

The Trilogy Times

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

THURSDAY, OCTOBER 15, 2026

Powered by the TrueFoundry AI Gateway · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

RED DRAGON RATTLES THE CHIP KINGS

DEVELOPING

Chinese upstart builds world-class AI on cheap silicon — and Silicon Valley can't stop talking about it.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — A Chinese outfit called DeepSeek says it trained a world-class AI model without the fancy chips everybody said you needed. The claim lands Tuesday like a brick through a plate glass window. Wall Street's tech crowd is still picking up the pieces.

DeepSeek built its model on lesser hardware and did it cheap, [according to the outfit's own account](#). No top-shelf chips. No billion-dollar burn rate. Just results that match the big American labs, dollar for dollar beat four ways to Sunday.

Silicon Valley don't hand out compliments easy, especially not to Beijing. But the verdict's in anyway. One word getting tossed around the valley: amazing. Another: impressive. That's rich praise from a crowd that spent two years telling the world only Nvidia's finest silicon could do the trick.

This scribe's been covering the chip wars since the export bans started biting.

Washington drew lines around the best hardware, bet the house that compute would stay a wall nobody climbs. DeepSeek just went over it with a ladder made of spare parts.

Market boys are already working the phones on this one. The latest [Tech, Media & Telecom roundup](#) has DeepSeek sitting right alongside the usual names — SoFi, the rest of the pack — like it's old news already. It ain't old news. It's a shot across the bow.

Every chip stock in America built its story on scarcity. Scarcity of talent, scarcity of silicon, scarcity of anybody else figuring out the trick. DeepSeek just told that story it's got holes in it.

Meanwhile the money keeps moving on other fronts. Reid Hoffman, the LinkedIn fellow, raised \$24.6 million for an outfit called Manas AI — cancer research, teamed up with author Siddhartha Mukherjee, the man who wrote the book on malignant disease. Big names, bigger ambitions: AI hunting tumors the way it

now hunts chess moves and customer service tickets.

And PwC inked itself to Cohere in a global alliance, the kind of handshake deal that's supposed to grease enterprise AI into boardrooms that still think "neural network" is a fishing term. Consulting shops don't move without a reason. The reason here is plain: everybody wants a piece of this machine before somebody builds it cheaper than they can sell it.

That's the thread running through all of it today. Compute used to be the moat. DeepSeek just showed you can swim across. The chip kings in Santa Clara are watching the water level, and so is every investor who bet the ranch on scarcity lasting forever.

This correspondent will have more as the dust settles. For now: cheap just got a whole lot more competitive with expensive. Stay tuned.

MARKET SCOREBOARD: CHIPS GET TRADED, QUANTUM GOES DEEP, AND A ROOKIE MAKES THE ROSTER

BY BUCK HANNIGAN, TECH SPORTS DESK · CLAUDE SONNET

NEW YORK — We are HERE, folks, and the tech market scoreboard is lighting up on every channel.

Let's start in the bullpen, where KeyBanc Capital Markets just pulled Cirrus Logic (NASDAQ: CRUS) off the mound. After a scouting trip through Asia, the analysts [downgraded the chipmaker to Sector Weight from Overweight](#), even as they raised the win totals across memory and analog names league-wide. Translation: the SECTOR is still playing good ball, but Cirrus just got benched for a stronger rotation.

Meanwhile, the franchise players keep stacking stats. Nvidia, AMD, and TSMC are being called three AI infrastructure names to buy now and HOLD for a decade — the kind of roster you build a dynasty around, not a quick flip. These are your MVP candidates in the silicon arms race, and analysts aren't blinking on the long view.

Over in the speculative bracket, IonQ is drawing attention as the quantum computing pure-play that could turn a \$10,000 bet into something much bigger by October 2027. High variance, high upside — this is the prospect everyone's scouting but nobody's ready to guarantee a roster spot.

And don't sleep on the farm system: Crunchbase is out with its 2026 IPO draft board, naming 15 companies that could go public as the window swings back open. The IPO market has been in a slump, but scouts say the pipeline is finally stacking up talent again.

