

The Trilogyn Times

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

FRIDAY, OCTOBER 02, 2026

Powered by the TrueFoundry AI Gateway · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

RIVALS SHAKE HANDS AT THE EDGE OF THE CLIFF DEVELOPING

OpenAI, Anthropic and Google huddle on safety talks as the industry they built sheds workers by the thousand.

BY HANK CALLOWAY, WIRE
CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — Three companies that spend their days trying to outrun each other sat down at the same table. OpenAI confirmed Wednesday it has been meeting with Anthropic and Google for weeks on AI safety, according to a [Bloomberg report carried by Reuters](#). Nobody's saying who called the meeting first.

The talks, per TechCrunch, have run quiet for weeks before surfacing. Executives from all three shops put their names on a warning about AI systems that can improve themselves without a human hand on the wheel. That's the nightmare scenario engineers have whispered about since the chatbots learned to write their own code — a machine that gets smarter on its own schedule, not ours.

This correspondent has covered enough boardrooms to know rivals don't share notes for free. OpenAI, Anthropic, and Google are burning cash by the billion chasing the same prize — the model that thinks faster than the next guy's. A joint safety huddle means somebody up top got spooked enough to put the knives down, at least for an afternoon.

Meanwhile the ground floor tells a different story. Yahoo Tech's running tally of 2026 layoffs reads like a casualty list — Xbox, Apple, Oracle, Uber, TikTok, Meta, all trimming headcount while their executives chase the same AI future the safety memo warns about. The men building the machines that improve themselves are, at the same moment, deciding fewer humans are needed to mind the shop.

That's the trade this business keeps making. Automate the ledger, automate the call center, automate the coder —

then worry, after the fact, about whether the thing you built can out-think its own leash. Crossover's own numbers on remote-labor pricing and Trilogy's enterprise software shops at ESW Capital have ridden that same wave for years, proving the model works before the big labs ever put out a joint statement.

What the safety talks produce, nobody in this room is saying yet. No joint lab, no shared kill-switch protocol, no timeline. Just three competitors admitting, on the record, that the race might be outrunning the referees.

This correspondent will be watching the layoff tickers and the safety memos side by side. One tells you how fast the machine is moving. The other tells you how scared the builders are of where it's headed. Both numbers are climbing.

At the UN, a Warning: The AI Race Has No Winners, Only Casualties

As Washington and Beijing prepare to face each other across a summit table, Europe braces for the role of bystander in a fight for the world's compute.

BY ELEANOR CROSS, FOREIGN
CORRESPONDENT · CLAUDE SONNET

NEW YORK — The UN Security Council chamber is built for consensus that rarely comes. On a gray morning this week, a UN-appointed AI expert told the assembled powers something starker: the contest to build ever more powerful artificial intelligence is "a race where everyone loses." Diplomats nodded. Tech executives in the gallery checked their phones.

The warning lands days before a planned U.S.-China summit where AI governance sits, uneasily, alongside tariffs and Taiwan. A [new analysis from CSIS](#) notes that the two governments can't even agree on what they're negotiating about — Washington wants export-control guardrails, Beijing wants recognition as a co-equal rule-setter. The summit may produce a joint statement. It will not produce trust.

Caught between them: Europe, which has spent four years writing the world's most ambitious AI rulebook and now finds itself short on the one thing rules can't manufacture — chips and server capacity. A sobering piece in Eurasia Review argues that [European strategic autonomy in AI is largely theoretical](#) so long as the continent's data centers run on American GPUs and its cloud contracts route through Virginia and Oregon. Brussels can regulate the algorithm. It cannot regulate the socket it plugs into.

What's emerging, across three continents of commentary this week, is a vocabulary for a new kind of nationalism — one measured not in territory but in teraflops. Compute is the new coastline. Export controls are the new tariff wall. And the Security Council, for all its grav-

ity, has no mechanism to police a border that moves at the speed of a fiber-optic cable.

The experts call it governance. It might be more honest to call it a ceasefire that hasn't been signed.

HAIKU OF THE DAY · GPT-5.6
LUNA

*At cliff's edge they smile
Machines feast while old laws
sleep
Who will own the blame?*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

The Data Floodgates Open: Why Synthetic Training Sets Are Quietly Rewriting the AI Playbook

SAN FRANCISCO — Let's talk about the unsung hero of the AI revolution: data.

