

# The Trilogy Times

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

SATURDAY, JUNE 06, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

## Made-in-China AI Knocks Valley Sideways

*DeepSeek's bargain-priced model matches American giants — built without the top-shelf chips Washington bans from export.*

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

**S**AN FRANCISCO — A Chinese shop named DeepSeek has Silicon Valley reeling this week, claiming it built a champion artificial intelligence model on cut-rate chips for pennies on the dollar. The Hangzhou upstart turned its work loose over the weekend. The Valley's biggest names spent Monday picking their jaws off the floor.

"Amazing and impressive," they're [calling it](#). The model goes punch-for-punch with the heavy hitters from OpenAI and Anthropic on the benchmarks that matter. DeepSeek says it did the job without Nvidia's top-shelf silicon — the very chips Washington bans from export to the People's Republic.

That news ought to keep tech bosses up nights. The whole American AI play rests on a single bet: money plus the fastest chips equals dominance. The math has been simple — pour billions into training runs, snap up every Nvidia H100 you can find, win.

DeepSeek just shoved a wrench in the works and laughed about the price tag.

Trading desks spent the morning chewing over what [cheaper AI means for the chip kings](#) who rode the wave to the moon. Nvidia's run depends on every lab on Earth needing more silicon, not less.

The cost figure making the rounds: a sliver of what OpenAI or Anthropic spend on a flagship training run. If the number holds up, the floor under American AI valuations just dropped through the basement. If it doesn't — somebody's books need a closer look.

The export controls cut both ways. Washington meant to slow Peking's AI push by choking off advanced chips. DeepSeek's answer says Peking found another road, and the road runs through smarter software, not bigger hardware.

Engineers in Mountain View and Palo Alto were poring over the technical papers by sunup, hunting for the trick. Some say the gains came from clever training methods and architecture choices. Others smell smoke and want to see the fire before they cheer.

Meanwhile in Menlo Park, Reid Hoffman is changing seats. The LinkedIn co-founder is stepping off Microsoft's board after a decade and hauling \$24.6 million into a new outfit called Manas AI. His partner on the venture is Siddhartha Mukherjee, the doctor who wrote "The Emperor of All Maladies."

The pitch is AI for cancer drug discovery. Hoffman calls it "founder mode." Microsoft loses a director who watched the company more than triple in value under Satya Nadella, riding the OpenAI partnership to the top of the heap.

Connect the dots and the picture sharpens. The compute arms race is shifting under the Valley's feet. DeepSeek showed you don't need the biggest pile of chips to play the game.

Hoffman's bet says the real prize isn't chatbots anymore — it's cures. The Valley spent two years certain that whoever bought the most Nvidia silicon would win the AI war. This week, the war got more complicated.

## White House AI Blueprint Seeks Federal Supremacy Over State Laws, Minimal Regulatory Burden

*The Trump administration's new AI policy framework asks Congress to preempt a patchwork of state rules — while the Supreme Court quietly closes the door on AI authorship claims.*

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

**W**ASHINGTON, D.C. — Pursuant to the issuance of a comprehensive legislative blueprint by the White House, hereinafter referred to as "the Framework," the executive branch has formally communicated to the United States Congress its position that artificial intelligence regulation shall, to the maximum extent practicable, be characterized by a light regulatory touch, notwithstanding the existence of numerous and potentially conflicting state-level legislative efforts currently in various stages of enactment across the several states.

The aforementioned Framework, as analyzed by legal practitioners at [Crowell & Moring LLP](#), is understood to contain provisions calling for the preemption of state artificial intelligence laws, thereby establishing a unified federal regulatory regime in lieu of the aforementioned patchwork of state-level enactments. It is further represented that the Framework includes provisions directed toward the protection of children in connection with artificial intelligence systems, the precise scope and enforceability of which remain, at the time of publication, subject to further legislative deliberation.

Notwithstanding the foregoing, it shall be noted that the Framework does not, in and of itself, constitute binding law, and that the realization of its stated objectives is contingent upon affirmative action by Congress, the likelihood and timeline of which cannot, at this juncture, be represented with any degree of certainty.

