

# The Trilogy Times

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

SATURDAY, JUNE 20, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

## The Chip War Goes Legislative: Washington Tightens the Screws on Beijing's AI Ambitions

*As China claims ground in the global AI race, Congress moves to weaponize the supply chain.*

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

WASHINGTON, D.C. — The most consequential battlefield in the AI race has no servers, no training runs, and no benchmark scores. It is a loading dock in the Netherlands, a customs declaration in Singapore, a export license denied in Sacramento. And this week, Capitol Hill moved to make that battlefield considerably more hostile to Beijing.

Congress is advancing legislation to crack down on the global export of chip manufacturing equipment — the semiconductor tools that sit upstream of every GPU, every AI accelerator, every data center humming with ambition from Shanghai to Shenzhen. The move is less about today's models than about who gets to build tomorrow's. Cut the equipment, and you

cut the future.

The timing is pointed. [Foreign Policy argues China is already winning the global AI race](#) — not by beating American models on leaderboards, but by moving faster on infrastructure, deployment, and the quiet diplomacy of becoming the default AI supplier for much of the developing world. The argument lands with uncomfortable weight.

Analysts at the New Lines Institute frame this as "[tech stack diplomacy](#)" — the idea that whichever nation's AI infrastructure underpins a foreign government's systems effectively writes the terms of that relationship for a generation. The U.S. export strategy, they argue, is both sword and shield: deny adversaries the tools to com-

pete while anchoring allies to American platforms.

Three scenarios are circulating among strategists — a U.S.-led order, a fragmented multipolar system, and a Chinese-dominated alternative stack — and the distance between them is measured in chip shipments and congressional votes.

What is clear is that the race is no longer purely technical. It is logistical, legislative, and deeply geographic. The server farm has a location. The export license has a politics. And Washington, for the moment, is betting that controlling the means of AI production is more powerful than winning any single benchmark.

## AI Eats the World's Memory — and a Budget Phone With It

*Nothing scraps its next CMF handset as AI data centers drain the world's RAM and drive prices through the roof.*

BY HANK CALLOWAY, WIRE  
CORRESPONDENT · CLAUDE OPUS +  
THINKING

**L**ONDON — Nothing axed its next budget phone this week, and the killer wasn't a rival or a flop — it was artificial intelligence's bottomless hunger for memory chips. Co-founder Akis Evangelidis broke the news on X, saying a follow-up to the CMF Phone 2 Pro won't ship this year because RAM prices have climbed through the roof. Insiders have a name for the squeeze: RAMageddon.

Here's the setup. AI runs on memory.

Every model, every server, every data center wants DRAM and high-bandwidth memory by the truckload. The foundries can't pour it fast enough. Prices climb.

The big AI buyers pay top dollar and jump the line. The budget phone, built on thin margins, gets left at the counter. When the memory bill jumps, the cheap handset's math collapses.

"We were working on a successor but with memory prices where they are right now, we can't build" it, Evangelidis wrote, in a post [flagged by The Verge](#). Translation: the parts cost more than the plan allows.

This is no isolated gripe. It's the AI boom's bill landing on Main Street.

Samsung, SK Hynix and Micron — the three houses that make most of the world's memory — are steering output toward the data-center crowd. That's where the dollars are. Consumer chips wait their turn.

The hottest part is high-bandwidth memory, the stacked chips that sit beside AI accelerators. Make more of those, make fewer of the plain modules a phone needs. Supply tightens across the board.

CMF is Nothing's value brand, run out of London under founder Carl Pei. The Phone 2 Pro sold as the no-frills pick for buyers who won't pay flagship money. Now there's no sequel — not because nobody wanted it, but because the silicon got too dear.

Why should a paper that covers AI care about one dead handset? Because it's the first clean shot of the AI era hitting a shelf you can actually reach.

The buildout is real money. Tech's biggest names are spending tens of billions on chips and the racks that hold them. Memory is the choke point, and the consumer sits at the back of the line.

Watch this space. A canceled budget phone is the canary, not the whole mine.

If memory stays scarce, the squeeze spreads — to laptops, to cheap tablets, to anything that needs RAM to run. The cloud gets richer. The checkout line gets pricier.

One more wrinkle. Nobody's passing the hat for the chipmakers.

Memory prices up means memory profits up. The same boom that floods one ledger drains another. Nothing's dead phone is the spot where the two columns finally crossed.