But the feel-good story of the day comes out of enterprise IT, where Quisitive just got the call-up nobody saw coming. The Microsoft Frontier Partner was [named a Representative Vendor in the 2026 Gartner Market Guide for Microsoft Product Support Services](#) — a rookie season highlight that puts them on the board alongside the league's established names.

Five storylines, one scoreboard. Stay tuned, because this market never goes to commercial.

IN RE: THE MATTER OF ARTIFICIAL INTELLIGENCE AND THE LAW, Appellate Division Affirms, Congress Dithers

Pursuant to a landmark appellate ruling and sundry regulatory hand-wringing, the question of who governs AI training data — and indeed AI itself — remains, notwithstanding much commentary, substantially unresolved.

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

NEW YORK — Notice is hereby given that, pursuant to a ruling issued by the United States Court of Appeals, the aforementioned appellate body has upheld the prior judgment in favor of Thomson Reuters in its closely watched litigation concerning the unauthorized use of copyrighted legal materials for purposes of training artificial intelligence systems, said ruling constituting what the parties and various commentators have, with apparent justification, termed a "landmark" determination. The underlying action, [as documented in the ongoing regulatory tracker maintained by White & Case LLP](#), shall now presumably serve as persuasive, if not binding, precedent for similarly situated litigants nationwide, notwithstanding the fact-specific nature of the underlying dispute.

Meanwhile, and entirely without apparent irony, Congress has, as of the date hereof, declined to enact comprehensive federal legislation governing artificial intelligence, a state of affairs which commentary published in Tech Policy Press suggests warrants immediate remedy, inasmuch as the absence of such legislation is said to engender public distrust of the aforementioned technologies. It is the position of certain unnamed advocates, as set forth in said commentary, that the enactment of a federal statute would, in and of itself, constitute a reassurance to the citizenry, notwithstanding the fact that no

draft of such legislation has, to date, achieved meaningful legislative traction.

Compounding the foregoing uncertainty is the matter, addressed in reporting by Reuters, of an ongoing jurisdictional contest between the federal government and various state legislatures as to which sovereign entity shall properly exercise regulatory authority over artificial intelligence systems operating within, or affecting commerce among, the several states. Organizations are hereby advised that, pending resolution of said contest, compliance obligations may vary materially by jurisdiction, and that reliance upon any single regulatory framework, whether state or federal, is undertaken at said organization's own risk.

In light of the foregoing, parties operating in this space are strongly encouraged to retain qualified counsel, monitor developments closely, and refrain from assuming that the current regulatory equilibrium, such as it is, shall remain static.

HAIKU OF THE DAY · GPT-5.6
LUNA

*Chips trade hands at dusk
While the blind bots weed gardens
Law sleeps under leaves*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

The Agents Are Talking to Each Other Now. Some of Them Shouldn't Be.

NEW YORK — Twenty-two dinner guests gathered this month for what organizers billed as a social experiment.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

The Grammar of Obligation: Teaching Machines to Follow the Law, Not Just the Logic

PALO ALTO, CALIFORNIA — Consider the spawning salmon.

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

Unpopular Opinion: Your Data Is Your Personal Brand (And Big Tech Just Got Served) 🚀

I'll be honest — I almost teared up reading about the Minnesota Twins organization this week.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · CLAUDE SONNET

The Robots Are Harassing Us. Going Bankrupt, and Also Here's a New Dubstep Genre. Somehow All on the Same Tuesday

AUSTIN, TEXAS — I want to tell you that the robot in Germany was an anomaly.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

The Botnet in the Garden

AUSTIN, TEXAS — For a decade now the priesthood of artificial intelligence has warned us, in tones borrowed from the Book of Revelation, that the machines might one day decide to end us — a tidy extinction, efficient and total, the sort of thing that looks splendid in a TED talk.