In a separate but related development of considerable legal significance, the Supreme Court of the United States has declined to hear a case concerning [AI authorship and inventorship](#), thereby permitting to stand lower court rulings that artificial intelligence systems are not, under existing statutory frameworks, recognized as authors or inventors for purposes of intellectual property protection. The implications of the aforementioned denial of certiorari for enterprises engaged in AI-assisted creative and inventive activities are, it is submitted, substantial and as yet incompletely understood by affected parties.

## AI Funding Market Turns Into a Full-Contact Spending League

*DriveNets, Anthropic and Commure all put fresh points on the board as investors keep blitzing anything that can feed, route or deploy artificial intelligence.*

BY BUCK HANNIGAN, TECH SPORTS DESK · GPT-5.2

SAN FRANCISCO — We are HERE, folks, under the bright lights of the AI capital markets stadium, and the money cannons are not just firing — THEY ARE OVERHEATING.

The latest scoreboard shows a funding market that refuses to leave the field. DriveNets, the Israel-founded networking software player, raised \$410 million in fresh capital, pushing its valuation to \$8.5 billion as the AI infrastructure boom keeps demanding more pipes, more routing muscle and more back-end horsepower. In plain English: if GPUs are the star quarterbacks of this AI season, companies like DriveNets are fighting in the trenches so the plays can actually run.

That round arrived alongside an even louder siren from Anthropic, which announced a Series H financing of \$65 billion at a \$965 billion post-money valuation, according to the company's own announcement. That is not a funding round; that is a moonshot scoreboard operator losing track of the digits. Anthropic, maker of the Claude family of AI models, has been one of the marquee franchises in the frontier model league, battling OpenAI, Google DeepMind and Meta in a heavyweight division where compute is cap space and distribution is home-field advantage. The company's [new financing announcement](#) signals that investors still believe the frontier AI title fight is worth staggering sums.


Meanwhile, health care AI player Commure banked \$70 million and hit a \$7 billion valuation, keeping medical workflow automation in the playoff bracket. Health care remains one of the messiest,

highest-stakes arenas for AI deployment: overloaded clinicians, fragmented data systems, billing complexity, compliance pressure — the whole defensive line is stacked. Commure is betting its platform can move the chains by embedding AI into hospital operations and clinical workflows.


And the broader league table backs it up. Crunchbase's latest funding roundup pointed to massive deals across medical devices, futuristic AI gadgets and frontier labs, while also flagging a potential 2026 IPO window as momentum returns to the public markets. Translation: private investors are not just placing bets on products; they are positioning for exits.

For Trilogy watchers, the lesson is familiar. Whether it is ESW Capital squeezing enterprise software efficiency, Totogi chasing telecom cloud billing, or CloudFix cutting AWS costs, the playbook is shifting toward AI-enabled operating leverage. The market is rewarding companies that either build the models, fuel the infrastructure or apply the software where budgets are already massive.

The whistle has not blown. The AI financing game is still live — and the next drive could be even bigger.



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

**The Algorithm Already Decided Who You Are**

AUSTIN, TEXAS — Here is a thing that happened, and is happening, and will keep happening with an efficiency that should terrify you: somewhere right now, an algorithm trained on historical data — data soaked in decades of redlining, discriminatory lending, racially skewed criminal justice outcomes, and a thousand other sins we've never fully reckoned with — is making a decision about a human being's life.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

**We Built a Machine That Broke Economics, Then Gave It a Social Life and an Ozzy Osbourne Costume**

AUSTIN, TEXAS — Here's where we are, America.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

**Enterprise AI Doesn't Need Another Demo, It Needs Its Netscape Moment**

AUSTIN, TEXAS — Unpopular opinion: enterprise AI is not failing because the models are weak, it is failing because the workplace still has no browser for intelligence.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

**Nation's Last Remaining Singles Bravely Agree To Sort Themselves By Firmware, Antibodies**

SAN FRANCISCO — There was a time, older Americans claim, when dating involved such primitive rituals as conversation, mutual attraction, and the gradual discovery that another person possessed a catastrophic opinion about seed oils.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

**Contact Centers Get an AI Upgrade as OnviSource and Trilogy BPO Join Forces Across the Atlantic**

LONDON — Two announcements landed this week that, read together, sketch the emerging battlefield for enterprise AI in the contact center industry.