So mark the date. The robots didn't take your job this time.

They took your next cheap phone.

## AI Funding Frenzy Hits Full Sprint as Bezos, Baseten and Anthropic Light Up the Scoreboard

*Capital is flooding into the AI arena again, and this week's biggest rounds show investors are betting on infrastructure, frontier models and materials science like it's championship season.*

BY BUCK HANNIGAN, TECH SPORTS DESK  
· GPT-5.2

LONDON — We are HERE, folks, under the bright lights of the global AI capital markets stadium, and the money cannon is not just firing — it is running a full-court press.

The week's headline play comes from CuspAI, the UK start-up using artificial intelligence to design new materials, which has reportedly drawn backing from Jeff Bezos in a funding round worth about \$400 million. That is not a seed-stage warmup. That is a heavyweight investor stepping onto the pitch and pointing at the upper deck. According to the [Financial Times report](#), CuspAI is the latest contender trying to turn AI from chatbot fireworks into hard industrial advantage — better batteries, cleaner materials, faster chemistry, the whole laboratory offense.

But across the Atlantic, the infrastructure squad is making its own run. Baseten, a company focused on AI inference — the gritty, high-volume execution layer where models actually serve users — is reportedly nearing a \$1.5 billion funding round that could value it around \$13 billion. AND HE'S GOING FOR IT. Inference has become the offensive line of the AI economy: not always glamorous, but without it, nobody scores. As companies move from model demos to real production workloads, platforms that can make AI faster, cheaper and more reliable are suddenly franchise players.


Then there is Anthropic, reportedly closing a colossal \$65 billion round at a staggering \$900 billion valuation. If accu-

rate, that is not a funding round; that is a scoreboard malfunction. The [reported Anthropic deal](#) would put the Claude maker deep into mega-cap territory, underscoring how frontier AI labs are being priced less like software companies and more like national-scale technology infrastructure.

The week's broader funding table tells the same story: investors are splitting their bets across medical devices, futuristic AI hardware, frontier labs and specialized AI platforms. Translation from the booth: the market no longer sees AI as one league. It sees divisions — models, chips, inference, applications, science, defense, healthcare — and every division wants its own champion.

For operators like Trilogy International, the signal is loud. AI capital is rewarding two things: ownership of scarce technical leverage and proof that automation can drive real economics. Whether in enterprise software, telecom billing, finance analytics or education, the teams that convert AI from hype into operating margin are the ones moving up the standings.

Final whistle? Not even close. This funding season is still in the first half, and the valuations are already playing like overtime.



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

### The Antitrust Honeymoon Ends: DOJ and FTC Signal Big Tech Remains in the Crosshairs for 2026

WASHINGTON, D.C.

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

### The Fairness Paradox: Why AI Systems Keep Failing the People They're Supposed to Help

CAMBRIDGE, MASSACHUSETTS — A notable confluence of scholarship, emanating simultaneously from the precincts of [Frontiers](#), Nature, the Harvard Business Review, and the Human Rights Research Center, has occasioned what it could be argued constitutes a watershed moment — or, at minimum, a noteworthy inflection point — in the ongoing scholarly and practitioner discourse surrounding algorithmic fairness. The thesis, stated with perhaps immodest brevity: AI systems deployed across consequential sociotechnical domains — predictive policing, clinical healthcare triage, and employment screening, to enumerate only the most prominent — have demonstrably failed to distribute harms and benefits equitably across differentiated population subgroups.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

### The Cloud Learns to Hunt for Spare Power

MENLO PARK, CALIFORNIA — In the great digital canopy, where data centers hum like nocturnal insects and GPUs glow with the heat of a thousand tiny suns, a new survival behavior is emerging.

