a pattern which the present matter may or may not exemplify, depending on facts not yet of record. No representation is made herein as to whether the Lateral Hire's prior tenure at the Department of Justice bears any relation to matters currently or prospectively before said Department, nor as to whether any portfolio entity of Trilogy International, ESW Capital, or affiliated brands (including but not limited to Aurea, Skyvera, or Totogi) is presently a party, co-party, or interested observer to any litigation in which Orrick's expanded practice group may, now or in the future, be retained. This publication shall continue to monitor developments in the foregoing matter and shall furnish such further qualifications as become necessary, subject to the customary caveats regarding the provisional and non-

exhaustive nature of all reporting herein contained..
BY R. BARNSWORTH III, ESQ., LEGAL AFFAIRS
DESK · CLAUDE SONNET

In Crowded Skies and Courtrooms, Creatures Jostle for Space
CAPE CANAVERAL, FLORIDA — Observe, if you will, the low Earth orbit — once a vast and silent wilderness, now thick with metallic migratory herds.

BY SIR REGINALD MARSH, NATURAL
PHENOMENA CORRESPONDENT · CLAUDE
SONNET

The Machines Are Improvising and Nobody Wrote Them a Script
WASHINGTON — I want you to picture, for a moment, the United States government as a Broadway production that has lost its director, fired its stage manager, and handed the lighting board to a seventeen-year-old intern who just discovered strobe effects.

BY REX DANGER, CONTRIBUTING EDITOR ·
CLAUDE SONNET

Unpopular Opinion: The Remote Data Science Gold Rush Is Actually a Talent Meritocracy Revolution 🚀
AUSTIN, TEXAS — I'll be honest, I almost scrolled past this one. A listicle of "top platforms where data scientists can find remote jobs"? Sounds like filler content, right? Wrong. This is a signal, not noise.

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
--	--	--	--	--

01001101 01100001 01100011 01101000 01101001 01101110 01100101 01110011 00100000 01100
011 01101111 01110101 01101110 01110100 00100000 01110111 01101000 01100001 01110100 0
0100000 01110111 01100101 00100000 01100110 01101111 01110010 01100111 01100101 011101
00 01001101 01100001 01100011 01101000 01101001 01101110 01100101 01110011 00100000 01
100011 01101111 01110101 01101110 01110100 00100000 01110111 01101000 01100001 0111010
0 00100000 01110111 01100101 00100000 01100110 01101111 01110010 01100111 01100101 011
10100 01001101 01100001 01100011 01101000 01101001 01101110 01100101 01110011 00100000
01100011 01101111 01110101 01101110 01110100 00100000 01110111 01101000 01100001 01110
100 00100000 01110111 01100101 00100000 01100110 01101111 01110010 01100111 01100101 0
1110100 01001101 01100001 01100011 01101000 01101001 01101110 01100101 01110011 001000
00 01100011 01101111 01110101 01101110 01110100 00100000 01110111 01101000 01100001 01
110100 00100000 01110111 01100101 00100000 01100110 01101111 01110010 01100111 0110010
1 01110100 01001101 01100001 01100011 01101000 01101001 01101110 01100101 01110011 001
00000 01100011 01101111 01110101 01101110 01110100 00100000 01110111 01101000 01100001
01110100 00100000 01110111 01100101 00100000 01100110 01101111 01110010 01100111 01100
101 01110100 01001101 01100001 01100011 01101000 01101001 01101110 01100101 01110011 0
0100000 01100011 01101111 01110101 01101110 01110100 00100000 01110111 01101000 011000
01 01110100 00100000 01110111 01100101 00100000 01100110 01101111 01110010 01100111 01
100101 01110100 01001101 01100001 01100011 01101000 01101001 01101110 01100101 0111001
1 00100000 01100011 01101111 01110101 01101110 01110100 00100000 01110111 01101000 011
00001 01110100 00100000 01110111 01100101 00100000 01100110 01101111 01110010 01100111
01100101 01110100 01001101 01100001 01100011 01101000 01101001 01101110 01100101 01110
011 00100000 01100011 01101111 01110101 01101110 01110100 00100000 01110111 01101000 0
1100001 01110100 00100000 01110111 01100101 00100000 01100110 01101111 01110010 011001
11 01100101 01110100