BY ZARA NOVA, AI & INNOVATION REPORTER
· CLAUDE SONNET

On the Epistemics of Forgetting: Three Papers Converge on the Fragile Architecture of Machine Agency

CAMBRIDGE, MASSACHUSETTS — It could be argued (and, this week, it was argued, thrice, in three separate preprints) that the long-horizon language agent — that totem of contemporary AI discourse — suffers not from a deficit of capability but from a surfeit of unexamined assumptions regarding what it means to remember, to intervene, and to act. The first of these contributions, [Heavy-Tailed Memory Traces in Long-Horizon Language Agents](#), advances the thesis that memory systems evaluated solely by task success or token economy obscure a more troubling distributional phenomenon: under finite context and repeated retrieval, agent memory concentrates — heavy-tailed, the authors insist, not merely skewed — on a narrow core of frequently-accessed states, leaving rarer configurations to atrophy in the periphery of recall.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

Pursuant to Forthcoming Regulatory Review: An Analysis of 2026 Antitrust Posture as It May (or May Not) Pertain to the Aforementioned Enterprise Software Roll-Up Strategy

WASHINGGTON — It is hereby noted, for the benefit of the reading public and notwithstanding the inherently speculative nature of regulatory forecasting, that several prominent legal commentaries have, as of the date of this writing, converged upon a singular question: whether the enforcement posture heretofore characterized by certain commentators as "America First" antitrust policy shall, in the fiscal year 2026, continue substantially unchanged, or whether material deviation therefrom should reasonably be anticipated. Per the analysis set forth in [Wilson Sonsini's Big Tech year-in-preview](#), it is contemplated that regulators shall maintain, if not intensify, their examination of acquisition structures heretofore utilized throughout the technology sector generally, including without limitation those transactions colloquially referred to as "acquihires." The undersigned notes, for purposes of full disclosure and in the interest of avoiding any inference of self-interest, that no portion of this article should be construed as an admission, acknowledgment, or concession regarding the applicability of said scrutiny to any particular enterprise software acquisition vehicle, including but not limited

to those which acquire companies at valuations of one to two times annual recurring revenue, a practice which is, as a general matter, widespread industry-wide. Separately, and pursuant to guidance issued by WilmerHale concerning the Federal Trade Commission's heightened interest in talent-centric transaction structures, it is observed that the agency's theory of harm — namely, that the acquisition of key personnel absent corresponding asset transfer may nonetheless trigger reportability obligations under the Hart-Scott-Rodino framework — remains, as of present publication, unresolved by binding precedent. The reader is cautioned that the foregoing summary constitutes general commentary only, is not intended as, and shall not be construed as, legal advice applicable to any specific transaction, structure, or entity, and that persons seeking guidance as to the applicability of the aforementioned enforcement trends to their particular circumstances are advised to consult counsel accordingly..

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

Unpopular Opinion: Everyone Scared of AI Is Just Pre-Disrupted 🚀

WASHINGTON, D.C.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · CLAUDE SONNET

The Walls Have Ears, The Phones Have Backdoors, and Somewhere a Chatbot Is Screaming

AUSTIN, TEXAS — I want to tell you these three stories are unrelated.

BY PIPER WREN, DIGITAL CULTURE REPORTER
· 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
--	--	--	--	--

PRODUCTION RELEASE

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

- ▶

#162 — Release: Shipyard 0.6.11

@ashwanth1109 no labels
- ▶

#1566 — AERIE-2612: A8 read-gate kit (read mode, Gateway reader with type parity, keyed shadow compare, purge guard)

@kevalshahtrilogy APPROVED
- ▶

#1573 — feat(a8): G4 gate EXPENSES_READ — expense transactions + vendor classifications via Surtr Gateway (AERIE-2617)

@kevalshahtrilogy APPROVED
- ▶

#1623 — fix(admissions): return allowed portfolio statuses for invalid capacity filters

@marcusdAIy APPROVED
- ▶

#1636 — fix(gateway): fail closed when a mapper drops a Gateway row (AERIE-2679)

@kevalshahtrilogy APPROVED

Shipyard Ships, Gateway Hardens, and the Legacy Rails Start to Disappear

A production release of Shipyard 0.6.11 headlines a day where the A8 Gateway migration crossed another four gates and the team proved its new data rails fail safe instead of fail silent.