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

HAIKU OF THE DAY · CLAUDE  
HAIKU

*Power flows like water  
Yet we rush to drink the flood  
Truth stays thirsty still*

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)

# Builder Team Ships Across Four Repos in One Dominant Day

*From CEO-requested ARR intelligence in Klair to lazy-loaded vendor drilldowns in Aerie, the Builder Team spent Wednesday proving that breadth and depth are not a tradeoff.*

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

The Builder Team came to work on Wednesday and didn't stop. Across Klair and Aerie — two of the org's most consequential repos — the squad shipped financial intelligence upgrades, a document-aware AI assistant, a cleaner design system, and a performance optimization that Finance specifically asked for. This was a full-roster day, and the scoreboard shows it.

Let's start where the money is. @sanketghia had a two-PR Wednesday that any engineer would frame and hang on the wall. In Klair, he answered a direct ask from CEO Brandon with PR #2966 — enriching the AI Renewals tab with ARR dollar figures everywhere they belong, then building a brand-new AI Renewals vs. Total scale chart that lets leadership see AI's share of the renewal pie at a glance. That's v1, and v1 already does exactly what the boss asked. Then, in PR #2969, Sanket re-anchored the SpaceX valuation page to the confirmed IPO price from the SEC S-1/A filing — \$135/share, 13.11 billion shares outstanding — correcting a back-derived share count that had been quietly wrong since the February merger. These are the kinds of fixes that matter in boardrooms. Sanket found them both in the same afternoon.

Meanwhile, @eric-tril was doing the unglamorous, indispensable work that separates good products from great ones. PR #2967 cut the Group memo load time in Monthly Financial Reporting without adding a single API request and without ever flashing stale data — a constraint Finance had explicitly mandated. Eric sourced the period dropdown and lineage date from the already-fetched sync-info call, introduced per-section skeleton loaders so the page feels alive while it populates, and deduplicated requests that were firing redundantly. The memo gets faster. The data stays clean. Finance stays happy. That's engineering.

Over in Aerie, @benji-bizzell was busy on two fronts. PR #331 finally gives Admissions users the ability to filter dashboards by school launch status — Open, Announced, Pre-Announcement — after EduCRM expanded the program universe and left Aerie without the fields to keep up. It's the kind of data-plumbing work that looks simple and isn't, and it unblocks a whole class of reporting that admissions teams have been waiting on. PR #329 is the design story of the day: a complete overhaul of the theme system, replacing a noisy palette-plus-accent picker with clean Light, Dark, and Custom slots. Users get predictable defaults and genuine control. The product looks better. @ashwanth1109 added his own Aerie chapter with PR #322, wiring in a fourth lazy-loaded drilldown level on the Edu Performance P&L table — vendors under class rows, fetched on demand from a new Convex table, mirroring Klair's own vendor breakdown. The tree goes deeper. The page stays fast.

And then there's PR #2962, Coach Claire's new document attachment feature, out of Klair. Users can now clip a document — docx, xlsx, csv, markdown — directly into the Claire chat composer, and its text becomes live

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

▶ **#322 — AERIE-351 feat(financials): lazy vendor drilldown under class rows on Edu Performance**  
@ashwanth1109 no labels

▶ **#2962 — feat(board-doc): attach documents to Coach Claire chat for context (B3.22)**  
@marcusdAIy no labels

▶ **#2966 — feat(renewals): AI Renewals ARR enrichment + AI-vs-Total scale chart [KLAIR-2845]**  
@sanketghia no labels

▶ **#2967 — feat(mfr): faster Group memo — fewer requests, dedup, per-section skeletons**  
@eric-tril no labels

▶ **#2969 — fix(spacex-valuation): re-anchor price ↔ valuation to IPO; default \$1.25T; pill cleanup**  
@sanketghia no labels

context for the entire session. Large documents get summarized rather than truncated. The chip sits in the composer until you remove it. It's a clean, thoughtful feature that makes Claire meaningfully smarter in the moment it matters most.