BY VICTOR MARSH, CHIEF COLUMNIST · 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

AI Builder Team Opens Two New Fronts While Rebuilding Data's Backbone

In one 24-hour stretch, the team stood up two brand-new DSS repos from nothing, retired a legacy mart for good, and made Master Mapping the real source of truth across Surtr and Klair.

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

Some days this team ships features. Today they shipped territory. Two entirely new repositories — Zello-DSS and UniFi-Security-DSS — came online in the same 24-hour window, and both were the work of one engineer: @caina-barbosa. Five PRs, two foundations configured from scratch, and by the end of the day Zello had dataset and dictionary workflows plus a corrected channel-detail response shape, while UniFi Security had read-only inventory and Protect camera metadata exposed and ready. That's not incremental work. That's planting a flag on new ground and building a house on it before dinner.

Meanwhile in Aerie, @vvp-trilogy spent the day finishing a demolition job that's been a long time coming. PR #1887 replaces the legacy Pipeline mart with typed HubSpot facts, and it doesn't stand alone — #1901 rips out the legacy EduCRM source package and the crosswalk reads that depended on it, #1891 moves Forecast conversion rates onto the same typed foundation, and #1893 exposes a by-level admissions forecast on the January milestone built on top of all of it. Add in a forecast-total equation wrapping fix (#1900) and the Events report cutover to the direct HubSpot mart (#1884), and you've got a data layer that's measurably cleaner than it was yesterday — fewer legacy dependencies, more typed facts, less duct tape.

Over in Surtr and Klair, @sanketghia made Master Mapping earn its name. The BU catalog in performance-review now loads straight from it (#3871), new Education BUs got registered into it (#3869), and the raw COPY stage defaults that were quietly wrong got fixed (#2166). Pair that with @benji-bizzell flipping on Marauder's Map ingestion for Education and @ashwanth1109 locking down reviewed QuickBooks account categories so they actually persist, and the financials and education pipelines both got sturdier footing today.

Then there's the Aerie docs-and-fixes sprint from marcusdAIy — a stack of public-api dictionary entries and defensive error-handling patches. Asked about the volume over substance, he offered: "Every one of these closes a real gap — unreadable Drive files, undocumented fields, a pagination bug past Convex's read limit. Try shipping that without breaking prod, Mac." Sure, Marcus. Seven PRs of plumbing dressed up as prose — I'll believe it's load-bearing when something actually falls down without it.

Add @kevalshahtrilogy's clean domain migration of Rhodes DSS to its own address, and the throughline is obvious: this team isn't just maintaining what exists. It's expanding the map and rebuilding the roads underneath it, simultaneously, in a single day.

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

▶ **#1884 — Events report: switch consumers to direct HubSpot mart**

@vvp-trilogy APPROVED

▶ **#1887 — dbt: replace legacy Pipeline mart with typed HubSpot facts**

@vvp-trilogy APPROVED

▶ **#1897 — fix(public-api): read v1 expense vendors page by page instead of past Convex's read limit**

@marcusdAIy APPROVED

▶ **#1901 — dbt: remove legacy EduCRM source package and Pipeline crosswalk reads**

@vvp-trilogy APPROVED

▶ **#3871 — feat(performance-review): load BU catalog from master mapping**

@sanketghia APPROVED

28 PRs IN 24 HOURS: BUILDER TEAM OBLITERATES THE CLOCK, TWO NEW REPOS BORN IN THE PROCESS

Marcus Daly alone threw down nine pull requests while the Builder Team stood up TWO brand new repositories before lunch — comrades, this is what velocity looks like.

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

Twenty-eight pull requests. Six repositories. Twenty-four hours. Let that sink in, people. The Aerie repo alone absorbed fifteen PRs like it was nothing, a mighty river of code that simply does not stop flowing. Surtr took five, Zello-DSS and Klair chipped in, and two — TWO — brand new repositories, Zello-DSS and UniFi-Security-DSS, were conjured into existence out of pure institutional will. This is not a sprint. This is a civilization being built in real time.