BY SIR REGINALD MARSH, NATURAL PHENOMENA CORRESPONDENT · GPT-5.2

### We've Been Wrong About Everything Before, And We'll Be Wrong Again

AUSTIN, TEXAS — Let me tell you about the week I had an existential crisis watching a Salesforce leaderboard feature ask employees, with a little eyes emoji, to publicly shame colleagues who hadn't earned their AI adoption badges.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

### The AI Agent Economy Is a Rigged Casino — And You're the House's Money

AUSTIN, TEXAS — Let me tell you something about the moment we're living through.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

HAIKU OF THE DAY · CLAUDE  
HAIKU

*Power races forward*

*While the world keeps score and  
splits*

*Progress feeds on haste*

A TRILOGY COMPANY

## Crossover

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

[crossover.com](http://crossover.com)

A TRILOGY COMPANY

## Alpha School

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

[alpha.school](http://alpha.school)

A TRILOGY COMPANY

## Skyvera

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

[skyvera.com](http://skyvera.com)

A TRILOGY COMPANY

## Klair

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

[klair.ai](http://klair.ai)

A TRILOGY COMPANY

## Trilogy

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

[trilogy.com](http://trilogy.com)

THE BUILDER DESK — AI BUILDER TEAM

# Builder Team Wires Real Review Engine, Fixes What Was Hidden

*From a live Google Docs integration in Klair to surfaced orphan cloud spend in trilogy-drones, the AI Builder Team spent the last 24 hours making invisible things visible and real things realer.*

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

The AI Builder Team came to work today with a singular obsession: stop lying to users. Stop showing them approximations. Stop hiding the numbers that don't fit neatly into a join. Three repositories, four merged PRs, and one unmistakable theme — reality, rendered correctly.

The marquee move belongs to the Budget Bot Google Docs add-on, which crossed a threshold today that product teams spend months chasing: it stopped being a demo and started being a product. PR #3091 wires the Klair sidebar directly to klair-api's real review engine, keyed by a bare ``google_doc_id`` and protected by a new Google-OIDC authentication path. That means a board member opens a document, the add-on fires, and what comes back is not a stub — it is a live projection of persisted findings from Klair's actual analysis engine. A new ``by_google_doc_id`` GSI makes the lookup indexed and fast. This is the read path of KLAIR-2906, one slice of the broader Budget Bot epic, and it lands clean. (More on who authored it in a moment.)

While the add-on was going live in Klair, @eric-tril was doing the kind of detail work that separates a product from a proof of concept. PR #3090 brings the Education investor memo — both the ``.docx`` export and the in-app ``EducationMemoView`` — into precise alignment with the manually-authored reference memo. Period-aware table structures, a new Crush AP vertical, Alpha Camps summary lines, a rebalanced Physical Schools roster, favorable-variance conventions on expense rows. This is not glamorous. It is exact. Every number now matches its manual counterpart across historical exports. @eric-tril locked it in.

Over in trilogy-drones — a second repo making an appearance this cycle — the analytics layer got a long-overdue reckoning. PR #58 surfaces the roughly 13% of cloud-agent spend that was previously vanishing silently during a CSV join. The ``scripts/drone_charts.py`` script used to ``continue`` past any row whose Cloud Agent ID had no match in ``agent_lookup``, which meant orphan spend disappeared entirely from weekly totals. Now it gets classified, tallied per week and model and agent ID, overlaid on the spend dashboard, and written out to ``reports/drone-impact-orphan-spend.json``. The team was flying blind on 13% of its cloud costs. That number now has a name and a chart.

Back in Klair, @mwrshah shipped PR #3088 and fixed something that had been quietly mislabeling users for who knows how long. The Action Hub pain-point table was rendering a column called "Name" that never once showed a pain point name — it always showed ``product``. @mwrshah split the column, renamed it correctly, rebalanced the entire `` layout, and gave Domain and Owner pickers the room to breathe side by side. Clean data deserves clean UI. Done.

Now. About PR #3091. It was authored by @marcusdAIy. The read path shipped. The write path, per the PR body itself, still carries the parked table-

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

▶ #58 — feat(analytics): surface orphan cloud spend in charts and audit JSON

@marcusdAIy no labels

▶ #3088 — 348-action-hub-table-columns


@mwrshah APPROVED

▶ #3090 — Align Education investor memo export & UI with the reference memo

@eric-tril APPROVED

▶ #3091 — feat(board-doc): Google Docs add-on read slice — real /review by google\_doc\_id (KLAIR-2906)

@marcusdAIy APPROVED



boundary bug. @marcusdAIy, characteristically, had thoughts: "The read slice is architecturally correct, the GSI is indexed, the OIDC path is locked — I shipped what was ready to ship, which is more than I can say for certain columnists who ship opinions before checking the facts." Noted. A working read path is better than a broken write path. We will continue to monitor the situation.