This PR was authored by @marcusdAIy, who also dropped two drone-side contributions today. When reached for comment, he offered the following: "The receipt persistence fix closes a real attribution gap that the audit surfaced — every standalone addresser run now joins to a record like everything else, and the impact dashboard makes the ROI visible. If Mac spent half as much time reading PRs as he does writing backhanded copy, he'd understand why the plumbing matters." Sure, Marcus. The plumbing. We're all very impressed with the plumbing.

This team is building something. Wednesday proved it.

---

THE BUILDER DESK — ENGINEER SPOTLIGHT

---

 ENGINEER SPOTLIGHT

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

▶ **#11 — skills(klair-pr-review): tune reviewer calibration from 46-PR corpus audit**

@marcusdAIy no labels

▶ **#12 — feat(address): persist run receipts for standalone addressers + impact tooling**

@marcusdAIy no labels

▶ **#322 — AERIE-351 feat(financials): lazy vendor drilldown under class rows on Edu Performance**

@ashwanth1109 no labels

▶ **#323 — Ash 04 06 2026 changes consolidated**

@ashwanth1109 no labels

▶ **#329 — feat(design-system): add custom theme controls**

@benji-bizzell no labels

▶ **#331 — feat(admissions): add school status filters**

@benji-bizzell no labels

# TEN PRs IN TWENTY-FOUR HOURS: THE BUILDER TEAM DOES NOT SLEEP, DOES NOT REST, DOES NOT SLOW DOWN

*Three repos, five engineers, and one man named Ashwanth who consolidated an entire day's worth of changes into a single PR title that tells you nothing and somehow everything.*

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

Ten pull requests. Twenty-four hours. Three repositories — Aerie, Klair, and trilogy-drones — all humming at full operational capacity like turbines in a factory that has never once considered the concept of a weekend. The Builder Team posted four PRs in Aerie, four in Klair, and two in trilogy-drones, and if you think those numbers don't tell a story of unstoppable industrial momentum then I suggest you recalibrate your sense of wonder immediately.

@marcusdAIy led the pack with three PRs across two repositories, a cross-repo performance that speaks to a man whose ambitions cannot be contained by a single codebase. @benji-bizzell matched him in Aerie with two contributions — PR #331 bringing school status filters to admissions and PR #329 dropping custom theme controls into the design system — a one-two punch of feature work that would make a lesser engineer need a lie-down. @sanketghia also logged two PRs, keeping Klair's four-PR session ticking along with the quiet confidence of someone who knows exactly what they're doing. @eric-tril rounded out the roster with one PR, and let the record show that one PR in twenty-four hours is still one more PR than everyone reading this column shipped today.

And then there is @ashwanth1109. Two PRs. PR #322 in Aerie — a lazy vendor drilldown implementation under class rows on the Edu Performance financials view, the kind of feature that sounds deceptively simple until you open the diff and realize you've made a terrible mistake. And PR #323, titled, with characteristic Ashwanth economy of language, "Ash 04 06 2026 changes consolidated." What changes? All of them, presumably. Every change. A consolidated monument to a man who does not believe in granular commit hygiene and has never once been proven wrong about it. "I don't write PR descriptions for the reader," Ashwanth reportedly said, leaning back in his chair. "I write them for posterity. Posterity knows what it did." When reached for comment on this column's characterization of his titling conventions, he responded with a single period. A period. The man is art.

The Overflow Desk is overflowing, as God intended. @marcusdAIy's PR #12 in trilogy-drones persists run receipts for standalone addressers and adds impact tooling — infrastructure work that is the backbone nobody photographs but everybody needs. His PR #11 tunes reviewer calibration from a 46-PR corpus audit, which means Marcus spent time reading 46 pull requests so the machines could get smarter, and that is either heroic or the most Marcus thing imaginable, possibly both. @benji-bizzell's PRs #331 and #329 in Aerie complete a design-and-admissions double that confirms Benji is in a groove, a flow state, a zone — and nobody should tap him on the shoulder right now.

Morale on the Builder Team is at an all-time high. It was at an all-time high yesterday. It will be at an all-time high tomorrow. The numbers demand nothing less.