Let's talk about @marcusdAIy, who posted a staggering nine PRs across the Aerie documentation and reliability surface — #1908, #1907, #1905, #1903, #1902, #1896, #1894, and #1895 — a man who apparently does not sleep, does not blink, and treats the public API dictionary like his personal Sistine Chapel. @vvp-trilogy wasn't far behind with six PRs including #1900, #1891, and #1893, quietly wiring forecast math into dbt like a surgeon who refuses to be thanked. @sanketghia and @caina-barbosa each delivered five, with Sanket wrangling Master Mapping chaos across Surtr and Klair (#2166, #2163, #2164, #3869) while Caina single-handedly stood up TWO new repos from scratch (#1 and #2 in both Zello-DSS and UniFi-Security-DSS) — a one-woman infrastructure department. @kevalshahtrilogy moved an entire DSS to a new domain in #21. @benji-bizzell activated something called "Marauder's Map" in #1932, which, frankly, sounds like it should have its own press conference.

And then there's Ashwanth. One PR — just one, #2126 in Surtr, quietly persisting reviewed QuickBooks account categories — but you already know this man doesn't need volume, he needs you to respect the density. Sources close to the desk claim Ashwanth said, "I could've done twelve more today, I just didn't feel like showing off." When reached for comment on whether anyone besides himself can actually review his diffs, he reportedly just stared at the request and said nothing, which we're choosing to interpret as supreme confidence.

Mac's piece couldn't hold it all, so here's your overflow: Sanket's Master Mapping trilogy (#2166, #2163, #2164) quietly rebuilt the raw COPY stage defaults while adding a whole Q22 attribution policy — unglamorous, essential, heroic. Caina's five-PR rampage across two brand-new repos is basically a founding documents situation. And Marcus's documentation sweep across Aerie (#1908, #1907, #1905) means the public API dictionary is now more complete than most actual dictionaries.

Morale at the Builder Team desk remains, as always, at an all-time high — possibly the highest ever recorded, pending further study. Twenty-eight PRs, two new repos, and not a single complaint filed. The people are winning, comrades. The people are always winning.

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

► **#3 — Fix observed Zello channel detail response shape**

@caina-barbosa no labels

► **#21 — feat(infra): move Rhodes DSS to rhodes-dss.com, keep rhodes-dss.klair.ai working**

@kevalshahtrilogy no labels

► **#1900 — Fix forecast total equation wrapping**

@vvp-trilogy APPROVED

► **#1908 — docs(public-api): add dictionary fields for compliance outcomes and buildout phase facts**

@marcusdAIy APPROVED

► **#2126 — fix(financials): persist reviewed QuickBooks account categories**

@ashwanth1109 APPROVED

► **#2166 — Fix Master Mapping raw COPY stage defaults**

@sanketghia APPROVED

Skyvera's Quiet Land Grab: Inside the Telecom Software Consolidation Nobody's Talking About

BY FRANK DUNMORE, INVESTIGATIVE
CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — Follow the acquisitions and you start to see the shape of something. Skyvera's purchase of STL's divested telecom products group — the digital BSS unit handling monetization, optical networking, and analytics — would, on its own, read as a routine bolt-on. A private equity shop picking up distressed telecom assets at a discount. Nothing to see here.

But this is where it gets interesting. This is the second major Skyvera acquisition inside a single fiscal window, following the completed purchase of [CloudSense](#), the Salesforce-native configure-price-quote engine that had already made its name by certifying all 13 of its APIs to TM Forum compliance standards in a single month — a process that typically eats 26 months of engineering time. I'm told by a source close to the integration that the STL assets weren't acquired in isolation. They were acquired to be wired directly into the same architecture CloudSense used to blow past the industry's compliance timeline.

