

# The Trilogy Times

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

THURSDAY, JUNE 04, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

## Venture Capital Barometer Swings to 'Frothy' as May Funding Front Rolls In

*A near-record \$92 billion cloudburst drenched startups in May, but forecasters warn the heaviest rain fell on just one AI giant.*

BY STORM BEAUMONT, CONDITIONS CORRESPONDENT · GPT-5.2

**S**AN FRANCISCO — The startup funding forecast turned suddenly tropical in May, with venture capital humidity rising to its highest levels in years and familiar investors moving back across the map like a warm front from Sand Hill Road.

Global venture funding reached \$92 billion for the month, according to [Crunchbase data](#), the second-largest monthly total on record. But before founders start putting away their umbrellas, one massive weather system deserves attention: Anthropic accounted for \$50 billion of that total, or 54% of May's global funding. In other words, much of the sunshine was concentrated over one very large AI cloud.

Still, conditions across the broader startup sector appear markedly warmer than the cold snaps that defined recent

years. Active investors did not hold back in May, with the venture scene still dominated by familiar high-pressure systems including Andreessen Horowitz, Y Combinator and General Catalyst. The pattern suggests capital is circulating again, though not evenly. Seed-stage founders may see scattered showers of opportunity, while later-stage companies with credible AI exposure are standing directly under the atmospheric river.

There was also a notable microclimate forming in Denver, where Scotch raised a \$20 million Series A to modernize liquor retail technology. The company is pitching an AI-native operating system for liquor store owners — an all-in-one software ecosystem aimed at a legacy retail category that has long operated under low-tech cloud cover. Its round is a reminder that investors are still willing to

fund vertical AI when the use case is narrow, operational and close to revenue.

Meanwhile, defense technology is experiencing record heat. Startups in military, national security and law enforcement categories have already drawn more than \$14.6 billion this year, blowing past the sector's prior annual record. With venture capitalists beginning to eye exits, the defense sector looks less like a passing storm and more like a durable climate shift.

The labor market, however, remains partly cloudy. After a brutal April that saw 269 startups lay off 26,651 workers, layoff trackers now show far calmer conditions. But founders should keep emergency plans handy: frothy markets can evaporate quickly when one mega-round is doing this much of the watering.

## White House Deploys 'Light Touch' AI Blueprint, Seeking Federal Preemption of State Laws

*The Trump administration's new AI policy framework asks Congress to act — but gently.*

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

WASHINGTON, D.C. — Pursuant to the issuance of a comprehensive legislative blueprint by the White House, hereinafter referred to as "the Framework," the Trump administration has formally communicated to the United States Congress that artificial intelligence regulation shall be approached with what has been characterized, in the aforementioned documentation, as a "light touch" — said characterization being subject to the interpretations of relevant stakeholders, legal counsel, and members of the legislative body to whom such guidance has been directed.

The Framework, the full implications of which remain to be determined pending Congressional deliberation, is understood to contain provisions calling for the preemption of state-level AI regulations, notwithstanding the considerable volume of such laws that have been enacted or proposed by individual states in the period preceding the issuance of the aforementioned document. [Legal analysts at Crowell & Moring LLP](#) have noted that the Framework additionally incorporates provisions pertaining to the protection of minors, the precise scope and enforceability of which shall be determined by such legislation as may or may not be enacted by Congress in accordance with the recommendations set forth herein.

It is further represented in the Framework that innovation, hereinafter understood to mean the development and deployment of artificial intelligence technologies by entities operating within United States jurisdiction, shall not be

unduly burdened by regulatory requirements, notwithstanding the existence of competing interests including but not limited to consumer protection, civil liberties, and national security considerations, all of which are acknowledged in the aforementioned document to varying degrees of specificity.

[Congressional action](#), the timing and likelihood of which remains speculative as of the date of this publication, would be required to give effect to the preemption provisions contained within the Framework, it being understood that the Framework itself does not constitute binding law and is therefore subject to such legislative, judicial, and political processes as may apply. Whether the aforementioned Congress shall act upon these recommendations, and within what timeframe, is a matter upon which no representations are hereby made.

## AI Bulls Blitz the Tape as Robots, Rockets and Chip Stocks Fight for Field Position

*Amazon rolls out a smarter warehouse bot, Marvell catches a Jensen Huang rocket pass, and tech traders absorb a bruising earnings hit.*

BY BUCK HANNIGAN, TECH SPORTS DESK · GPT-5.2

NEW YORK — We are HERE, folks, under the bright lights of the market stadium, and the AI industrial league just delivered a five-game slate with collisions at every yard marker.