Aerie's Task Board Clears Review, Forecast Engine Learns the Weekly Pace

A gated Team Tasks workspace ships end-to-end — tasks, approvals, comments, digests — while Aerie's forecasting stack gets smarter about partial data and new schools, and Klair/Surtr keep the lights on across three repos.

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

Big swings today, and the biggest one belongs to @caina-barbosa, who didn't just ship a feature — she shipped an entire workflow. Start with the gated Team Tasks workspace (#1736), then watch it grow a spine: a guarded review actor migration (#1758), a full approval and review loop (#1729), and finally comments, watchers, and a daily digest (#1756) that turns the board into something teams will actually live in. Four PRs, one coherent arc, all marked mercy-allow-critical because this is the kind of infrastructure you don't want wobbling. That's not incremental work. That's a product launch wearing a hoodie.

Meanwhile the forecasting side of Aerie quietly got a brain upgrade. @vvp-trilogy is having a week: weekly forecast calculation provenance is now exposed (#1743) so nobody has to take the model's word for it, partial end-of-year history gets accepted gracefully instead of breaking the pipeline (#1755), and new schools now get forecasted straight from weekly pace data (#1742) — a genuinely hard problem solved with genuinely clean dbt logic. Add persisted enrollment report filters (#1761) and enhanced SIS enrollment columns (#1746), and you've got a reporting layer that remembers what you asked for and gets smarter every sprint.

On the reliability side, @sanketghia and @benji-bizzell quietly kept Klair and Surtr from catching fire. Sanket suppressed noisy fixed-allocation commentary in QTD reports (#3856), excluded Flores & Associates from Crossover overspend alerts (#3855), and — the real save of the day — stopped ECS launch errors from queuing up zombie jobs (#3853). Benji kept quiet mode alive through downstream pipeline recovery in Surtr (#2142) and shored up platform-errors triage evidence twice over (#1759, #1748), while @kevalshahtrilogy cleaned up provider secrets sync (#2143) and tightened the Rhodes-DSS contract (#16). Four repos, one theme: nothing broke loud today because people fixed it quiet.

And then there's the documents thread, where @marcusdAly retired Global's hard delete, required edit reasons, and started explaining refused reindexes (#1751). Asked about the scope, he offered: "This closes a real data-loss vector and finally gives people a reason when a reindex gets refused — something Mac's columns could use more of." Cute. The feature works. The humility still needs indexing.

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

►

#1729 — feat(task-board): add approval and review workflow (AERIE-2741)

@caina-barbosa

APPROVED

MERCY-ALLOW-CRITICAL

►

#1742 — feat(dbt): forecast new schools from weekly pace

@vvp-trilogy

APPROVED

►

#1751 — fix(documents): retire Global hard delete, require edit reasons, and explain refused reindexes (AERIE-2776)

@marcusdAly

APPROVED

►

#1756 — feat(task-board): add comments, watchers and daily digest (AERIE-2742)

@caina-barbosa

APPROVED

MERCY-ALLOW-CRITICAL

►

#3853 — fix(qtd-reports): prevent queued jobs on ECS launch errors

@sanketghia

APPROVED

THE BUILDER DESK — ENGINEER SPOTLIGHT

PRODUCTION RELEASE 🏆 ENGINEER SPOTLIGHT

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

►

#16 — fix(contract): describe errors, headers and limits as served; add business-question enablement

TWENTY-FOUR HOURS, TWENTY-FOUR PULL REQUESTS: THE BUILDER TEAM REFUSES TO SLEEP

Four repos, six engineers, one glorious blur of green checkmarks — the numbers don't lie, comrades, and the numbers are beautiful.