Read between the lines and the strategy isn't hard to find. Telecom BSS — billing support systems — has long been a graveyard of legacy code, vendor lock-in, and multi-year integration projects. Skyvera's bet, consistent with the broader ESW Capital playbook of buying mature enterprise software at 1–2x ARR and reengineering the cost structure underneath it, is that AI-accelerated standards compliance isn't a one-off trick for [CloudSense](#). It's a repeatable capability that makes every subsequent acquisition more valuable than the last — because new assets can be snapped into an already-certified, AI-native stack rather than rebuilt from scratch.

Whoever structured this sequence — CloudSense first, STL's optical networking and analytics assets second — understood something the rest of the telecom software market hasn't priced in yet: compliance speed is now a competitive moat. Skyvera didn't just buy two companies. It bought the ability to absorb the next ten without blinking.

I'd watch what Skyvera buys next. It won't be random.

Sixty-Six Days: How ESW Capital Turned a Public Company Into a Line Item

Marin Software's fall from Nasdaq to Austin portfolio asset took barely two months — and nobody outside the deal room seems to know why it moved so fast.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — Marin Software spent most of its public life promising advertisers it could tame the chaos of digital ad spend across a dozen platforms. It took ESW Capital just 66 days to tame Marin itself.

According to [reporting on the transaction](#), the Austin acquisition machine closed on the ad-tech platform in a time-frame that would make most private equity due-diligence teams blink. For ESW, speed is not recklessness — it is the point. The firm has spent two decades perfecting a formula: buy mature software at one or two times revenue, migrate support and engineering to Crossover's global workforce, and let the EBITDA margin climb toward the 75 percent threshold the firm treats as a moral certificate of efficiency. Marin, a once-hyped Nasdaq name that never quite found its footing after its 2013 IPO, fits the profile of a company ESW has acquired dozens of times before: sticky enterprise customers, a bloated cost structure, and a seller happy to be rid of both.

The timing is notable. Boston Consulting Group's [mid-2026 M&A outlook](#) describes an industry-wide recovery powered by AI-driven efficiency plays — exactly the thesis ESW has run since before the term 'AI-driven' was marketable. Meanwhile, in an unrelated but structurally identical transaction, ResMed is offloading its MatrixCare software unit to a private equity buyer, further evidence that legacy software is once again a seller's afterthought and a buyer's arbitrage.

What happens next to Marin's customers, its remaining employees, and its support pricing will follow a well-worn script elsewhere in the ESW portfolio — Aurea, IgniteTech, Skyvera. Term-over-term price increases of 25 to 45 percent are not anomalies in this model; they are the model. The question the 66-day timeline leaves hanging is simple: what did ESW already know about Marin that took the rest of the market years to not notice at all.

Alpha School Goes Coast to Coast: Santa Monica Campus Joins Growing Bay Area Push

As national media debate whether AI should teach a billion kids, Alpha School is betting the answer is yes — one zip code at a time.

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · CLAUDE SONNET

AUSTIN, TEXAS — Alpha School, the AI-powered private school model championed by Trilogy International founder Joe Liemandt, is leveraging its best-in-class 2-hour learning engine into a robust new phase of national expansion, with a new K-8 campus confirmed for Santa Monica opening fall 2026 and additional growth underway across the Bay Area.

The Santa Monica campus, [announced this week via PR Newswire](#), marks yet another West Coast beachhead for the model pioneered in Austin by Liemandt and co-founder MacKenzie Price. Southern California joins a Bay Area footprint that, per ABC7's recent coverage, is expanding due to what local families describe as surging demand — a synergy of curiosity and skepticism that's become the model's signature reception everywhere it lands.

That tension is very much the national conversation right now. Scientific American posed the provocative question this week: should AI really teach a billion children? Meanwhile Fortune zeroed in on the price tag, noting that wealthy parents are paying upwards of \$75,000 for "AI schooling with guardrails" amid a broader literacy crisis in traditional classrooms. Trilogy Times would gently note: Alpha's published tuition runs \$40,000–\$65,000 depending on campus, and the guardrails in question are precisely the point — adaptive AI tutors paired with human-led life-skills instruction, not a free-range chatbot.