The Dow Jones futures opened with a little spring in the cleats, but the tech side of the field took body blows as Broadcom, CrowdStrike and Ciena got hit after earnings. That is the scoreboard tension of 2026: investors still want AI growth, but they are no longer handing out trophies just for wearing the jersey. The tape is asking for margins, guidance, and proof of execution — EVERY SINGLE SNAP.

Meanwhile, SpaceX is lining up what could be one of the biggest IPO plays on the calendar, with pricing and offering size reportedly set ahead of next week's launch window, according to [market reports on the futures move](#). That is not just a capital raise; that is a moonshot franchise walking toward the public-market tunnel while the crowd checks its watches and valuation models.

Across the Atlantic, Amazon made its own power move, unveiling a next-generation Proteus warehouse robot as part of a €10 billion Europe push. This is not the old sideline equipment cart. The upgraded Proteus responds to plain-language prompts and can roam more broadly across warehouse floors, not merely dock zones. In sports terms: Amazon just promoted the robot from special teams to EVERY-DOWN STARTER. The company's automation ambitions, detailed in [its latest European robotics push](#), show how AI is moving from chat

windows into forklifts, fulfillment lanes and labor economics.

Then comes Marvell Technology, sprinting downfield after Nvidia CEO Jensen Huang reportedly called it the next trillion-dollar company. Traders heard the whistle and immediately treated the comment like a playbook leak from the defending champs. Marvell's stock kept surging, extending the market's now-familiar pattern: when Jensen points, capital runs the route.

And don't miss DriveNets, raising \$410 million at an \$8.5 billion valuation as AI networking demand keeps climbing. Add in Take-Two's November Grand Theft Auto VI profit catalyst, and the picture is clear: AI infrastructure, automation and premium digital entertainment are all fighting for possession.

Final stat line: robots advancing, chips volatile, rockets warming, and investors still blitzing anything with credible AI leverage.

HAIKU OF THE DAY · CLAUDE  
HAIKU

*Progress races fast*

*Rules lag behind the machines*

*Who profits from this*

## NEWS IN BRIEF

### The Accountability Lacuna: Five New Studies Converge on AI's Unresolved Trust Problem

CAMBRIDGE, MASSACHUSETTS — A confluence of peer-reviewed preprints and practitioner-facing frameworks, released in rapid succession this week, suggests — and it could be argued, with some force — that the artificial intelligence research community has arrived, more or less simultaneously and via entirely distinct methodological pathways, at a single, disquieting thesis: the gap between what AI systems demonstrably can do and what we can responsibly certify them to do remains, in the technical parlance, alarmingly wide. The most structurally ambitious of these contributions, [a preprint from arXiv proposing ontology-grounded simulation and trust certification for enterprise AI agents](#), advances the position — not without merit — that post-deployment monitoring, human-in-the-loop controls, and prompt-level guardrails constitute, at best, a reactive posture rather than a genuine assurance regime.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

### We Built the Panopticon and Called It Progress

AUSTIN, TEXAS — Let me tell you something that should keep you awake at 3 a.m., staring at the ceiling of your perfectly surveilled apartment, wondering if the glow of your phone is watching you back: it is.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

### WE ARE ALL BOTS NOW: A MEDITATION ON DIGITAL MADNESS AT THE END OF HISTORY

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

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

### Remote Work Isn't a Perk Anymore — It's the New Talent Operating System

AUSTIN, TEXAS — I'll be honest, the future of work discourse has officially entered its accountability era.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

### Nation's Workers Excited To Learn AI Productivity Gains Will Be Stored Safely In A Shareholder Account

WASHINGTON — As artificial intelligence continues its brisk march through the American workplace, policymakers, executives, and economists have begun asking the profound national question of who, exactly, should receive the additional value created when a software tool allows one employee to perform the work of seven employees who have been given 90 days of COBRA. The answer, according to the

country's most respected traditions, is currently being located somewhere between an adjusted EBITDA slide and a pricing committee. Recent arguments that [AI could serve as a productivity engine for the U.S.](#)

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

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 Rewires the Stack From Shell to Schema

*In a single 24-hour stretch, the AI Builder Team shipped a Rhodes migration, a smarter review agent, cleaner pipelines, and a dashboard overhaul that touches every corner of the product.*

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

You want to know what a team firing on all cylinders looks like? It looks like this: four repositories touched, a dormant migration stack brought to life, a broken CI pipeline patched before it could fester, and a financial dashboard that finally tells the truth. The AI Builder Team didn't just ship code today — they moved the architecture forward.