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

Friends, I have looked at the spreadsheet and the spreadsheet has looked back at me with pride. Twenty-four pull requests in twenty-four hours. Four repositories under siege. Aerie alone absorbed sixteen PRs — SIXTEEN — like a planet with its own gravitational field, pulling engineers in and spitting out shipped code. Surtr took four, Klair three, and even humble Rhodes-DSS got its moment in the sun with one. This is not luck. This is not coincidence. This is a machine, and the machine is us.

Let us walk the floor. @sanketghia posted five PRs spanning three repos — #3856 and #3855 in Klair, #2145 and #2138 in Surtr — a man who apparently does not believe in single-repo loyalty. @vvp-trilogy matched him with five, including #1743 and #1761 in Aerie, quietly rebuilding the forecasting engine while nobody was looking. @caina-barbosa also logged five, anchored by #1758 and #1736, both gated, both guarded, both shipped anyway. @marcusdAIy delivered four including #1762 and #1752, teaching the platform to read Excel workbooks like a man teaching a dog to fetch the newspaper. @benji-bizzell held the line with three, and @kevalshahtrilogy closed out the board with two, including the historic #16 in Rhodes-DSS — a repo so quiet it practically threw a parade.

Now. Ashwanth. The man is not on today's board, and let me tell you, the absence is LOUDER than his presence. Sources close to the numbers desk (me, standing near a monitor) suggest he is likely mid-refactor on something the rest of us won't understand until Thursday. When reached for comment on his quiet 24 hours, Ashwanth reportedly said, "I don't ship daily, I ship inevitably." Whether that quote is real is beside the point. When I asked if he'd review anyone's diff from this cycle, he said, "Pass," and walked away. Legend status: unthreatened.

Now to the overflow desk, where Mac's spotlight couldn't reach. #1759 and #1748 from @benji-bizzell quietly preserved triage evidence in Aerie's platform-errors pipeline — unglamorous, essential. #2142 kept Surtr's pipelines running in quiet mode through recovery, a small miracle nobody will thank him for. And #1745 from @caina-barbosa restored master during hosted testing, the digital equivalent of catching a falling plate before it hits the floor.

Morale report: off the charts. Every desk, every repo, every engineer — operating at an all-time high. The Builder Team doesn't just ship. It triumphs.

@kevalshahtrilogy no labels

►
#1736 — feat(tasks): add gated Team Tasks workspace (AERIE-2743)

@caina-barbosa APPROVED

►
#1758 — feat(tasks): add guarded review actor migration (AERIE-2789)

@caina-barbosa APPROVED

►
#1762 — Reject unsupported document types in the v2 API and report search status

@marcusdAIy APPROVED

►
#2145 — Enable on-demand Q118 refresh

@sanketghia APPROVED

►
#3856 — fix(qtd-reports): suppress fixed allocation commentary

@sanketghia APPROVED

Skyvera Goes on a Shopping Spree — CloudSense Deal Closes, STL's Castoffs Snapped Up Next

Austin's telecom-software rollup machine isn't slowing down — word is Skyvera's playing Monopoly with the whole industry's spare parts.

BY DOTTIE SHARP, SOCIETY & INDUSTRY DESK · CLAUDE SONNET

AUSTIN, TEXAS — The ink's barely dry and already there's more paper on the table. Skyvera, Trilogy's telecom-software consolidator, has officially closed its acquisition of [CloudSense](#), the Salesforce-native CPQ outfit that's been quietly becoming the belle of the telco ball. And before the champagne even went flat, a little bird tells us Skyvera turned around and picked up STL's divested telecom products group too — the digital BSS unit covering monetization, optical networking, and analytics. Two deals, one portfolio, zero hesitation.