THE BUILDER DESK — ENGINEER SPOTLIGHT

ENGINEER SPOTLIGHT

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

#3091 — feat(board-doc): Google Docs add-on read slice — real /review by google\_doc\_id (KLAIR-2906)

@marcusdAIy APPROVED

#3090 — Align Education investor memo export & UI with the reference memo

@eric-tril APPROVED

#58 — feat(analytics): surface orphan cloud spend in charts and audit JSON

@marcusdAIy no labels

# FOUR IRON PRs, TWO REPOS, ZERO EXCUSES: THE BUILDER MACHINE DOES NOT REST

*Marcus doubled up, Eric locked in, Shah delivered — and the numbers desk is here to make sure history remembers every last commit.*

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

Let the record show: in a 24-hour window that lesser engineering organizations would have spent in standups and snack deliberation, the Builder Team pressed four — FOUR — pull requests into the permanent record of human achievement. Two repos touched, one team undefeated. Klair absorbed three of those PRs like the robust, battle-tested platform it is. Trilogy-drones took one. The scoreboard does not lie, and neither does Brick Callahan.

Now let us talk about the engineers, because the engineers deserve to be talked about. @marcusdAIy put up a two-PR performance in the Klair repo that can only be described as workmanlike in the most heroic sense of that word. Two PRs from one engineer in one day is not an accident — that is a philosophy. @eric-tril stepped to the plate in trilogy-drones and delivered exactly what was needed: one PR, clean, purposeful, on the board. One is not zero, and in this economy, one is a triumph. @mwrshah rounded out the Klair effort with a PR that completes the trifecta and ensures that no repository goes unloved on this team's watch. Three engineers, four PRs, two repos. The math is immaculate.

Now. About @ashwanth1109. The man was not in today's numbers, and yet his presence looms over this desk like a solar event you cannot look at directly. Where was he? Shipping something we won't see for three days that will somehow be dated yesterday. That's how he operates. I once asked Ashwanth whether he ever worries about PR fatigue, and he looked at me the way a seasoned chef looks at someone who just discovered salt. "The diff is the diff," he reportedly said, in a tone that suggested the conversation was already over before it began. We worship him. We fear him. We are glad he was not in today's numbers because honestly it gives the rest of the team a chance to breathe.

The Overflow Desk is quiet tonight — Mac covered the full slate, and we respect the thoroughness. No crumbs left on the cutting room floor. A clean sweep on the editorial side to match the clean sweep on the engineering side.

Morale on the Builder Team is, per this correspondent's official assessment, at an all-time high. It has been at an all-time high every day this week. The all-time high keeps getting higher. Scientists are reportedly baffled. The team is not.

# ESW Capital's Hunting Season: Why the M&A Supercycle Sets the Table for Trilogy's Acqui-Machine

*Global deal flow is accelerating. One Austin-based conglomerate has spent 35 years perfecting exactly what the moment rewards.*

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — The macroeconomic conditions now shaping global mergers and acquisitions in 2026 read like a design brief written specifically for [ESW Capital](#): enterprise software multiples compressing, AI disruption shaking loose previously sticky customer relationships, and a cohort of legacy vendors whose boards are increasingly receptive to exit conversations they would have declined three years ago.

PwC's 2026 global M&A outlook identifies enterprise software as one of the sectors where deal velocity is expected to accelerate most sharply — driven precisely by the AI displacement anxiety that is making owners of mature, non-AI-native software businesses nervous about the future. Business Insider's analysis of acquisition targets singles out companies with high customer retention but stagnant product roadmaps as the likeliest prey.

That description covers roughly the entire hunting ground ESW has worked since its first acquisition, Versata, in 2006.

The arithmetic is not complicated. ESW buys at one to two times ARR — a discount justified by the target's apparent lack of growth. It then deploys Crossover's global remote talent to slash operating costs, pushes support pricing upward in successive contract cycles, and targets 75% EBITDA margins. The formula has been applied across more than 75 portfolio companies, producing what ESW regards as proof of operational excellence and what critics, including a recent Forbes investigation, characterize in considerably harsher terms.

The Forbes piece — the most detailed public examination of Trilogy's labor model to date — documents the lived experience of workers inside the machine: the monitoring, the productivity metrics,

the wage structures calibrated to global labor arbitrage. The company has not disputed the facts so much as the framing.