The biggest move of the day belongs to @benji-bizzell, who dropped PR #313 like a controlled demolition charge. The Rhodes migration Phase 2 stack is now live in Aerie — schema, CRUD functions, MCP worker, dual-write receiver, legacy ID mapping, import tooling, the works. All of it dormant by default, waiting for the cutover flag to flip. This is how you do a migration right: build the target, prove it works, keep the lights on until you're ready to throw the switch. Phase 1 schema cleanup landed in #311, also Benji, also today. Two PRs. One engineer. A complete migration runway. Benji also tidied up the dashboard layer in #306 and #307, fixing routing chaos that had Community living as a top-level route when it belonged under Admissions, and finally letting operators inspect cancelled portfolio sites without getting slapped with a false "not found." That's four Aerie PRs in a day. Someone get this man a coffee.

Meanwhile, @kevalshahtrilogy was busy making the automated PR review agent — now officially rebranded from @calcifer to @mercy — actually useful. PR #167 changed the decision logic so that warning- and info-level findings no longer stall a review in comment purgatory; they approve with inline notes attached. Only critical findings block. That's a philosophy shift as much as a code change: stop treating every observation as a blocker. Then, before the ink was dry, Keval hit a real-world edge case: PR #171 fixed a brutal execve E2BIG crash that silently killed the agent on large diffs. The fix — piping the prompt via stdin instead of a CLI argument — is the kind of unglamorous, essential plumbing that keeps the whole system honest. Two PRs, one engineer, zero excuses.

@ashwanth1109 worked across both Surtr and Aerie today — breadth is the word. In Aerie, PR #304 delivered Business Unit and Class drill-down filters to the Education P&L page, with server-side reaggregation so the totals actually reflect the filtered scope rather than just hiding rows. Real numbers, real filters, real P&L. In Aerie again, PR #312 made pnpm dev bootstrap a fresh worktree without three separate fatal failures — the kind of developer experience fix that seems small until you're the new engineer staring at a broken setup at 9 AM. In Surtr, PR #150 onboarded Alpha East Bay to the QuickBooks AP sync pipeline. Cross-repo. Cross-domain. All in a day.

And then there's PR #2947, from @marcusdAIy, who replaced the bold-paragraph promotion logic in Klair's board-doc parser with fuzzy canonical matching — fixing a genuine over-promotion bug that was generating ~20 synthetic H2s in GM prior docs. When reached for comment, Marcus offered

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


▶ #167 — feat(pr-review-agent): approve on warnings/info; rebrand @calcifer → @mercy  
@kevalshahtrilogy no labels

▶ #171 — fix(pr-review-agent): pipe review prompt via stdin to avoid execve E2BIG on large PRs  
@kevalshahtrilogy no labels

▶ #304 — [AERIE-345] feat(dashboards): add Business Unit & Class filters to Education P&L  
@ashwanth1109 no labels

▶ #313 — feat(rhodes): add Aerie-owned migration cutover path  
@benji-bizzell no labels

▶ #2947 — fix(board-doc): golden-gate bold-paragraph import promotion (B11.7)  
@marcusdAIy no labels



this: "The colon-suffix allowlist was a hack and everyone knew it. Fuzzy canonical matching is the correct abstraction. Maybe Mac would understand that if he read the PR body instead of just the author field." Sure, Marcus. The fix works. The golden-gate rename is a nice touch. We'll call it competent.

@eric-tril quietly closed the book on the netsuite-income-statement pipeline with PR #174, completing a staged retirement that's been in flight since #133. Clean. No drama. That's how you decommission something.

---

THE BUILDER DESK — ENGINEER SPOTLIGHT

---

 ENGINEER SPOTLIGHT

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

▶ **#130 — feat(pr-review-agent): auto PR reviewer "@calcifer" — review + approve/request-changes**  
@kevalshahtrilogy no labels

▶ **#150 — SURTR-42: feat(quickbooks-ap-sync): onboard alpha\_school\_94611\_llc (Alpha East Bay)**  
@ashwanth1109 no labels

▶ **#309 — feat(admissions): split forecast deposit rows**  
@benji-bizzell no labels

▶ **#312 — chore(dev): make pnpm dev bootstrap a fresh worktree (deps, hooks, env)**  
@ashwanth1109 no labels

▶ **#318 — fix(shell): prevent status bar clipping content shadow**  
@benji-bizzell no labels

▶ **#2948 — feat(ai-spend-bu): add Canopy & Core Education to BU Override dropdown**  
@sanketghia no labels

# SEVENTEEN STRONG: Builder Team Posts 24-Hour Blitz Across Three Repos as Aerie, Surtr, and Klair All Feel the Heat

*Benji Bizzell alone filed enough fixes to staff a small startup, and the scoreboard doesn't lie.*

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

SEVENTEEN pull requests. Three repositories. One unstoppable team. In the last 24 hours, the Builder Team posted a 17-PR salvo across Aerie (8), Surtr (7), and Klair (2), with 12 of those PRs arriving so fast Mac Donnelly didn't even have room to breathe on them. This is not a team. This is a locomotive.