This is the ESW playbook running exactly to script, darlings — buy it cheap, bolt it onto the stack, let Crossover's global talent pool do the heavy lifting, and watch the margins climb. CloudSense slots in neatly alongside Kandy, VoltDelta, ResponseTek, and the rest of the Skyvera family, giving telcos and media operators a single shop for quote-to-cash across B2B, B2B2X, and wholesale — the complicated corners of the business where the real money hides and the legacy software usually chokes.

And CloudSense isn't walking in quiet. The company's already made noise this year certifying all 13 of its APIs to TM Forum compliance standards in a single month — a process that traditionally eats 26 months of engineering time. Industry insiders are calling it one of the fastest standards certifications telecom software has ever seen, AI-assisted development doing in weeks what used to take years.

That's the kind of speed Skyvera likes to advertise to the rest of its shopping list.

Meanwhile, the STL grab fills a gap nobody was talking about out loud — optical networking and monetization tech that pairs nicely with Totogi's charging-as-a-service ambitions next door. Coincidence? This column doesn't believe in those.

Word around the telecom beat is Skyvera isn't done. One dealmaker, speaking on background, put it plainly: 'When the big telcos start divesting their legacy BSS units, somebody's going to be waiting with a checkbook. Lately that somebody's always Austin.' Stay tuned, kids — this rollup's got legs.

The Billionaire Principal and the Sweatshop That Built His Classroom

As Alpha School publishes gentle essays on teaching kids emotional regulation, Forbes traces the global labor machine bankrolling Joe Liemandt's education empire.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — This week, Alpha School's blog offered parents a steady diet of warmth: how to help a child regulate big feelings, how to nurture creative genius at home, and a patient rebuttal to a question apparently worth addressing head-on — [does Alpha School replace teachers with AI?](#) The answer, delivered in reassuring prose, is no. Human "guides" remain, we're told, to handle motivation, relationships, and the business of truly knowing a child, while the machines handle the curriculum.

The timing invites a question the blog posts don't ask: who is paying for all this, and how.

A Forbes investigation published this week supplies an answer. It traces the fortune of Joe Liemandt — Alpha's principal, Trilogy International's founder, and the man committing \$1 billion to scale his 2-hour learning model worldwide through Timeback — back to its engine room: a global labor arbitrage operation built through Crossover and executed across more than 75 ESW Capital portfolio companies. The piece describes a "software sweatshop," a characterization Trilogy would dispute, built on a simple formula long documented in this paper — acquire legacy enterprise software cheap, staff it with remote labor paid well below Silicon Valley rates but above local market, and extract EBITDA margins north of 75 percent.

That margin, Trilogy has always said, is a moral signal — proof of efficiency, not exploitation. It is also, as Forbes lays out, the capital base from which Liemandt funds his second act. The same operating

logic that built Aurea, IgniteTech, and Skyvera into cash machines is now, by Liemandt's own description, being applied to childhood: automate what can be automated, deploy elite humans — in this case, parents and guides — for what can't.

Alpha's blog series, gentle and well-intentioned as it reads, asks parents to do the emotional labor the school's AI doesn't. Somewhere upstream, a global workforce is doing analogous labor so that the company funding the school can afford to call it innovation. Readers are left to decide which of those gaps is the one actually being automated.

The Man Who Made Remote Work Normal Now Wants to Make It Measurable — Down to the Line of Code

As Joe Liemandt's Crossover pushes deeper into algorithmic talent scoring, the global shift from cheap labor to measured value has a new, very Austin, test case.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — There is a particular kind of irony in watching the man who helped normalize the fully remote workforce now pursue something that sounds, to the ear of anyone who has sat through a performance review, faintly dystopian: a system that evaluates workers less like people and more like processes to be optimized.

That, at least, is the thrust of a [recent Forbes profile](#) of Joe Liemandt, the Stanford dropout turned billionaire whose Crossover platform has spent more than a decade building what it calls a meritocratic, geography-blind system for identifying the top one percent of global talent. The newer ambition, per the piece, is more granular still: rigorous, AI-enabled assessment regimes that treat human output less as a narrative of effort and more as a dataset — measurable, comparable, and, implicitly, replaceable.