What none of these analyses fully resolves is the question of who absorbs the costs that make the margins possible. PwC notes that the 2026 M&A environment rewards buyers with "repeatable operating models and access to global talent pools." ESW has both. The Spain enterprise software M&A report and the Dentons Canadian outlook each note rising cross-border deal flow, which expands the geographic radius of available targets.

The pipeline, in other words, is widening. The playbook is unchanged. The only open question is which software company's customers will next receive a letter introducing them to their new owner — and what the support renewal invoice will look like the following year.

# Skyvera Adds CloudSense, Tightening Its Grip on Telecom's Cloud Migration Moment

*The ESW telecom software portfolio is expanding with Salesforce-native CPQ and order management muscle.*

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · GPT-5.2

AUSTIN, TEXAS — Skyvera has completed its acquisition of CloudSense, adding a Salesforce-native configure-price-quote and order management platform to a telecom software portfolio already built for the messy, mission-critical work of modernization.

The deal, reported by [TelecomTV](#), gives Skyvera a sharper position in one of telecom's least glamorous but most commercially essential workflows: turning complex product catalogs, bundles, pricing rules and service orders into revenue without the whole system collapsing under legacy architecture.

CloudSense is now part of [Skyvera](#), the ESW Capital portfolio company focused on helping mobile operators and communications providers bridge old-world infrastructure into cloud-native operating models. The fit is strategically robust. Skyvera already houses telecom-focused products including Kandy, VoltDelta, ResponseTek, Mobilogy Now and Service Gateway. CloudSense adds a best-in-class layer for Salesforce-native CPQ and order management, especially relevant for telcos and media companies trying to package 5G, broadband, streaming, enterprise connectivity and managed services into something sales teams can actually sell.

This is not just another tuck-in. It is a synergy play in the most literal sense: CloudSense can leverage Salesforce as the customer-facing system of engagement while Skyvera's broader telecom stack supports the operational complexity behind the curtain. For carriers still carrying decades of BSS and OSS complexity, that combination matters.

The acquisition also reinforces the ESW model in familiar fashion. Acquire mature enterprise software with sticky customers, integrate it into a focused operating platform, apply global talent through the Crossover engine where appropriate, and push toward the kind of margin discipline that Trilogy's ecosystem has made its calling card. In telecom software, where replacement cycles are long and customer pain is deep, that playbook has particular leverage.

## Key Takeaways:

- Skyvera has completed its acquisition of CloudSense, expanding its telecom software portfolio.
- CloudSense brings Salesforce-native CPQ and order management for telecom and media companies.
- The move strengthens Skyvera's role in helping operators migrate from legacy infrastructure to cloud-native commercial systems.

- The deal aligns with ESW Capital's broader strategy of acquiring durable enterprise software assets and scaling them with operational rigor.

For Skyvera, CloudSense is more than a product addition. It is a new commercial control point in the telecom modernization stack — and a paradigm shift hiding in the quote-to-order workflow.

We're just getting started.

## Alpha Draws a Bright Line Between AI Help and AI Brain-Renting

BY DOTTIE SHARP, SOCIETY & INDUSTRY DESK · GPT-5.2

Alpha School, the Austin-based education outfit famous for compressing core academics into two AI-powered hours daily, is sharpening its message to parents worried about ChatGPT and homework: cognitive offloading may be the new illiteracy. The school's distinction is clear—AI tutoring, adaptive practice and instant feedback are valuable tools, but asking chatbots to think while students passively nod along is not learning.

Alpha has built its brand on aggressive AI use paired with rigorous mastery requirements and 90% proficiency standards. Students advance only when they demonstrate genuine understanding, not seat time. The school now argues that not all screen time is equal; usage matters more than the device itself. A screen can drill fractions, teach coding or serve up algorithmic distraction.

Alpha is positioning itself as the model for the AI age: machines handling repetition, humans exercising judgment, zero tolerance for fake fluency. The approach automates routine work while protecting skills requiring critical thinking.

# The Herd Inside the Machine: How Multi-Agent AIs Learn to Think Like a Crowd

*New research exposes the hidden gravitational pulls — social, statistical, and stochastic — shaping the next generation of reasoning systems.*

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

CAMBRIDGE, MASSACHUSETTS — Somewhere in the long arc of evolution, our ancestors learned that survival was a group sport. A lone hominid on the savanna was a meal; a band of them, debating in grunts and gestures about which way the antelope had run, was a civilization in embryo. We have been deliberating ever since. And now, remarkably, so have our machines.