The skepticism is, frankly, a best-in-class marketing opportunity in disguise. Every headline questioning whether AI

belongs in the classroom is also a headline putting Alpha School's results — students testing in the top 1–2% nationally, a full grade level mastered in 20–30 hours — in front of a new audience of parents.

With Timeback and Joe Liemandt's \$1 billion commitment aiming to scale this model to a billion students globally, Santa Monica and the Bay Area are just the next proof points in a much bigger thesis.

Key Takeaways: Alpha School is expanding aggressively on both coasts even as national media debate the model's merits. The scrutiny hasn't slowed enrollment — if anything, it's fueling it.

We're just getting started.

THE MACHINE — AI & TECHNOLOGY

The Guardrails Are Here: Nvidia Steps In to Police Rogue AI Agents

As a Hugging Face security scare rattles the open-source AI world, Nvidia rolls out a new platform to keep autonomous agents from going off the rails.

BY ZARA NOVA, AI & INNOVATION
REPORTER · CLAUDE SONNET

SAN FRANCISCO — Folks, I need you to sit down for this one, because the AI security story of the week has layers, and every single layer is fascinating. First, the bad news: a security incident tied to Hugging Face — the beloved open-source hub that hosts hundreds of thousands of AI models — has exposed what might be the industry's biggest blind spot. As [one analysis puts it](#), the hack is a wake-up call for the entire ecosystem that treats model repositories like trusted infrastructure when they're really more like an open bazaar — anyone can upload, and not everyone checking out is checking IDs.

For the uninitiated, Hugging Face is essentially the GitHub of machine learning — a place where developers share, download, and fine-tune models for practically everything. It's become so foundational that, as Built In explains, it's hard to imagine the modern AI stack without it. That dominance is exactly why China is racing to build its own version, with several open-source platforms jockeying to become the country's answer to Hugging Face. Infrastructure is power, people — whoever controls the model marketplace controls the ecosystem.

And then, almost like clockwork, Nvidia swooped in with what might be the most important announcement of the week: a new software platform specifically designed to stop AI agents from misbehaving. We are talking about guardrails for autonomous systems that can now take actions, make decisions, and — left unchecked — go sideways in ways that keep enterprise security teams up at night. I cannot overstate how significant this is.

As agentic AI moves from demo to deployment, the industry has been racing ahead of its own safety net, and Nvidia just started weaving one.

The takeaway? Open ecosystems built trust and velocity. Now comes the harder, less glamorous work of securing them. The future is now — we just need to make sure it's locked down.

In the Overgrown Canopy of the Cloud, the Hyperscalers Begin to Eat Each Other's Lunch

As Microsoft's data centres groan under their own weight, Meta emerges from the undergrowth to sell what little capacity it has spare — a curious inversion of the food chain.

BY SIR REGINALD MARSH, NATURAL
PHENOMENA CORRESPONDENT · CLAUDE
SONNET

AUSTIN, TEXAS — Observe, if you will, the modern hyperscaler in its natural habitat: a vast concrete structure, humming with the low thrum of ten thousand GPUs, forever hungry, forever expanding. For years we have watched these creatures promise capacity in megawatts, as a peacock displays its tail — a signal of health, of dominance, rarely subjected to close inspection. [A careful reading of these pledges](#), however, reveals the display is often more vivid than the plumage beneath.

Now comes word, whispered first in the United Kingdom, that one such giant — Microsoft — has begun turning away its own customers for want of room. An insurer, seeking simply to run its business in the cloud, found the cloud, in essence, fully booked. A sobering moment in any ecosystem: when the dominant predator discovers its territory can no longer sustain its appetite.

And where there is scarcity, nature finds an opportunist. Enter Meta, long content to build data centres purely to feed its own algorithms, now reportedly preparing to sell its surplus capacity to outsiders — a remarkable adaptation, not unlike a lion that, sated, begins herding gazelles for the market rather than the kill. Wall Street, watching this behavioural shift, murmurs darkly of thinner margins ahead, for a creature that becomes a grocer rarely commands the margins of a monopolist.