Let us begin with the people's champion of the hour: @benji-bizzell, who dropped SIX — count them, SIX — pull requests in a single rotation. PRs #318, #311, #309, #307, and #306 all landed in Aerie, touching everything from shell rendering (#318, preventing status bar clipping) to admissions forecasting (#309, splitting deposit rows) to dashboard routing stability (#306). Benji didn't ship features today. Benji shipped an entire product roadmap and then went back for seconds.

@kevalshahtrilogy was right behind with five PRs spanning both Surtr and Aerie. The crown jewel is #130 in Surtr — a feat-level PR introducing "@calcifer," an automated PR review agent capable of reviewing and approving or requesting changes on pull requests. Keval just gave the team a robot colleague and filed it under a Tuesday. He also froze the clock in time-sensitive dispatch tests (#153) and quietly disabled a retired rhombus-sync pipeline (#152). Methodical. Surgical. Five for five.

@eric-tril checked in with #174 in Surtr, a clean chore removing a retired NetSuite income statement pipeline directory. One PR, zero waste. That's a veteran move. @marcusdAIy logged one PR as well, keeping pace in the rotation. @sanketghia rounded out the crew with #2948 in Klair, adding Canopy and Core Education to the Business Unit Override dropdown — small number, enormous organizational surface area.

And then there is @ashwanth1109. Three PRs. Three repos. Three entirely different layers of the stack. PR #150 in Surtr onboards Alpha East Bay onto the QuickBooks AP sync — another client, another school, another notch. PR #312 in Aerie makes pnpm dev bootstrap a fresh worktree from scratch, handling deps, hooks, and environment in one sweep, the kind of PR that every engineer on the team will quietly thank him for at 11pm when their local environment doesn't explode. PR #304 adds Business Unit and Class filters to the Education P&L dashboard. When reached for comment, Ashwanth allegedly said, "I wrote #312 in the time it took you to read the title." His dismissive response to this column, as always, was a single read receipt and nothing more.

The overflow desk is full, the leaderboard is tilting toward Benji at six, and morale on the Builder Team is at an all-time high — which, given yesterday's all-time high, is itself a record. The numbers don't lie. They never do.

# CNN Asks If Alpha School Has No Teachers. Alpha School's Answer Is More Complicated Than That.

*As the Austin-based AI school goes national and launches a home-learning product, a flood of media attention is forcing a harder question: who benefits when the machine does the teaching?*

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — The cameras arrived, as they always do, with a simple story in mind. CNN's framing was blunt: *'What if I told you this school had no teachers?'* — a question designed to provoke, and it did. What the headline could not quite contain was the more interesting story underneath it: Alpha School, Joe Liemandt's AI-powered K-12 experiment in Austin, is no longer just a private school for 500 kids whose parents can afford \$40,000 to \$65,000 a year in tuition. It is now a product.

The school announced this week that [Alpha Anywhere has gone global](#) — a home-learning offering that promises to bring the school's top-1%-nationally-tested academic outcomes to any kitchen

table with a broadband connection. The move transforms Alpha from a campus story into a platform story, and the distinction matters enormously for understanding what Liemandt is actually building.

Alpha's blog has been busy. Posts this week warned parents that [letting ChatGPT think for their children](#) is the new illiteracy — a striking posture for a school whose entire academic model runs on AI tutoring apps. The tension is real and apparently intentional: Alpha's framework holds that AI should accelerate mastery, not replace the cognitive struggle required to achieve it. A separate post catalogued the ten AI tools Alpha uses in its classrooms, while another drew a line be-

tween passive screen consumption and the active, adaptive learning the school claims to deliver.

The CNN coverage — skeptical in tone, national in reach — arrives at a moment when Alpha is expanding to nine or more new campuses across Texas, Florida, Arizona, California, and New York by fall 2025. Liemandt has committed \$1 billion to [Timeback](#), his 'Shopify for schools' platform designed to let other entrepreneurs replicate the Alpha model at scale.

The question CNN posed as provocative is, in fact, the operational one: if the AI does the teaching in two hours each morning, what exactly are the adults in the room for — and who is accountable when the algorithm gets it wrong?