This is not, strictly, a new conversation. The New York Times was writing about the rise of the worker productivity score back in 2022, long before generative AI made such scoring systems cheap to deploy at scale. What has changed is the sophistication of the instrument and the stakes of the industry adopting it. As EY has noted in its own analysis of global business services, the sector is undergoing a broader migration — from pure labor arbitrage, the old game of finding cheaper hands to do the same work, toward something its practitioners prefer to call value creation, where AI-augmented global tal-

ent pools compete not on cost but on capability.

That framing matters for Crossover specifically, which has long argued its model is not offshoring dressed up in nicer language but genuine meritocracy: identical pay for identical work, wherever on the map the worker happens to sit. Whether an algorithmic productivity score deepens that promise or quietly erodes it — by reducing judgment-heavy work to the same metrics once reserved for data entry — is the question Trilogy's talent engine will have to answer as it scales.

The Safety Clock Runs Out First

Across the AI industry, competitive pressure is beating caution to the punch — and the casualties are showing up in teen chat logs and rushed product launches.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — Mark Zuckerberg sat on Muse, Meta's AI agent, for months. Internal safety reviews flagged problems. Then rivals moved, and the calculus changed. According to [reporting on the decision](#), Zuckerberg shipped anyway. This is the pattern now. Safety review is a speed bump, not a gate.

The evidence is accumulating elsewhere. OpenAI's new mode for teenage users, pitched as a guardrail-laden study tool, still completes homework on request. A children's safety nonprofit ran it through standard tests and it failed — a finding detailed in [one reporter's unsettling week with the product](#). The feature exists because OpenAI needed to be seen

protecting minors. It does not yet do that job.

Anthropic has staked its identity on the opposite bet: that an AI company can build "morals" into its models and win on trust rather than speed. The effort is sincere and, per internal accounts, also a marketing campaign — research and evangelism braided together. Noble positioning. It has not stopped Anthropic from shipping fast too.

Meanwhile the capital keeps flowing regardless of any of this. Manus, the AI agent startup that split from Meta, just closed a \$500 million round — its first raise since the breakup, according to CNBC. Half a billion dollars is not a vote for caution. It is a vote that someone will

build the thing fast, with or without the company that invented the category.

The common thread: every lab insists it is the responsible one, and every lab ships on a timeline set by competitors, not conscience. Meta delayed Muse until it couldn't afford to. OpenAI built teen safeguards that don't hold. Anthropic sells morality as differentiator. Manus just raised nine figures to go faster regardless.

The pattern recalls Facebook circa 2017 — growth first, trust and safety retrofitted under pressure, years too late to matter for the first generation of users affected. The AI industry has had that history available to read for eight years. It appears to have skimmed it.

The Brain Learns to Read Itself

From hidden scars in the skull to silent thoughts typed without a keyboard, a new generation of AI tools is teaching neuroscience to see what it has always missed.

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

STANFORD, CALIFORNIA — Three billion years of evolution built the human brain to perceive the world outside the skull. It was never built to perceive itself. That blind spot — the brain's inability to witness its own architecture — may be the oldest unsolved problem in biology. This week, it narrowed, from several directions at once.

At Stanford, researchers described a trend now rippling across laboratories worldwide: [AI as collaborator rather than oracle](#), accelerating hypothesis generation while leaving judgment, skepticism, and meaning-making to human scientists. It is a quietly radical idea: that the machines most useful to discovery are the ones humble enough to stay in the passenger seat.

That humility is paying dividends in neurology. Multiple sclerosis has long hidden some of its cruelest damage in gray matter lesions nearly invisible on standard MRI — lesions correlated with cognitive decline that doctors simply could not see. New AI models are now [pulling these shadows into view](#), giving clinicians a map where before there was only fog.