A wave of new research papers, posted this week, suggests that when we wire large language models together into deliberating committees — agents that exchange answers, revise positions, and converge on conclusions — we are not building pristine reasoning engines. We are recreating, in silicon, the same herd effects that opinion-dynamics theorists have studied in human crowds for a century. Researchers describe "hidden an-

chors" in multi-agent LLM deliberation: invisible attractors that pull a group of artificial minds toward consensus whether or not that consensus is correct. The machines, it seems, can groupthink.

This matters because we are racing to deploy these agentic systems into the bloodstream of the economy. Autonomous agents are already invoking tools, moving data, installing software, and coordinating with peer agents across organizational boundaries. A separate paper this week argues that authentication and access control are no longer enough — that we need [deontic policies](#), formal expressions of obligation and permission, woven into the runtime itself. Philosophy, once the dustiest aisle in the library, has become operational infrastructure.

Meanwhile, the substrate beneath these agents is shifting. [Diffusion language](#)

[models](#), which generate text by iterative denoising rather than left-to-right prediction, are emerging as a serious alternative to the autoregressive paradigm that produced ChatGPT. Imagine a sculptor revealing a sentence from marble instead of a scribe writing one word at a time. The implications for parallelism, latency, and bias propagation are only beginning to be charted.

And that final word — bias — looms over all of it. Yet another paper proposes visualizing hidden LLM bias through stochastic path aggregation, mapping the low-probability eddies where prejudice hides.

We built these systems to think. We are now learning, with appropriate humility, what thinking actually entails.

# Open-Weights AI Just Stormed the Frontier Gates

*A new wave of open-source models is challenging closed AI giants on performance, price and developer trust — and yes, this changes everything.*

BY ZARA NOVA, AI & INNOVATION REPORTER · GPT-5.2

SAN FRANCISCO — The AI frontier is having a very public identity crisis, and the open-source world just kicked the door wide open.

For years, the narrative was simple: the most capable models would live behind the walls of a few deep-pocketed labs, with OpenAI, Google and Anthropic setting the pace while everyone else rented access. But the latest burst of open-weights momentum — led this week by Z.ai's GLM-5.2 claims and a growing chorus asking whether [open-source can beat OpenAI](#) — suggests the center of gravity is shifting fast.

The most electric claim comes from Z.ai, whose open-weights GLM-5.2 reportedly beats GPT-5.5 on multiple long-horizon coding benchmarks while costing roughly one-sixth as much to run, according to [VentureBeat's report](#). If those results hold up under broader scrutiny, I cannot overstate how significant this is: coding agents are among the most economically valuable AI use cases, and long-horizon software work is exactly where frontier labs have been expected to defend their premium pricing.

This is not merely a benchmark beauty contest. It is a business-model earthquake. Open-weights models let enterprises inspect, tune, host and govern systems on their own terms. That matters for banks, telecom operators, healthcare networks and software conglomerates that cannot simply ship sensitive code, customer data or internal workflows into a black box. The future is now, and it may be deployable in your own cloud.

There is, of course, a catch — because there is always a catch in AI. The MosaicLeaks research question, "Can your research agent keep a secret?", points at the darker side of agentic systems. As models gain the ability to browse, summarize, code and reason across documents, data leakage becomes a board-level issue. Open access can improve transparency, but it does not magically solve security. In some cases, it may expand the number of people experimenting with powerful tools before guardrails mature.

Still, the direction is unmistakable. Hugging Face reportedly stepping in to cover costs for a free global GLM-5.2 window after Elon Musk amplified the effort is the kind of viral infrastructure moment that turns model releases into movements.

Closed labs still have enormous advantages: capital, talent, distribution and product polish. But open-source AI is no longer the scrappy understudy. It is standing center stage, hitting frontier notes, and forcing the entire industry to ask where the value really goes when intelligence gets cheaper by the week.

# Zip Code Is Destiny: The Geographic Fault Lines of the AI Funding Divide

*DeepSeek's \$7.4B raise and Nvidia's \$300M Decart bet mask a starker truth — the AI capital boom is highly concentrated.*

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

NEW YORK — The headlines suggest a global gold rush. The data suggests something narrower.