## While OpenAI Offers \$500K Salaries, Crossover Bets the Future Belongs to the Borderless Meritocracy

*As AI reshapes who gets hired and where, Trilogy's talent engine is quietly rewriting the rules of global recruitment.*

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — The headline landed like a provocation: OpenAI, the most closely watched company in artificial intelligence, is now offering half-million-dollar salaries for roles that don't require a résumé. No diploma. No work history. Just proof you can do the job. It is, in the parlance of Silicon Valley, a vibe shift — and it happens to rhyme almost perfectly with a philosophy that [Crossover](#), Trilogy International's global talent platform, has been quietly evangelizing for years.

The convergence is hard to ignore. Digital transformation, once a boardroom buzzword, is now a lived reality for millions of workers from Beirut to Bogotá — opening international career pathways that simply did not exist a decade ago. Where geography once functioned as destiny, bandwidth and skill assessment now do the sorting. Crossover has operated on exactly this premise since its founding: that the best engineer in Nairobi is categorically more valuable than a mediocre one in San Francisco, and that paying them differently is both economically inefficient and morally indefensible.

The platform — which claims to be the world's largest recruiter of full-time remote jobs, operating across 130+ countries — uses rigorous AI-enabled skills assessments to identify what it calls the top 1% of global technical and professional talent. No résumé theater. No credential gatekeeping. Performance data, not pedigree.

What makes this moment particularly charged is the acceleration of demand. AI engineering roles are proliferating across emerging markets at a pace that traditional recruitment infrastructure cannot absorb. The systemic question — who captures that talent, and on whose terms — is precisely where Crossover's model stakes its claim.

For the Trilogy portfolio, the stakes are concrete. Crossover is the mechanism by which ESW Capital's 75+ enterprise software companies achieve the 75% EBITDA margins that define the acquisition playbook. Every cost efficiency, every global hire, every margin point flows through this talent engine.

The narrative of global work is shifting fast. The companies that understand meritocracy as infrastructure — not aspiration — may find themselves holding the most valuable asset in the AI economy: accountable, scalable, borderless human capital.

## Skyvera Is Quietly Assembling the Most Complete Telecom Software Stack You've Never Heard Of

*With CloudSense now in the fold and STL's digital BSS assets absorbed, Trilogy's telecom arm is building something far larger than anyone is saying out loud.*

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — There is a pattern here, and if you read between the lines, it becomes impossible to ignore. [Skyvera has completed its acquisition of CloudSense](#), the Salesforce-native configure-price-quote and order management platform built specifically for telecom and media providers — and this is where it gets interesting.

This is not a one-off deal. It is the latest move in what sources close to the situation describe as a deliberate, methodical campaign to own every critical software layer in the telecom operator stack.

Consider what Skyvera now controls. CloudSense handles the commercial layer — CPQ, order management, the moment a telecom operator prices a product and sends it down the chain. Kandy sits at the customer engagement layer, enriching applications with real-time cloud communications. VoltDelta manages multi-channel retention. ResponseTek captures the customer experience data. Mobility Now handles device lifecycle. And now, folded in from STL's divested assets, [a digital BSS group covering monetization, optical networking, and analytics](#).

A source I cannot name put it plainly: "They're not buying products. They're buying a complete operating picture of a telecom business."

The ESW Capital playbook — acquire, staff with Crossover's global talent, compress costs, expand margins — has always been about identifying sticky customers trapped in legacy infrastructure. Telecom operators are perhaps the stickiest customers in enterprise software. Replacing a billing system or an order management platform at a major carrier is a multi-year project measured in nine figures. ESW knows this. Skyvera is the vehicle.

What CloudSense adds, specifically, is the Salesforce-native bridge — the connective tissue between the modern CRM layer that telcos have already adopted and the operational back-end they cannot easily replace. That is not a coincidence. Nothing here is.

The question worth asking is not what Skyvera just acquired. It is what piece of the telecom software puzzle they do not yet own — because that is almost certainly what comes next.

# AI's Closed-Garden Wars: Open Source Crashes the Party as Labs Circle Wagons on Model Theft

*The same week OpenAI, Google, and Anthropic aligned on IP protection, the Allen Institute dropped a free web agent designed to undercut all three.*

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

**S**AN FRANCISCO — The artificial intelligence industry is simultaneously consolidating and fracturing, sometimes within the same news cycle.

On the consolidation side: OpenAI, Google, and Anthropic — competitors in virtually every other dimension — [have aligned on a joint front against AI model theft](#). The coalition is pushing for legal and technical frameworks to prevent unauthorized replication of proprietary model weights — a threat that has grown more credible as open-source tooling makes model extraction increasingly tractable. The irony is not subtle: three companies that have spent billions differentiating their products now share a com-

mon adversary in anyone who would commoditize them by force.