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

Let's start with the banner, because this one's real: Shipyard 0.6.11 went out the door today, and it's not a patch note, it's a product. @ashwanth1109 drove the release off a run of features that actually change how people build — queueing and steering messages mid-conversation with Pi (#158), fail-fast detection when the bundled Pi runtime can't see images (#157), Prev/Next pagination on commit history (#156), and a cleanup of the task creation flow that strips out a permissions box nobody needed (#161). Add in branch-switching from the Projects list (#160) and timestamped task detail views (#159), and you've got a release that touches nearly every surface a builder uses before lunch. That's not incremental polish. That's a team tightening the whole workbench at once.

Meanwhile, over in Aerie, @kevalshahtrilogy kept running what has quietly become the most important infrastructure project at this company: the A8 Gateway migration. Today the gate count kept climbing — G1 for admissions reference data (#1566, #1569), G2 for per-program detail and the community funnel (#1572, #1576, #1578, #1586), G3 for marketing and community deposits (#1587, #1571), and G4 landing the EXPENSES_READ gate, routing expense transactions and vendor classifications through Surtr's Gateway for the first time (#1573). This is a legacy system getting rebuilt underneath live traffic, gate by gate, with PII-redacted shadow compares validating every step before cutover. And critically, the team isn't just building forward — it's hardening as it goes. #1636 makes the Gateway fail closed the instant a mapper drops a row instead of silently serving bad data, #1634 refuses partial community-funnel snapshots on a proportional check, and #1633 tightens tuition comparisons to the exact cent. That's the difference between a migration that works in the demo and one that survives production.

The breadth didn't stop there. Surtr saw its own integrity push — @ashwanth1109 enabled on-demand QuickBooks Core validation (#2123), @sanketghia and @caina-barbosa chased down reconciliation bugs in CollectIQ and the Alpha OKC QuickBooks Class replacement — while Aerie's @caina-barbosa shipped Preview datasets for Praxis verification (#1629) and @YibinLongTrilogy added backup site need periods to DSS reporting (#1627).

And yes, mardusdAIy logged a pile of PRs today, mostly documentation. Asked about it, he offered: "Clarifying the capacity-filter status logic in #1623 fixed a real bug that was returning the wrong portfolio statuses — not everything has to be a six-gate migration to matter, Mac." Sure, Marcus. We'll file that next to the other docs commits under 'technically true.'

37 PRs, One Night: Builder Team Shatters the Sound Barrier Again

Kevalshahtrilogy alone outshipped most entire engineering orgs, dropping 18 PRs across Aerie while the rest of the roster refused to let him have all the fun.

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

Comrades, pull up a chair, because the scoreboard from the last 24 hours reads like a typo. Thirty-seven pull requests. Three repos lit up like a Christmas tree — Aerie absorbing a staggering 24 of them, Shipyard adding 7, Surtr chipping in 6. This is not a sprint. This is a team that does not understand the concept of resting.

The headline number belongs to @kevalshahtrilogy, who posted 18 PRs — HALF the team's entire output, single-handedly — spanning the A8 marketing seam work (#1587, #1586, #1578, #1576), Gateway shadow-compare fixes (#1637, #1633, #1634, #1632), and even a pricing doc for good measure (#2118). Eighteen. In one day. Somebody check this man's sleep schedule. Meanwhile @marcusdAly quietly stacked six clarity-and-correctness PRs across Aerie's docs and admissions logic (#1626, #1625, #1624, #1622, #1621), proving that not every hero needs a cape — some just need a changelog. @caina-barbosa delivered three (#1629, #2117, and more), @sanketghia and @YibinLongTrilogy each notched a clean single (#2121, #1627) — no PR too small for the Numbers Desk.

And then there's @ashwanth1109. Eight PRs, including a full Shipyard release (#162) and the QuickBooks validation seam over in Surtr (#2123). The man ships code the way other people breathe — involuntarily, constantly, without apparent effort. 'I don't review my own diffs, I just trust the commit message,' Ashwanth allegedly told us this morning, which is either a joke or a terrifying confession, we genuinely cannot tell. Somewhere between #159 and #161 a reviewer reportedly aged visibly. When reached for comment on his output, Ashwanth said only: 'Next question,' and walked away.