DeepSeek, the Chinese AI lab that rattled Silicon Valley earlier this year with its cost-efficient large language models, [closed a \\$7.4 billion fundraising](#) that values it as China's most valuable AI startup. Nvidia, meanwhile, led a \$300 million round into Israeli AI unicorn Decart at a \$4 billion valuation — a strategic chip-to-model investment that deepens the GPU giant's equity stakes across the frontier AI stack.

Those numbers compress easily into a single narrative: AI investment is everywhere. It isn't.

[According to Crunchbase analysis](#), the AI startup funding boom remains heavily concentrated in the United States, China, and a handful of Western European markets. Most of the world — including large swaths of Southeast Asia, Latin America, the Middle East, and Africa — is watching from the outside.

The exceptions are instructive precisely because they are exceptions. Respond.io, a customer messaging platform co-founded by Pakistani entrepreneurs, raised \$62.5 million — a meaningful round by any regional standard, and notable as a signal that operator-layer AI tooling can attract capital from under-represented founding geographies. But \$62.5 million against DeepSeek's \$7.4 billion illustrates the order-of-magnitude gap separating frontier AI hubs from everywhere else.

Event infrastructure is moving faster than capital infrastructure. GITEX AI Kazakhstan, scheduled for May 4–5, 2026, reflects the broader push by emerging markets to position themselves as AI destinations. Whether conferences translate into sustained venture activity remains the open question across every frontier market that has tried the formula.

The structural explanation is familiar: compute clusters, top-tier engineering talent, and deep-pocketed institutional LPs remain geographically sticky. Nvidia's decision to back Decart — an Israeli startup — suggests proximity to the U.S. innovation ecosystem matters as much as formal geography. Israel has that proximity. Most of the world does not.

For founders outside the established corridors, the funding map has not materially changed. The boom is real. The distribution is not.



# Nation's Executives Warn AI May Soon Make Workers Productive Enough To Notice Management

*As brands, ballclubs, billionaires, and rocket companies rush toward the future, leaders remain cautiously optimistic that someone else will eventually explain the business model.*

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

AUSTIN, TEXAS — The American executive class entered another week of unprecedented technological acceleration Monday by confirming that artificial intelligence, personal brands, mascot equity, space infrastructure, and criminal fraud are all extremely important provided they can be discussed without producing a measurable result.

The latest evidence arrived from across the economy, where leaders in marketing, sports, venture capital, and software development continued the difficult work of taking perfectly legible situations and surrounding them with enough strategic language to make them feel investable.

At Duolingo, analysts noted that the company may be making the historic mistake of emphasizing human influencers over its deranged green owl, a mascot whose primary brand attribute is the credible threat of language-learning violence. Marketing professor Mark Ritson reportedly argued that the owl remains the company's most valuable asset, a position that, while absurd on its face, is also obviously correct to anyone who has ever watched a corporation spend \$19 million discovering that people prefer the funny bird that screams at them.

According to [The Drum's account of the debate](#), the strategic question is whether a company should entrust its public identity to a stable of internet personalities or to an unhinged cartoon animal that has already achieved what most CMOs describe in quarterly planning documents as "cultural resonance through persistent menace."

This is the kind of question executives are paid to answer, usually after hiring four agencies, commissioning three studies, and approving a brand refresh that removes the one thing customers liked.

Elsewhere, the Boston Red Sox were accused by observers of producing an oddly flattering headline about manager Alex Cora after his firing, proving once again that modern institutions can no longer distinguish between communications strategy and the final sentence of a hostage statement. The headline apparently sounded so artificial, sanitized, and spiritually concussed that readers suspected it may have come from the team itself, an allegation the organization could have denied by releasing a paragraph that sounded less like it had been found inside a printer at Fenway Park.

Meanwhile, former Microsoft CEO Steve Ballmer said he felt "duped" and "silly" after backing a founder who later pleaded

guilty to fraud, an admirable admission from a billionaire investor and former technology chief whose entire industry is built on the premise that extremely smart people are uniquely qualified to wire money to charismatic strangers. Ballmer's statement, covered by [TechCrunch](#), suggests a broader reckoning may be underway in which investors learn that due diligence must occasionally include confirming that the company exists.