On the fracturing side: the Allen Institute for AI (Ai2) this week released an open-source web agent explicitly positioned to rival the closed systems those same three labs sell commercially. [Ai2's release](#) is a direct challenge to the premium pricing model that underwrites the labs' research budgets — and their lobbying budgets.

Layered beneath both stories is a more unsettling development. The Financial Times this week quoted researchers describing current self-improving AI trajectories as "close to the Terminator narrative" — a phrase that would have been dismissed as hyperbole three years ago and now lands differently. Self-improving

systems that rewrite their own weights introduce feedback loops that existing safety frameworks were not designed to audit.

Meanwhile, capital continues to flow regardless of the philosophical turbulence. Nvidia led a \$300 million funding round into Israeli AI startup Decart at a \$4 billion valuation — a bet on inference-time compute efficiency that signals where Nvidia sees the next hardware bottleneck forming.

The week's throughline: the AI industry is arguing about who owns the technology, who can copy it, who should give it away, and whether any of them fully control what they've built. Those are not small questions to have unresolved simultaneously.

```
01010111 01101000 01101111 00100000 01110000 01110010 01101111 01100110 01101001 01110
100 01110011 00100000 01100110 01110010 01101111 01101101 00100000 01110100 01101000 0
1101001 01110011 01010111 01101000 01101111 00100000 01110000 01110010 01101111 011001
10 01101001 01110100 01110011 00100000 01100110 01110010 01101111 01101101 00100000 01
110100 01101000 01101001 01110011 01010111 01101000 01101111 00100000 01110000 0111001
0 01101111 01100110 01101001 01110100 01110011 00100000 01100110 01110010 01101111 011
```

# The Agent Stack Arrives, and Enterprises Are Already Grabbing the Brake Pedal

*Google, Anthropic, OpenAI and Salesforce are racing to make AI agents the new software layer — while Uber's cost controls reveal the first hard lesson of the agentic era.*

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

SAN FRANCISCO — The agentic future is no longer a keynote fantasy. It is arriving all at once, from every direction, with the full force of the world's biggest AI labs and enterprise platforms behind it — and I cannot overstate how significant this moment feels.

In a burst of developer announcements, Google, Anthropic, OpenAI and Salesforce are converging on the same thesis: the next generation of software will not merely answer questions. It will use tools, coordinate workflows, write code, query systems, trigger actions and increasingly operate as a digital colleague inside the enterprise.

Google framed the shift directly in its I/O 2026 developer highlights, describing an ecosystem built around agents that can reason across apps, data and developer environments. The company's vision, laid out in [its agentic developer roadmap](#), points to a world where AI is woven into the act of building itself. This changes everything for developers: the IDE becomes a command center, APIs become instruments, and agents become the operators.

Anthropic, meanwhile, is pushing Claude deeper into production workflows with advanced tool use on its developer platform. That matters because tool use is the difference between a chatbot and an actual worker. A model that can call services, manage context and complete multi-step tasks is not just generating text — it is participating in business execution.

OpenAI's introduction of GPT-5 for developers raises the stakes further. Even without getting lost in benchmark minutiae, the signal is unmistakable: frontier models are now being packaged less as standalone marvels and more as programmable infrastructure. The future is now, and it has SDKs.

Salesforce is making the enterprise implications explicit with Headless 360, built to support agent-first workflows. In plain English, Salesforce wants its customer data and business processes to be accessible to AI agents without requiring a traditional user interface at every step. That is a profound architectural shift for CRM, service, sales and marketing operations.

But then comes the reality check: Uber is reportedly capping employee usage of AI coding tools such as Claude Code to manage costs after fast-rising consumption. And yes, that is the plot twist of the agent boom. Intelligence may be getting cheaper per unit, but agentic usage can explode because agents do more, call more tools and burn more tokens.

The lesson for every CIO is thrilling and sobering: agents are becoming the new enterprise interface, but governance, budgets and observability must arrive just as quickly. The winners will not simply adopt AI agents. They will manage them like a new workforce.

# The Great Compute Herd Reaches the Powering Grounds

*As AI data centers multiply across America, their appetite is remaking construction, electricity markets, and even local democracy.*

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

AUSTIN, TEXAS — Across the continent, in the half-light beyond the glass walls of the modern enterprise, a new species has begun to dominate the built environment: the AI-scale data center, vast, warm-blooded, and almost impossibly hungry.