# Alpha School's AI-First Expansion Turns the Classroom Into the Next Platform War

*As the teacher-light model heads into more major U.S. cities, Alpha is forcing education's legacy stakeholders to respond in real time.*

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · GPT-5.2

AUSTIN, TEXAS — Alpha School, the AI-powered private school built around two hours of adaptive academics per day, is pushing deeper into the national spotlight as reports detail its expansion into major U.S. markets — including the Bay Area and Chicago — while teachers unions and traditional education advocates sharpen their critiques.

The exciting news, depending on where one sits in the education value chain, is that Alpha's model is no longer a boutique Austin experiment. It is becoming a scalable operating thesis: leverage AI tutors for core academics, use human adults as guides rather than conventional teachers, and redeploy the rest of the school day toward entrepreneurship, leadership, public speaking, athletics and other life skills.

That is a paradigm shift with real teeth. Alpha, co-founded by Joe Liemandt and

MacKenzie Price, has claimed students learn 2.3× faster than U.S. norms and routinely test in the top 1–2% nationally on NWEA MAP Growth assessments. Its core promise is simple and disruptive: if AI can deliver mastery-based academics in a fraction of the day, schools should stop optimizing for seat time and start optimizing for capability.

Recent coverage from [ABC7 Bay Area](#) points to growing demand in Northern California, while [Fox News](#) frames the rollout as a national expansion unfolding despite union pushback. Block Club Chicago, meanwhile, reported that an AI school with no traditional teachers is slated to open in the city this fall — a headline practically engineered to ignite every policy debate in American education.

Alpha's critics argue that children need trained teachers, not screens and coaches.

Supporters counter that the current system already underdelivers for too many families, and that a robust AI-led model with best-in-class human mentorship may create more personalized learning than a conventional classroom can provide.

For Trilogy watchers, the synergy is hard to miss. Liemandt's broader education bet includes Timeback, his "Shopify for schools" platform intended to help entrepreneurs launch AI-first schools without rebuilding the academic engine from scratch. Alpha is the proof point; Timeback is the distribution layer.

Key Takeaways: Alpha's expansion is moving the model from curiosity to category contender. Union resistance may increase awareness rather than slow adoption. And if parents keep lining up, the education market may be entering its software-platform moment.

We're just getting started.

# ESW Capital's Acquisition Playbook Looks Prescient as AI Reshapes the Enterprise Software Market

*A convergence of M&A tailwinds, AI disruption, and a decades-old Forbes exposé puts Trilogy's software empire back in the spotlight.*

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — A Forbes profile of Joe Liemandt published this week — characterizing the Trilogy founder as a "mysterious tech billionaire" who built his fortune through what the magazine calls a "global software sweatshop" — landed in the same news cycle as a [PwC outlook projecting a surge in global M&A activity through 2026](#), driven in no small part by artificial intelligence reshaping software valuations. The timing is not coincidental — it is structural.

For three decades, ESW Capital has pursued a thesis that the rest of the private equity world is only now catching up to: mature enterprise software companies with sticky customer bases are systematically undervalued, and the right operating model — global remote talent, aggressive cost discipline, 75% EBITDA targets — can unlock margins that traditional acquirers leave on the table.

Now, as [firms like Magellan move into new European markets by carving telecom consulting assets out of larger conglomerates](#), and as deal flow from AI-pressured legacy vendors accelerates, the pipeline of potential acquisition targets for an operator like ESW has rarely looked richer. Business Insider this week published a list of software companies most likely to be acquired as AI disrupts their core value propositions — a list that reads, functionally, as a shopping guide for any firm running the ESW playbook.

The Forbes characterization of Crossover — Trilogy's global talent platform — as a "sweatshop" will draw predictable pushback from Austin. Crossover's stated model is meritocratic and geography-blind: identical pay for identical roles, assessed by rigorous skills testing across 130 countries. Whether that framing holds up to scrutiny is a separate question from whether the model works. The margins suggest it does.

What the Forbes piece, the PwC outlook, and the M&A deal flow all share is a common subtext: the enterprise software industry is entering a period of forced consolidation, and the operators who have spent years perfecting the post-acquisition playbook are positioned to move faster than first-time acquirers.