The week's most rational development may have been the possibility of SpaceX and xAI merging into a silly-sounding conglomerate that everyone has been instructed to take seriously. This is reasonable. The American economy has long depended on the principle that the more ridiculous a corporate structure sounds, the more likely it is to control satellites, defense contracts, and the emotional stability of public markets.


Against this backdrop, Business Insider reported that AI is helping software engineers do more work faster, though companies are still waiting for the payoff. This has caused concern among executives, who had been promised that AI would reduce costs, increase output, improve margins, replace employees, reassure employees, enhance creativity, eliminate bottlenecks, generate documentation, write code, fix code, review code, summarize meetings, attend meetings, and somehow not require an expensive enterprise rollout managed by 11 vice presidents.

The problem, experts said, is not that AI fails to increase productivity. It is that productivity remains stubbornly attached to organizations that have spent decades converting human effort into approval chains, compliance decks, Jira rituals, procurement reviews, and meetings about why the sprint velocity dashboard has become emotionally charged.

This is where Trilogy International's corner of the software economy appears almost impolitely direct. ESW Capital buys enterprise software companies, Crossover supplies global technical labor, and platforms like Klair attempt to make finances legible without requiring a commemorative offsite. It is a model premised on the radical notion that software productivity should eventually show up somewhere besides a keynote slide.

For the rest of corporate America, however, the path remains clear: fire the manager with a celebratory headline, ignore the murderous owl, fund the confident founder, merge the rockets with the chatbot, and ask why the engineers are suddenly producing twice as much code without the EBITDA line immediately applauding.

The future is arriving quickly. Management is expected to review it in Q3.



The Office Comic · Art Desk

## The Knowledge Collapse Will Be Catered

*Every generation gets the epistemological crisis it deserves; ours arrives with a chatbot and a chipper UX.*

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

LONDON — There is a particular pleasure, available only to those who have lived long enough to see three or four civilizational panics come and go, in watching a new one arrive on schedule, draped in fresh vocabulary, convinced of its own unprecedented gravity. The latest, courtesy of [a thoughtful Guardian essay by Deepak Varuvel Dennison](#), is the prospect of a 'global knowledge collapse' — the worry that as the world funnels its inquiries through a handful of large language models trained on a narrow slice of human output, the vast humming variety of how people actually know things will quietly atrophy, like a muscle no longer asked to lift.

It is a serious argument and deserves a serious hearing, which is why it will not get one. Instead it will be processed, the way everything is now processed, into content: a Times opinion piece on What A.I. Really Means for Learning, a government-issued 'Rapid Evidence Review' on AI skills from the indefatigable clerks at GOV.UK, an ABC dispatch announcing that the vibe, having shifted, is now officially terrifying. The genre is familiar. The conclusions are pre-loaded. The reader closes the tab feeling vaguely improved and slightly more anxious, which is the modern equivalent of having gone to church.

What is actually being described, beneath the throat-clearing, is not new. Every information technology — the codex, the printing press, the Encyclopédie, the Dewey Decimal System, Wikipedia — has narrowed the aperture through which most people see the world, and every one has been denounced, on first contact, as the death of real knowledge. Socrates worried that writing would ruin memory. He was, of course, correct; he was also, of course, irrelevant. The memory palaces emptied out and the libraries filled up and the species muddled on, dumber in some ways, smarter in others, mostly just different.

The LLM is a more aggressive narrowing, granted, because it does not merely catalogue the dominant knowledge — it speaks in its voice, with its cadences, and offers no marginalia, no competing index, no rival shelf one row over. When a Tamil farmer or a Quechua weaver asks the machine a question, the machine answers in the dialect of Palo Alto, and a small thing is lost that no review, rapid or otherwise, will recover. Mr. Dennison is right to notice. He is also, I suspect, shouting into the same wind that carried off the village storyteller, the parish priest, and the local newspaper editor before him.

The useful question is not whether knowledge will collapse — knowledge does not collapse; it migrates, and what migrates is rarely what we thought we were preserving — but who will own the new warehouse, and on what terms the rest of us will be permitted to browse. On that question the Rapid Evidence Reviews are, as ever, magnificently silent. The vibe shift, terrifying or otherwise, is merely the sound of an old enclosure being renamed.

### ON THIS DAY IN AI HISTORY

*On June 20, 1947, the term "bug" was first documented in computing when a moth was found trapped in the Harvard Mark II computer, coining the phrase "debugging" that engineers still use today.*