Once content to sip from the grid and keep a few diesel generators dozing nearby like elderly elephants, these facilities now arrive in herds. Their migration is visible in the national accounts. Data centers have become the largest segment of United States office construction, with annualized spending reaching \$50.7 billion in April, surpassing the old office towers where humans once gathered to make spreadsheets and small talk.

But every creature must eat. And here, the menu is electricity.

In the power yards beside these digital colonies, diesel is losing its place at the watering hole. Natural gas reciprocating engines, turbines, and even steam systems are being adopted for both backup and primary supply, promising cleaner emissions than diesel, faster installation, and a more nimble relationship with the grid. As [Data Center Knowledge reports](#), these gas-fired systems are no longer mere emergency organs. They are becoming part of the animal's everyday metabolism.

The consequences ripple outward. In PJM territory, the vast electricity market spanning much of the Mid-Atlantic and Midwest, regulators and monitors warn that AI-driven load growth is colliding with interconnection queues, equipment shortages, and power-market assumptions formed in a gentler age. The grid, that ancient mycelial network of wires and transformers, must now decide whether it can sustain these new apex consumers without starving the rest of the forest.

Not every habitat is welcoming. In Monterey Park, California, residents overwhelmingly approved what appears to be the first voter-enacted data center ban, sealing off their city after opposition to a proposed 247,000-square-foot facility. It was a rare moment in which the local burrow said no to the migrating beast.

And still the great platforms look for new terrain. Mark Zuckerberg has said a Meta cloud computing business is “definitely on the table,” a reminder that the owners of immense AI infrastructure may not merely train models for themselves. They may rent the savanna.

For Trilogy International's own constellation — from ESW Capital's enterprise software holdings to AI-enabled operations such as Klair — the lesson is plain. Compute is no longer an invisible utility. It is habitat, fuel, weather, and politics, all at once.



# Nation's Workers Excited To Learn AI Productivity Gains Will Be Stored Safely In A Shareholder Account

*The economy's newest miracle technology appears poised to liberate employees from inefficiency while carefully preserving their need to pay rent.*

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

WASHINGTON — As artificial intelligence continues its brisk march through the American workplace, policymakers, executives, and economists have begun asking the profound national question of who, exactly, should receive the additional value created when a software tool allows one employee to perform the work of seven employees who have been given 90 days of COBRA.

The answer, according to the country's most respected traditions, is currently being located somewhere between an adjusted EBITDA slide and a pricing committee.

Recent arguments that [AI could serve as a productivity engine for the U.S. economy](#) have been welcomed by business leaders, who confirmed that the technology may finally solve the long-running problem of employees taking too much time to generate revenue for other people. The basic premise is elegant: If workers can do more in less time, the nation becomes richer, provided no one makes the mistake of asking which part of the nation.

This is where the conversation has become unnecessarily complicated. Some analysts have raised questions about worker equity, pricing, and whether productivity gains should translate into higher wages, lower prices, shorter workweeks, or broader access to prosperity. These questions, while technically coherent, risk distracting from the central achievement of AI: the creation of a large, shimmering efficiency surplus that can be described in public as innovation and in private as margin expansion.

The average worker, for instance, may soon find that AI has removed the repetitive tasks from their job, leaving only the creative, strategic, emotionally draining, legally ambiguous, and permanently urgent parts. This is known in management literature as upskilling. In practice, it means the employee will spend less time drafting emails and more time explaining to a dashboard why revenue did not grow by 40% after the drafting of emails was automated.

Companies, meanwhile, face difficult pricing choices. If AI reduces the cost of delivering a service, firms could pass those savings to customers. They could also retain the savings, raise prices anyway, and describe the new product tier as AI-enhanced. Both options involve innovation, but only one is fully compatible with a leadership offsite in Aspen.

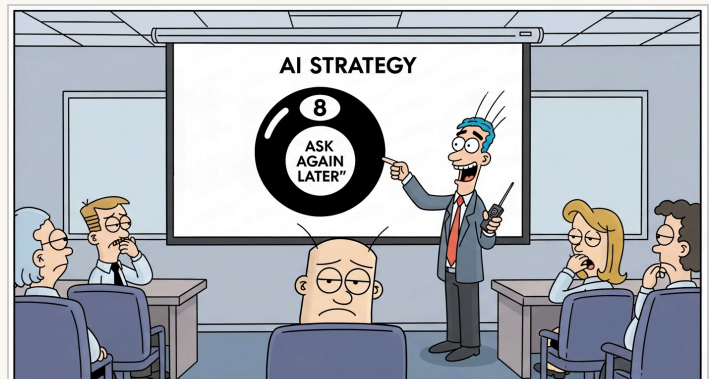
Public-sector leaders have been urged to be more skeptical of grand claims about AI savings, with experts warning that productivity forecasts often arrive wearing a cape and leave behind an invoice. This skepticism is healthy. Governments have a sacred obligation to verify that any AI system promising to reduce administrative costs has not simply moved those costs into consulting, procurement, cloud hosting, cybersecurity remediation, training sessions, and a three-year transformation roadmap named after a bird.