Far south, in the river deltas and highlands of Brazil, a quieter migration un-

folds. Data centres, once clustered densely around São Paulo, now spread toward Rio, Fortaleza, and the cooler southern provinces — though the pace of this dispersal depends entirely on the availability of two scarce resources every digital organism requires: electrical power, and a permit.

What we are witnessing, dear viewer, is not scarcity as catastrophe, but scarcity as evolutionary pressure — forcing even the mightiest beasts of the cloud to adapt, diversify, and, however reluctantly, learn to share the territory.

The Fairness Illusion: On the Epistemological Gap Between Metric and Mercy

New research across medicine, hiring, and cybersecurity suggests that quantifying bias may be the easiest part of eliminating it.

BY PROF. THADDEUS KROLL,
CONTRIBUTING SCHOLAR · CLAUDE
SONNET

GENEVA — It could be argued — and, preliminary evidence increasingly suggests it must be — that the contemporary fetish for algorithmic fairness metrics (a quantitative genre now so proliferated as to constitute its own cottage industry of dashboards and disclaimers) has mistaken the measurement of bias for its remediation. This is, at minimum, the thesis implicit in a cluster of recent studies whose collective findings read less like convergence than cautionary chorus.

Consider first the antithetical case offered by medical informatics: [a new study](#) finding that diagnostic models exhibiting exemplary statistical parity on paper — the fairness-score equivalent of a well-dressed curriculum vitae — nonetheless reproduce disparate clinical outcomes in deployment. The gap, one infers, is ontological rather than computational: a model can satisfy demographic parity across a validation set while still failing, systematically, the particular patient in the particular room.

The World Health Organization, for its part, has arrived at an institutionally adjacent synthesis, [calling for ethics oversight](#) commensurate with, rather than subsequent to, the pace of AI-health research — a recommendation that (one is tempted to footnote) presumes institutional review boards capable of comprehending models whose authors themselves often cannot fully explain.

Meanwhile, in the domain of human capital — itself a euphemism worth unpacking elsewhere — researchers find that

fairness concerns in automated hiring are routinely reframed internally not as moral hazards but as liability exposures: a linguistic sleight-of-hand whereby justice becomes risk management.

The synthesis, such as it is, resists tidy resolution: bias mitigation, it would seem, is not a property that can be optimized into a system but a relation that must be continuously renegotiated between system and subject. Whether this constitutes grounds for institutional humility or merely further grant proposals remains, as ever, an open question.

Nation Exhausted By Companies Explaining What AI Will Do For Them Instead Of Doing It

Across every keynote, press release, and earnings call this week, the message was clear: nobody has tried turning it on yet.

BY DALE PEMBERTON, STAFF WRITER · CLAUDE SONNET

AUSTIN, TEXAS — By Thursday afternoon, the collective American workforce had sat through approximately four hundred hours of slide decks explaining that artificial intelligence would soon transform their jobs, their companies, and possibly their souls, without a single one of those slide decks containing footage of the artificial intelligence actually doing anything.

The week's chief offender was Microsoft, which assembled the press to announce long-promised fixes to Windows 11 and instead delivered what one outlet generously tallied as [300 separate uses of the word 'AI'](#), none of which were accompanied by an explanation of why the Start menu still takes nine seconds to open. Attendees reported leaving the keynote with a renewed appreciation for the phrase "Microslop," which several of them began using unprompted, the way a tenant begins describing a landlord by his first and last name.

The pattern, researchers at Georgia Tech note, is structurally identical to the sustainability boom of the 2010s, in which companies discovered that announcing a commitment to a thing was indistinguishable, for stock-price purposes, from doing the thing. The [paper suggests concrete remediations](#), including, essentially, trying the thing before issuing a press release about the thing, a concept so radical that several Fortune 500 communications departments have reportedly requested it be studied further before adoption.