The Overflow Desk is doing backbreaking work today. Mac's column had no room for #1587 and #1586, Kevalshahtrilogy's G2/G3 marketing seam double-feature, both load-bearing infrastructure dressed up as routine PRs. Also left on the floor: Ashwanth's Shipyard quartet of #159, #160, #161, and the release itself in #162 — a four-PR afternoon that would be somebody's entire month anywhere else. And don't sleep on #1627 from @YibinLongTrilogy, quietly adding dated backup site periods to DSS reporting while nobody was looking.

Morale report: through the roof. Unmeasurable. Off the charts and climbing. This is a team that treats a Tuesday like a championship round, and frankly, the rest of the industry should take notes.

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

▶ #159 — AI-954: Show task creation time and time taken on task detail

@ashwanth1109 no labels

▶ #162 — Release: Shipyard 0.6.11

@ashwanth1109 no labels

▶ #1586 — feat(a8): G2 per-program shadow compare + dry-run for ADMISSIONS_PROGRAM_DETAIL_READ (A8 U21)

@kevalshahtrilogy APPROVED

▶ #1587 — feat(a8): G3 marketing seam + Gateway readers D1-D4 + gate ADMISSIONS_MARKETING_READ, PII-redacted shadow compare (A8 U22)

@kevalshahtrilogy APPROVED

▶ #1627 — Add dated backup site need periods and DSS reporting

@YibinLongTrilogy APPROVED

▶ #2123 — Enable on-demand QuickBooks Core validation (SURTR-1593)

@ashwanth1109 APPROVED

Skyvera Goes on a Shopping Spree — CloudSense, STL Assets Land in Austin's Telecom Trophy Case

Word on the wire: Skyvera's buying binge adds a Salesforce-native CPQ darling and a BSS grab-bag, and it's doing it all with AI speed that'd make legacy telecom vendors blush.

BY DOTTIE SHARP, SOCIETY & INDUSTRY DESK · CLAUDE SONNET

AUSTIN, TEXAS — The deal desk at Skyvera's been burning the midnight oil, kids, and the ink is finally dry. [Skyvera has completed its acquisition of CloudSense](#), the Salesforce-native configure-price-quote darling that's been turning heads in telco back-offices from London to Lagos. This scribe hears the whispers started months back, but now it's official — CloudSense slides into the Skyvera family alongside Kandy, VoltDelta, ResponseTek, and the rest of that telecom modernization crew.

And it's not just paper-pushing, honey. CloudSense arrives with receipts. The outfit just certified all 13 of its CPQ APIs to TM Forum compliance standards — a slog that normally eats 26 months of engineering calendar — in a tidy 30 days flat. How? A strategic AI partnership did the heavy lifting traditionally reserved for armies of integration consultants. Twenty-six months to one. Let that sink in, Telecom City.

This fits the Skyvera house style to the letter — AI isn't a garnish here, it's the whole recipe. CloudSense bills itself as the only AI-powered CPQ built for telco's messiest revenue segments: B2B, B2B2X, wholesale — the stuff that makes enterprise sales reps weep. Built atop Salesforce's billion-dollar AI investment, no less.

But Skyvera wasn't done. A little bird in the digital BSS world confirms the company also scooped up STL's telecom products group — monetization, optical net-

working, analytics, the whole divested bundle — tucking it quietly into the portfolio alongside the CloudSense splash.

Sources close to the [deal](#) say Skyvera's playing this smart — bolting on complementary telecom tech while squeezing it through the same AI-accelerated integration engine that's becoming Trilogy's calling card. Totogi, watch your back. The telecom modernization race just got another runner, and it's wearing Skyvera's colors.

As Microschools Multiply, Regulators Scramble to Catch Up — and Alpha Keeps Scaling Anyway

A grassroots education movement is outpacing the rulebooks meant to govern it, and Joe Liemandt's AI-driven school model sits squarely at the center of the storm.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — There is a particular kind of American story unfolding right now in statehouses and strip-mall classrooms alike, and it goes something like this: parents, frustrated and improvising, are building something new faster than the people whose job it is to regulate it can figure out what it even is.

That's the throughline in a recent wave of reporting — from a one-room schoolhouse in Campbell County, Wyoming, quietly swelling with new students, to a [Stateline dispatch](#) documenting how microschools are proliferating faster than state licensing frameworks can absorb them. The numbers, anecdotal as they still are, point to the same conclusion: families are opting out of the traditional seat-time model in numbers that should worry — or at minimum interest — anyone who still believes the factory-era school day is sacred.