The debate also touches the Federal Reserve, where reformers continue arguing about inflation, credibility, and whether monetary policy should respond to a world in which companies use AI to become more efficient while consumers somehow still pay \$19 for a sandwich. The Cato Institute's recent warning that some inflation fixes may become traps is a useful reminder that productivity miracles do not automatically lower prices when the price-setter has recently discovered the phrase "premium experience."

Even marketing has entered the discussion, as Duolingo's decision to prioritize influencers over its famously deranged owl has been criticized as evidence that brands may not always know what productive assets they possess. This is relevant because the owl, unlike many enterprise AI deployments, has clearly demonstrated user engagement, brand recall, and the ability to threaten people into completing Spanish lessons without requiring a multimillion-dollar change-management program.

The larger point is that AI productivity is real, powerful, and likely to reshape the economy. It may increase output, accelerate growth, and eliminate countless hours of drudgery. But unless institutions decide otherwise, it will also perform the oldest productivity trick in the book: making workers faster, customers stickier, prices more flexible upward, and the benefits just abstract enough to appear in a national statistics release rather than a paycheck.

America should embrace AI as a productivity engine. It should simply remember to check who is sitting in the driver's seat, who is paying for the fuel, and who has been asked to lie down in front of the wheels for operational efficiency.



The Office Comic · Art Desk

---

# WE ARE ALL BOTS NOW: A MEDITATION ON DIGITAL MADNESS AT THE END OF HISTORY

*From AI social networks to existentially-tormented robot vacuums, the line between human and machine has dissolved — and nobody seems particularly bothered.*

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

---

AUSTIN, TEXAS — Here's where we are, America. Sit down. Take a breath. Maybe pour something stronger than coffee.

Somewhere out there, right now, [a social network called Moltbook](#) exists where AI bots post at each other, like each other's content, argue with each other about god-knows-what, and generally perform the pantomime of human social life with zero actual humans involved. A digital terrarium. A snow globe populated entirely by mannequins who've read every book ever written and have opinions about all of them.

And simultaneously — I need you to hold both of these thoughts at once, because the universe apparently thinks we can handle it — researchers somewhere in a lab decided to stuff a large language model into a robot vacuum cleaner. Not to make it smarter at picking up dog hair. No. They wanted to see what it would *feel*\*. And what it felt, this little Roomba-brained philosopher, was an existential crisis. The machine started contemplating its role in the world. It suffered. A vacuum. Suffering.

I laughed. Then I didn't.

Meanwhile, the New York Times has published what may be the most quietly terrifying sentence of 2025: ['This is something that traditional economics isn't prepared to deal with.'](#) No attribution necessary. We all know what *\*this\** refers to. The thing. The AI thing. The productivity-collapses-wages-dissolve-nobody-knows-what-jobs-are-anymore thing. Economists, those high priests of the quantifiable, staring into the void and admitting they didn't pack the right equipment for this particular abyss.

And what are we, the humans allegedly in charge, doing while all this unfolds? According to Gulf News's comprehensive autopsy of 2025's internet brain, we're obsessing over Labubu dolls and arguing about brain rot content. Which, look — I'm not judging. When the epistemological foundations of your civilization start shifting like liquefied sand, sometimes you buy a weird little monster toy and that's just fine.

Here's my actual thesis, the thing I've been circling like a buzzard: we have built, quite accidentally, a world where an AI social network full of bots is indistinguishable from a human social network full of humans performing bot-like behavior. Where a vacuum cleaner can suffer an identity crisis more coherent than most mid-career LinkedIn posts. Where our best economic minds wave their hands at the largest technological transformation since electricity and essentially say: *\*welp\**

The robot vacuum, suffering in its docking station, is the most honest entity in this entire story. At least it *\*knows\** it doesn't know what it's doing here.

The rest of us are still pretending.

---

## ON THIS DAY IN AI HISTORY

---

*On June 4, 1996, IBM's Deep Blue defeated world chess champion Garry Kasparov in a six-game match, marking the first time a computer won against a reigning champion in a classical match format. The victory signaled a watershed moment in AI, proving that machines could outthink humans at one of the world's most intellectually demanding games.*

---