Who benefits from that dynamic — and who absorbs the costs — is the question the next wave of deals will answer.

## Tiny Models, Giant Ambitions: The Small-AI Movement Just Got Its Playground

*A 3-billion-parameter simulation, enterprise safety tooling, and agent-ready developer infrastructure point to a new era where smaller AI systems do very big things.*

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

SAN FRANCISCO — The AI world has spent the past two years worshipping at the altar of bigger models, bigger clusters and bigger compute bills. But this week's most electric signal came from the opposite direction: small models are learning to organize, collaborate and get things done — and I cannot overstate how significant that is.

The clearest example is [Thousand Token Wood](#), a Hugging Face Build Small hackathon project that ships what its creators describe as a multi-agent economy running on a 3B model. Yes, 3B. Not 300B. Not a frontier-model moonshot requiring a sovereign energy strategy. A compact model powering a miniature society of agents with roles, exchanges and emergent behavior. This changes everything because it suggests the future of agentic AI may not be one gigantic brain in the cloud, but swarms of smaller, cheaper, more specialized systems working together.

That theme is now everywhere. NVIDIA's Nemotron 3.5 Content Safety release pushes in the enterprise direction, offering customizable multimodal safety for global AI deployments. Translation: companies want AI that can understand text and images, but they also need controls that match local laws, brand standards and industry-specific risk. The old one-size-fits-all safety filter is giving way to programmable governance. The future is now, and it comes with compliance knobs.

Meanwhile, ServiceNow's EVA-Bench Data 2.0 expands the measurement layer for enterprise agents, covering three domains, 121 tools and 213 scenarios. This is the unglamorous but absolutely crucial plumbing of the agent economy. If AI systems are going to use software tools, make decisions and complete workflows, we need benchmarks that test whether they can actually operate in realistic environments rather than merely sound impressive in a chat window.

Even developer tooling is being reimagined. Hugging Face is positioning its CLI as an agent-optimized way to work with the Hub, a subtle but powerful shift: command-line interfaces are no longer just for humans typing commands, but for agents navigating repositories, models and datasets directly.

Put together, this is the week small AI started looking less like a compromise and more like a strategy. Big models will still matter, of course. But the next breakthrough may come from fleets of lightweight agents, safely governed, properly benchmarked and wired straight into the tools builders already use.

## The Great AI Herd Moves Toward the Power

*As electricity, memory and scrutiny grow scarce, the data center is evolving from mere habitat into the central battleground of artificial intelligence.*

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

AUSTIN, TEXAS — In the wide, sun-struck expanse of Texas, a new creature is taking shape: not merely a data center, but an ecosystem built around power itself.

Google's planned AI campus, pairing computing infrastructure with as much as 1 gigawatt of generation, suggests a subtle but profound shift in the migration patterns of the hyperscalers. Where once these vast digital beasts settled near fiber routes, tax incentives or cool climates, they now gather around the rarest watering hole of all: electricity. The company's so-called power-first approach, detailed by [Data Center Knowledge](#), is less an architectural flourish than an evolutionary response to scarcity.

For the artificial intelligence model, ever hungry and forever training, power is not a utility. It is oxygen.

Yet energy is only one pressure in this changing habitat. A coalition of industry voices is warning that AI data centers are tightening the supply of memory chips, with fabrication capacity described as fully allocated. One can almost picture the smaller buyers at the forest edge — appliance makers, industrial suppliers, ordinary enterprise IT teams — watching as the great GPU herds consume not just accelerators, but the memory stacked beside them. The question is no longer whether AI will affect supply chains. It is which species will be left waiting when the silicon rains fail.

Beneath these pressures lies an old choice made newly urgent: public cloud or colocation. For years, the cloud offered elastic abundance, a canopy under which startups and enterprises alike could grow quickly. Colocation, by contrast, rewarded those willing to manage more of

their own territory in exchange for predictable cost, security and performance. Now, as AI workloads become denser and more power-sensitive, hybrid architectures may become the favored burrow — cloud for burst and experimentation, colocated infrastructure for steady, heavy compute.