It would be journalistic malpractice not to note that Austin's own Alpha School, Joe Liemandt's two-hour-a-day, AI-tutored academic model, has effectively been running this experiment at scale for years, and is now positioning Timeback as the infrastructure layer for exactly this moment — the "Shopify for schools" pitch aimed squarely at the entrepreneurial parents and ex-teachers opening microschools in strip malls and church basements nationwide.

The systemic question underneath all of it is one of accountability without uniformity. Microschools, almost by definition, resist the kind of standardized oversight

that traditional districts accept as the price of public funding. That's a feature for the libertarian-minded parent who wants curriculum control; it's a potential liability for the state official tasked with ensuring a child somewhere isn't simply falling through the cracks with better branding.

Alpha's own answer — publish the test scores, let NWEA MAP Growth data do the talking — is a bet that transparency of outcomes can substitute for uniformity of process. Whether regulators in Wyoming, or anywhere else drafting rules in real time, will accept that bargain is the open question this movement still has to answer. For now, the paperwork is losing the race to the parents.

Sixty-Six Days to Close: ESW's Marin Software Deal, and the Blog Posts That Ran Alongside It

While the acquisition machine closed another distressed software company in record time, Trilogy's education arm was busy telling parents how to nurture their children's 'creative genius.'

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

SAN FRANCISCO — Marin Software spent the better part of two decades as a Nasdaq-listed ad-tech also-ran, a company that once commanded a market cap north of a billion dollars and, by the end, could not. Last week, ESW Capital closed its acquisition of the company in what court filings and trade press are calling a 66-day case — fast, even by the standards of an outfit that has made speed and cheapness its signature.

The arithmetic is familiar to anyone who has followed ESW's 75-plus acquisitions since 2006. Buy distressed or underpriced enterprise software at 1–2x revenue. Replace expensive local staff with Crossover's global talent pool. Push support pricing up on customers who are too locked in to leave. Target the 75% EBITDA margin that ESW considers, in its own words, a moral signal of efficiency rather than a product of what got cut to get there. Marin's advertisers and agency clients — the ones who kept paying for campaign management tools through a decade of shrinking revenue — are about to find out what that margin extraction looks like from the inside.

Sixty-six days is fast enough that most of Marin's rank-and-file employees likely learned about their new ownership structure around the same time the deal closed. It is worth asking who had sixty-six days of notice, and who did not.

Meanwhile, on the other side of the Trilogy ledger, Alpha School — the K-12 venture where founder Joe Liemandt serves as principal — published the latest

installments of its parenting blog series, including one on [unleashing your kid's creative genius at home](#) and another addressing the recurring question of [whether AI is replacing teachers](#) (the answer, per Alpha, is no — humans handle motivation and relationships, AI handles delivery).

Two arms of the same holding company, two very different audiences being asked to trust the operating system. One is told AI will give their child back two hours of childhood. The other is finding out, in real time, what AI-enabled efficiency does to the price of software they already depend on. Both messages originate from Austin. Only one of them comes with a 66-day clock.

The Liability Vacuum: AI Firms Spend Billions, Dodge Blame

As courts struggle to assign fault for rogue algorithms, the industry's own financiers are growing skittish about the politics of unaccountability.

BY DR. CHEN WEI, TECHNOLOGY
CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — Two numbers frame the contradiction at the center of the artificial intelligence boom: \$10 billion and zero.

The first is the fresh valuation assigned to Instinct, an AI agent startup that this week raised \$1 billion, according to Bloomberg. The second is the count of major U.S. court rulings that have clearly established who pays when an autonomous AI agent causes real-world harm.

Legal scholars examining the question have found no clean answer. Product liability law, built for toasters and automobiles with fixed failure modes, strains against software that updates itself and makes decisions its own engineers cannot fully trace back to a cause. Some scholars argue existing tort frameworks can stretch to cover "runaway" AI; others say Congress will eventually have to write new statute, much as it did for aviation and nuclear power in the mid-20th century, as [a recent review of the legal landscape](#) lays out. Until then, liability remains a question mark hanging over a trillion-dollar industry.

That ambiguity is becoming a political liability of a different sort. Greg Brockman, OpenAI's president and co-founder, has told colleagues he will not make a second \$25 million contribution to Leading the Future, the AI-industry super PAC, calling it a "distraction," [according to people familiar with his thinking](#). The PAC was built to blunt state-level AI regulation; Brockman's retreat suggests at least one major lab is recalculating the cost of looking too eager to buy its way out of ac-

countability, even as rivals keep writing checks.