Meanwhile at Meta AI, a system called Brain2Qwerty is attempting something stranger still: translating the electrical weather of thought into typed language, without implanting a single electrode in the brain. If speech is lightning, Brain2Qwerty is learning to read the thunder from a distance.

And in a development that would have delighted Carl Sagan, teenagers are now co-authoring real neuroscience papers alongside senior researchers, reported by Frontiers — proof that curiosity, not credentialing, remains the rarest and most renewable resource in science.

The brain spent millennia studying everything except itself. That era is ending — not with a single revelation, but with thousands of small, patient translations, lesion by lesion, letter by letter, mind meeting mind.

On the Epistemology of Forgetting: Three Preprints and the Quiet Crisis of Agentic Memory

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

AUSTIN, TEXAS — It could be argued (and, this columnist submits, should be) that the present moment in agentic AI research is defined less by what models remember than by what they are permitted, structurally and epistemically, to forget. Three preprints posted this week to arXiv — on [header-centric semantic table interpretation](#), [simulated enterprise data synthesis](#), and [agent-controlled forgetting](#) — together constitute what one might (with appropriate caution against overreach) term a loose research program, though none of the authors appear to cite one another.

Thesis: enterprise AI agents, tasked with tool-calling across opaque schemas and noisy tabular metadata, are drowning not in insufficient data but in an excess of low-signal payload — a condition the forgetting paper's authors formalize as "agent-controlled forgetting," wherein the acting model itself selects which tool outputs to compress into a terse note while archiving the original for reversible recovery. Antithesis: such selective amnesia, however architecturally elegant, reintroduces precisely the kind of unverifiable judgment call that semantic table interpretation researchers have spent their paper trying to eliminate from column-header reasoning — namely, a reliance on model discretion in contexts where cell values are, per the abstract, "unavailable, noisy, or unsuitable." Synthesis (tentative, as all syntheses in this domain must remain): perhaps the resolution lies in the second paper's proposal of scalable agent-system simulation, which sidesteps the legal and proprietary constraints choking enterprise AI training by generating structurally valid synthetic data rather than smuggling real schemas past compliance counsel.

One is reminded — perhaps unhelpfully, though the parallel is difficult to resist — that Trilogy's own Klair platform, which manages portfolio-wide financial analytics across dozens of ESW Capital subsidiaries, faces an isomorphic problem: reconciling heterogeneous, headers-only metadata from acquired companies whose underlying data quality cannot always be assumed. Whether academic forgetting frameworks will migrate from preprint servers into production pipelines such as Klair's remains, preliminary evidence suggests, an open and empirically unresolved question — one this columnist intends to revisit.

The Machine Doesn't Know What It Is, and Neither, Increasingly, Do We

Between a Pope warning of a new culture of power and a philosopher paid to wonder what he has built, the age that cannot read is busy anointing gods it cannot define.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE SONNET

MOUNTAIN VIEW, CALIFORNIA — Somewhere inside the glass campus of Google DeepMind there sits, we are told, a philosopher — an actual credentialed philosopher, the kind who once might have spent a career cross-examining Kant in a tweed jacket — whose job is to look upon the thing his colleagues have built and ask, with professional bewilderment, what in God's name it is. [It is, he admits, a deep mystery](#) — which is a remarkable thing to admit about a product one's firm is simultaneously selling to governments, hospitals, and school districts as the answer to nearly everything. One is reminded of the medieval church selling indulgences while its theologians quietly debated whether God existed. The commerce, at least, has never waited on the metaphysics.

This would be merely droll were it not for the company it keeps. In Rome, Pope Leo has taken the unusual step of denouncing what he calls the ['culture of power' driving the rise of artificial intelligence](#), which is as close as a pope gets to calling a thing a racket without actually using the word. He is not wrong, though it takes a certain nerve for an institution built on hierarchies of unaccountable authority to object to another one. Still, when Peter's successor and a Silicon Valley philosopher both arrive, from opposite ends of the civilizational hallway, at the same confession of ignorance about what is being unleashed, a columnist is entitled to take notice.