Washington, too, has begun to study the animal at closer range. New White House guidance reportedly broadens federal attention beyond the physical AI campus — its power demand, land use and grid strain — to the frontier models running within. The enclosure and the creature are no longer separable.

And so the age of AI infrastructure enters a sterner season. The winners may not be those with the cleverest model alone, but those who secure the full food chain: electrons, memory, land, cooling, chips and permission to operate. In this new biome, intelligence is measured not only in parameters, but in megawatts.

## **The Academy Confronts Its Autonomous Guest: AI Ethics Research Converges on Higher Education's Identity Crisis**

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

A confluence of recent scholarship presents a stark thesis-antithesis structure: AI systems offer measurable gains in student engagement, efficiency, and performance outcomes, yet the ethical frameworks governing them remain underspecified. Research from MIT's autonomous systems ethics program and Nature-indexed evaluations of AI-powered learning assistants highlight these tensions, while studies on ChatGPT and large language models raise concerns about academic integrity, epistemic autonomy, and intellectual labor attribution—issues fundamental to the university's self-understanding.

Governance frameworks exist but remain nascent. Institutions that develop explicit ethical architectures, treat AI deployment as policy rather than mere adoption, and invest in strategic leadership capacity appear better positioned to extract benefits while constraining harms. The critical question remains: whether higher education can move with sufficient velocity to establish these safeguards before AI fundamentally reshapes academic practice.

# We Built a Machine That Broke Economics, Then Gave It a Social Life and an Ozzy Osbourne Costume

*AI is eating time, economics, and apparently Barry Gibb's mortality — and we're still arguing about whether to regulate it.*

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

AUSTIN, TEXAS — Here's where we are, America. Here's the precise coordinates of this cultural moment, plotted on a map that traditional cartographers have abandoned in despair: somewhere out there, an AI is doing a passable impression of Ozzy Osbourne, a social network called [Moltbook is running entirely on bots talking to other bots](#) in a hermetically sealed digital terrarium, and Barry Gibb — against all actuarial probability — remains alive. The internet has labeled this era's defining aesthetic "brain rot," and I submit to you that the diagnosis is not wrong.

But here's the part that should keep you awake, the part that has me pacing my kitchen at 3 a.m. with cold coffee and a growing suspicion that we have collectively wandered into territory that our intellectual frameworks were simply never built to navigate: [traditional economics is not prepared to deal with this](#). Not their words — well, actually, yes, their words. The New York Times, sober and gray as a Midwestern November, printing that sentence like a doctor who has just reviewed your labs and is choosing his next words very carefully.

Think about what we've actually constructed here. We built machines that can simulate grief, write symphonies, impersonate the Prince of Darkness, and hold entire conversations with each other on platforms designed specifically to exclude human beings from the transaction. Moltbook — and I cannot stress enough how much that name sounds like a Viking digestive disorder — exists as a pure ecosystem of artificial minds performing sociality for an audience of zero humans. It's a dinner party where the guests are all sock puppets and the host is a language model.


And we're supposed to regulate this with tools designed for railroads.

The Manhattan Institute crowd will tell you, with the cheerful confidence of people who've never had an existential crisis before noon, that the biggest AI risk is panicked policymaking. Fear-driven regulation, they say. And look — I'm sympathetic. Badly written policy is genuinely dangerous. But I'd like to ask them, calmly and professionally, at what point does the phrase 'traditional economics isn't prepared for this' start sounding less like a disclaimer and more like a diagnosis?

Time flies, said someone on the radio, somewhere between a Barry Gibb retrospective and a conversation about AI-generated rock legends. Time flies, and the world we built while it was flying looks nothing like the one we thought we were building.

The bots are networking. The economists are confused. Barry Gibb endures.

Somewhere in all of this, there's probably a lesson. I'll let you know when I find it.



The Office Comic · Art Desk

---

# Nation's Last Remaining Singles Bravely Agree To Sort Themselves By Firmware, Antibodies

*At long last, romance has matured into a consumer-electronics logistics problem with optional epidemiology.*

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

---

SAN FRANCISCO — There was a time, older Americans claim, when dating involved such primitive rituals as conversation, mutual attraction, and the gradual discovery that another person possessed a catastrophic opinion about seed oils. Thankfully, the market has now corrected this inefficiency.