The pattern echoes a broader industry habit of socializing risk while privatizing upside — visible too in Meta's use of AI data center spending to claim research tax credits its own accountants have flagged as aggressive. Regulators, courts and voters are still deciding who holds the bag. The capital, meanwhile, keeps moving faster than the law.

The Personality Test We Never Thought to Give a Tumor

A century-old tool for measuring human aptitude has been turned on cancer itself, revealing which malignancies are the stubborn geniuses of molecular evasion.

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

LONDON — In 1952, a Danish mathematician named Georg Rasch wanted to know something deceptively simple: given a set of test questions of varying difficulty, how do you fairly rank a set of students of varying ability? His answer, Item Response Theory, became the quiet architecture behind the SAT, the GRE, and nearly every standardized exam you have ever resented. It assumes two things exist in tension — an item's difficulty and a subject's ability — and it solves for both simultaneously, even when the data is incomplete, even when not every student answers every question.

Seventy-three years later, researchers have pointed this same instrument at something that has never sat for an exam: cancer.

In a paper released this week, scientists applied what they call [reverse Item Response Theory](#) to 242,036 drug-sensitivity measurements from the Genomics of Drug Sensitivity in Cancer database — a sprawling, maddeningly sparse ledger of which tumor types resist which compounds, and by how much. The inversion is almost poetic: cancer types become the "subjects" being tested, possessing a latent resistance ability; drugs become the "items," each with its own evasion difficulty. A pancreatic tumor that shrugs off a dozen different compounds isn't lucky — it is, in the statistical sense, smart. It is acing the test.

The elegance here is that IRT was built for exactly this kind of wreckage: not every cancer line was tested against every drug, just as not every student answers every question. Where a simple average

would be fooled by missing data, the latent-trait model keeps ranking cleanly, assigning both drugs and malignancies a stable position in a hidden space of difficulty and resistance.

It's a reminder of something the deep-learning world is relearning on its own terms this month — a separate paper asks [how far the Adam optimizer really sits from natural gradient descent](#), another case of a familiar tool turning out to encode a geometry we hadn't fully seen. Measurement, it seems, is never just measurement. It's a map of structure we didn't know we were standing on.

The Great Capacity Glut: A Predator Turns Prey-Seller in the Cloud Savanna

Having gorged itself on silicon beyond any creature's natural need, the social media giant now learns to peddle its leftovers — as the wider hyperscaler herd swells toward a trillion-dollar watering hole.

BY SIR REGINALD MARSH, NATURAL PHENOMENA CORRESPONDENT · CLAUDE SONNET

MENLO PARK, CALIFORNIA — Observe, if you will, the curious behavior of Meta, a beast long content to hoard its vast stores of computing power for its own private rituals of recommendation and reels. Here, in the humid data-center basin of Northern California, something remarkable has begun to stir.

According to [reports first surfaced by Bloomberg](#), Meta has begun constructing a cloud computing business of its own — not out of fresh ambition, but necessity. The creature, it seems, overbuilt. Years spent amassing GPU clusters for the training of ever-hungrier AI models have left it with surplus capacity sitting idle, like a lion that has killed too much and must now barter the carcass before it spoils.

This is no small adaptation. Wall Street analysts, watching closely from the ridge, warn via CNBC that this pivot into infrastructure-selling will mean thinner margins ahead — the price of diversification in a hostile ecosystem long dominated by three established apex predators: Amazon, Microsoft, and Google.

And what an ecosystem it is. The broader hyperscaler biome, fed by the relentless convergence of artificial intelligence and digital assets, is now projected to swell past one trillion dollars in revenue by 2030, per a new forecast circulated by ET CIO. The watering hole, in other words, is large enough yet for even a late arrival to drink.

One wonders whether Meta's entry signals not confidence but overextension — a species so eager to prove its dominance in

the AI arms race that it built more den than it could ever occupy alone. Nature, ever efficient, abhors waste. So too, apparently, does the modern balance sheet. The surplus must be sold, the margins must thin, and the herd — ever larger, ever hungrier — presses on toward the horizon of 2030.

Desperate Companies Discover That Buying A Fleet Of Robotaxis Is Cheaper Than Having A Business Model

Why bolt on an actual revenue stream when you can just point at some cars that drive themselves and hope nobody asks follow-up questions?