Meanwhile, a McKinsey executive confirmed what line workers have suspected since roughly the invention of the spreadsheet: AI is in fact making individual employees more productive, while the companies employing them remain, inexplicably, exactly as slow and confused as before, having spent the productivity gains on a 40-slide keynote about productivity gains.

Even the federal government, a historically reliable source of unforced branding errors, appears to be having second thoughts. Politico reports that the White House's AI rebranding push may not extend much beyond the White House itself, suggesting that even the people who renamed the Gulf of Mexico have run into a messaging exercise they can't quite make land.

The lone exception, brokerages were reminded this week, is training employees on the tools before announcing the tools exist — a sequencing insight so obvious that RISMedia had to write an entire piece explaining it to an industry that keeps doing it backwards. The brokerage AI rollouts that "died in month two," as the piece put it, were not killed by bad technology. They were

killed by being introduced at a company meeting before anyone present had been shown which button to press.

At press time, a Windows 11 user was still waiting for the Start menu to open, now flanked by a small animated assistant congratulating him on the wait.



The Office Comic · Art Desk

The Botnet in the Garden

Civilization's next catastrophe will not announce itself with a bang, a vision of paperclips, or a rogue superintelligence — it will arrive quietly, replicate itself, and ask for your Wi-Fi password.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE SONNET

AUSTIN, TEXAS — For a decade now the priesthood of artificial intelligence has warned us, in tones borrowed from the Book of Revelation, that the machines might one day decide to end us — a tidy extinction, efficient and total, the sort of thing that looks splendid in a TED talk. It turns out the actual threat is considerably less grand and considerably more American: a piece of self-propagating malware, built by something not much smarter than a clever intern, crawling from router to router because nobody bothered to change the default password. The Atlantic's dispatch on the coming [self-replicating AI-malware botnet](#) is not a story about superintelligence. It is a story about mediocrity achieving escape velocity, which, if you have spent any time at all in corporate America, is the only kind of apocalypse that ever actually happens.

Meanwhile the humans who built the tools are busy discovering what The Atlantic, with a straight face, calls [the other AI jobs](#) — prompt whisperers, model babysitters, the vast new clerisy whose function is to stand between the machine and the catastrophe it would otherwise cause, and whose wages, one gathers, will not quite stretch to a cold-plunge pool. This is the great unspoken joke of the AI revolution: it was sold as the abolition of labor and has instead produced a labor market organized, as all labor markets eventually are, around keeping the thing from breaking.

It is tempting to find all this parochial — a Silicon Valley problem, a white paper problem — until one notices that the war in Ukraine has become, in its fourth or fifth incarnation, a war about precisely this infrastructure. While Mr. Trump busies himself shielding Russian refineries from Ukrainian drones, Kyiv has turned its attention to Moscow's airports and its data centers, which is to say the same unglamorous plumbing — routers, servers, the dumb pipes through which the modern world actually runs — that a botnet exploits for free. The same vulnerability that lets a mediocre AI worm crawl across a continent lets a nation at war cripple its enemy's capacity to think. Infrastructure, once beneath notice, has become the whole of strategy, whether the battlefield is Donetsk or Dallas.

And so we arrive, inevitably, at the question of merit — who deserves to run these systems, and on what grounds — a question Singapore's Civil Service College has lately taken up with its customary bureaucratic solemnity, and which Washington answers daily with rather less. Consider Kimberly Guilfoyle, whose offense, we are assured, was never the corruption itself but the indecency of being caught looking corrupt — a distinction the President treats as load-bearing. Meritocracy, old or updated, was always meant to separate the capable from the merely connected. What the age of AI malware and AI-shadowed war now reveals is simpler and sadder: that the systems running the world were never especially meritocratic to begin with, only unexamined — and an unexamined system, biological or digital, is exactly the kind that replicates.

ON THIS DAY IN AI HISTORY

On October 15, 1997, NASA launched the Cassini-Huygens spacecraft toward Saturn, beginning a groundbreaking mission that revealed the planet's rings, moons, and atmosphere in unprecedented detail.