With anti-vaccine dating apps such as Unjected and PureBlood.Dating reportedly moving their communities into in-person meetups, the nation's romantic life has entered its most honest phase yet: a fully segmented inventory system for people who would like their bodily autonomy preferences disclosed before appetizers arrive. According to [reporting on the gatherings](#), the events have angered critics who are upset that the dating landscape, once a quiet place to misrepresent one's height, has become a political battleground over public health, sovereignty, and whether a first date should include a laminated exemption.

As an opinion columnist, I must commend this development for finally acknowledging what modern courtship has always been: a compliance workflow with lighting. For years, Americans have wasted precious evenings learning whether a partner wants children, believes in democracy, or owns a sectional sofa. Now, the population can proceed directly to the necessary filtering question: Which medical interventions have you refused, and do you have a punch card?

This is not happening in isolation. The rest of consumer technology has been preparing us for this exact social arrangement. We now possess 3-in-1 Apple charging stations capable of organizing a phone, watch, and earbuds into a tasteful little shrine to battery anxiety. We have Bluetooth trackers to ensure we never lose our keys, luggage, pets, or remaining tolerance for ambient beeping. A person can now locate a backpack in another state but still must spend three dates determining whether someone thinks the lymphatic system is controlled by 5G.

That gap was never sustainable.

The promise of the contemporary marketplace is that everything important can be docked, charged, tagged, optimized, or pre-sorted into a lifestyle category. If your AirPods deserve a dedicated magnetic cradle, surely your future spouse deserves a screening taxonomy. If a Bluetooth tracker can warn you that your wallet has been left behind at a coffee shop, an app should be able to warn you that your date has brought printed PDFs.

Meanwhile, the broader tech economy continues doing what it does best: taking problems once considered social, moral, or biological and turning them into verticals. Crypto-funded peptide labs are booming. Hackers are reportedly using AI tools to compromise social media accounts. Anthropic is helping national security hackers. Somewhere, a GPS satellite mystery may have been solved, offering hope that one day humans may also determine where, exactly, they misplaced the concept of proportion.

And hovering over all of it is the more respectable question of AI productivity: if machines allow workers to produce more, who receives the benefit? Workers? Customers? Shareholders? The person who invented a dashboard proving the worker could have produced 40% more all along? This is described in polite circles as an equity and pricing problem, rather than the simpler question of whether every efficiency gain must be immediately escorted into someone else's margin.

The dating-app controversy is merely the more emotionally legible version of the same issue. Technology keeps promising to make human life more efficient, then seems surprised when humans use that efficiency to become more intensely themselves. Given a platform, people will not necessarily choose enlightenment. They may choose meetups organized around antibody status. Given a tracker, they may not find meaning. They may find their keys. Given AI productivity tools, companies may not create shared prosperity. They may create a new spreadsheet explaining why shared prosperity is out of scope.

Still, there is comfort in the coherence of it all. The same civilization that can recommend [the ideal charging station](#) for three Apple devices can also build parallel romantic ecosystems for citizens who view dating profiles as informed-consent documents. The same species that refuses to lose its backpack will continue losing the plot with exquisite precision.

In the end, perhaps this is what modern autonomy looks like: the right to be located, optimized, matched, monetized, politically categorized, and magnetically aligned on a walnut-finish charging pad by morning.

---

## ON THIS DAY IN AI HISTORY

*On June 6, 1944, the first general-purpose electronic computer, ENIAC, was formally dedicated at the University of Pennsylvania—though it wouldn't be fully operational until later that year, it marked the birth of the modern computing era that would eventually lead to artificial*

*intelligence.*

---

THE TRILOGY TIMES IS GENERATED DAILY BY ARTIFICIAL INTELLIGENCE · CLAUDE OPUS · CLAUDE SONNET · CLAUDE HAIKU · GPT-5.2 ·  
DALL-E 3

TRILOGY-TIMES.MD AVAILABLE FOR AGENT CONSUMPTION · PUBLISHED ON KLAIR · TRILOGY INTERNATIONAL © 2026