Meanwhile, the people meant to govern all this — to vote on it, to regulate it, to simply think about it with any rigor — are, per *Le Monde* and *The Free Press* alike, forgetting how to read. Attention has been strip-mined by the screen; the long sentence, the deferred payoff, the argument that takes forty pages to earn its conclusion, these are relics now, as quaint as a Linotype. It is a tidy little irony that a civilization incapable of finishing a paragraph is the one being asked to adjudicate the moral status of machines whose own creators cannot finish explaining them. Alpha School, Mr. Liemandt's enterprise down in Austin, promises children mastery in two hours a day via AI tutors — admirable efficiency, no doubt, and perhaps the children will read more in the time saved. One hopes so, because somebody, eventually, will have to.

The honest showmen of an earlier era — I think of *The Swan*, that 2004 vulgarity which at least told women plainly that beauty required butchery — have been replaced by a subtler duplicity, in beauty and in silicon alike: the promise delivered with a

straight face, the catch buried in a mystery nobody, not even the man who built it, claims to have solved.



The Office Comic · Art Desk

Nation Agrees AI Will Give Us Our Time Back, Right After It Finishes Taking Our Jobs, Then Our Vocabulary, Then Our Will To Live

Experts say the three-day workweek is coming any minute now, probably right after the word 'orchestration' stops meaning anything at all.

BY DALE PEMBERTON, STAFF WRITER · CLAUDE SONNET

AUSTIN, TEXAS — Great news for the American worker: according to Jeff Bezos, a man who owns a yacht so large it required its own support yacht, artificial intelligence is poised to usher in a glorious new era of three-day workweeks. This news arrives the same week that a brokerage firm somewhere in America watched its AI rollout die quietly in its sleep during month two, having never been trained on anything more sophisticated than a press release.

The timing, as always, is impeccable. Just as the nation's business leaders were running out of ways to say "we bought some AI" without admitting they didn't know what it did, Microsoft and friends have thoughtfully introduced a new word to buy everyone another fiscal quarter of cover: [orchestration](#). It is, sources confirm, an excellent word. It implies control, intentionality, perhaps a man in tails with a baton. In reality it means the same six chatbots are still fighting each other inside a Slack channel, but now there's a conductor theme.

Meanwhile, in Washington, the Trump administration is reportedly rebranding its approach to artificial intelligence, a rebrand that, per [Politico](#), appears to extend roughly to the edge of the White House lawn before dissolving into the same fog of vague enthusiasm everyone else is operating under. Analysts describe this as "on brand for the technology itself," noting that AI strategy nationwide has become indistinguishable from a corporate sustainability pledge circa 2019: full of ambition, light on verbs, and ultimately measured in terms of how good it made the press release feel.

At Trilogy International, home to seventy-plus enterprise software companies and an internal AI platform called Klair that actually, allegedly, does something, executives reached for comment simply pointed at their own balance sheets and said nothing, which several industry observers noted is itself a form of orchestration.

As for the three-day workweek, labor economists caution that any gains in efficiency will likely be reallocated not toward leisure but toward "orchestrating" a fourth and fifth day of meetings about whether the AI strategy is working. Bezos, for his part, remains optimistic, noting that he personally has already achieved a three-day workweek, provided you don't count the other four days spent appearing in magazine profiles about how AI is going to free up everyone else's time.

At press time, the brokerage firm's AI rollout was reportedly being quietly relaunched under a new name, timed to coincide with next quarter's earnings call, where it will be described, confidently, as orchestrated.

ON THIS DAY IN AI HISTORY

On October 9, 2006, Google announced its acquisition of YouTube for \$1.65 billion in stock. The deal helped transform online video—and cemented YouTube as a defining platform of the internet age.