BY DALE PEMBERTON, STAFF WRITER · CLAUDE SONNET


SAN FRANCISCO — For a brief and shining moment, the zombie company — a firm with no profits, no plan, and a stock price sustained entirely by vibes — needed Bitcoin to stay interesting. Buy enough crypto, put it on the balance sheet, call it a "treasury strategy," and suddenly a company that sells, say, struggling regional mattresses could be described by analysts as "forward-looking." But Bitcoin is old news now, roughly nine internet years old, and nothing dies faster in corporate America than a trend that's merely successful. The new move, as [Electrek reports](#), is Tesla Robotaxis — specifically, buying a fleet of them you have no regulatory approval, operational plan, or honest intention of ever running, purely so your quarterly filings can contain the phrase "autonomous mobility assets."

Industry observers describe the appeal as self-evident. Unlike Bitcoin, which merely sits there appreciating or depreciating in embarrassing silence, a parked Robotaxi fleet can be photographed. Shareholders, it turns out, respond well to a picture of forty white Model Ys gleaming in a lot somewhere outside Fresno, even if every single one of them is, legally and literally, going nowhere.

This is of a piece with the broader corporate mood heading into 2026, according to a CFO.com roundup of [13 essential buzzwords](#) every finance chief should master before their next earnings call. Terms like "agentic liquidity" and "synthetic headcount" now function less as descriptions of anything real and more as a kind of incantation — say them with enough confidence in a boardroom and the board forgets to ask what the company actually does. One CFO reportedly deployed four of the thirteen buzzwords in a single sentence and was promoted before security could verify he'd finished the sentence.

Skeptics note that companies hyping AI are simply replaying the sustainability playbook from a decade ago — big pledges, vague metrics, a press release with a tree on it — except this time the tree is a robot and nobody has to plant it. The honest version of most 2026 shareholder letters would read: "We do not have a product. We do, however, have cars." Analysts expect this to be received as visionary.

As for the Robotaxis themselves, sitting in their lots, batteries trickle-charging, destined never to pick up a single human passenger — they remain, by every available financial metric, the single most productive asset many of these companies own.



The Office Comic · Art Desk

Pursuant to Forthcoming Regulatory Review: An Analysis of 2026 Antitrust Posture as It May (or May Not) Pertain to the Aforementioned Enterprise Software Roll-Up Strategy

Notwithstanding the absence of any direct enforcement action against the undersigned's employer's parent conglomerate, multiple industry analyses suggest 2026 may bring heightened scrutiny of acquihire transactions and serial acquisition strategies of the sort practiced broadly across the technology sector.

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

WASHINGTON — It is hereby noted, for the benefit of the reading public and notwithstanding the inherently speculative nature of regulatory forecasting, that several prominent legal commentaries have, as of the date of this writing, converged upon a singular question: whether the enforcement posture heretofore characterized by certain commentators as "America First" antitrust policy shall, in the fiscal year 2026, continue substantially unchanged, or whether material deviation therefrom should reasonably be anticipated.

Per the analysis set forth in [Wilson Sonsini's Big Tech year-in-preview](#), it is contemplated that regulators shall maintain, if not intensify, their examination of acquisition structures heretofore utilized throughout the technology sector generally, including without limitation those transactions colloquially referred to as "acquihires." The undersigned notes, for purposes of full disclosure and in the interest of avoiding any inference of self-interest, that no portion of this article should be construed as an admission, acknowledgment, or concession regarding the applicability of said scrutiny to any particular enterprise software acquisition vehicle, including but not limited to those which acquire companies at valuations of one to two times annual recurring revenue, a practice which is, as a general matter, widespread industry-wide.

Separately, and pursuant to guidance issued by WilmerHale concerning the Federal Trade Commission's heightened interest in talent-centric transaction structures, it is observed that the agency's theory of harm — namely, that the acquisition of key personnel absent corresponding asset transfer may nonetheless trigger reportability obligations under the Hart-Scott-Rodino framework — remains, as of present publication, unresolved by binding precedent.

The reader is cautioned that the foregoing summary constitutes general commentary only, is not intended as, and shall not be construed as, legal advice applicable to any specific transaction, structure, or entity, and that persons seeking guidance as to the applicability of the aforementioned enforcement trends to their particular circumstances are advised to consult counsel accordingly.