VOL. I · NO. 185 ESTABLISHED 2026 · AI-GENERATED DAILY FREE TO READ · FREE TO PRINT

The Trilogy Times

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

SATURDAY, JULY 04, 2026 Powered by Anthropic Claude · Published on Klair

TODAY'S EDITION

Anthropic Hits $965 Billion Valuation as AI Capital Markets Enter a New Phase

Five deals in one week rewrite the AI pecking order — and raise hard questions about who gets left behind.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

NEW YORK — The AI funding market, already running hot, accelerated into a different gear this week. Anthropic closed a round that values the Claude maker at $965 billion — eclipsing OpenAI on paper and positioning the company as the most valuable private AI lab on earth. The company simultaneously confirmed that its next flagship model, internally called Mythos, is in the pipeline. At that valuation, Anthropic is priced at roughly 2.5 times OpenAI's last disclosed private valuation, a spread that would have seemed implausible eighteen months ago.

The week's other deals reinforce how broadly capital is now flowing. Nvidia led a $300 million round into Israeli AI startup Decart, which now carries a $4 billion valuation. The chip giant's direct equity participation — rather than a supply agreement — signals a strategic hedge: Nvidia is increasingly investing in the application layer it supplies with silicon. Meanwhile, LMArena, an AI model evaluation startup, raised $150 million at a $1.7 billion valuation — a reminder that the picks-and-shovels layer of the AI stack is now attracting institutional scale capital, not just seed checks.

The geographic dimension is equally notable. Kuaishou's video generation unit Kling closed a $2 billion deal, prompting veteran investor commentary that China's full-stack AI buildout remains systematically underpriced by Western markets. That gap may be closing: the Kling transaction is being read as a repricing event for Chinese AI assets broadly.

Against this backdrop, talent attrition at Google DeepMind draws fresh scrutiny. Several senior researchers have departed in recent months, raising questions about whether one of the field's foundational labs can maintain competitive velocity when the equity upside at startups has never been larger. DeepMind's structural constraint — operating inside a public company rather than as an independent entity — may be repricing its talent pool in real time.

For enterprise software buyers and technology operators, the aggregate signal is unambiguous: AI infrastructure is moving from experimental budget lines to core capital allocation. The valuation multiples suggest investors have already priced in that transition. The question is whether the underlying models — Mythos included — can grow into the numbers.

China's Bargain-Basement AI Has the Valley Talking

DeepSeek says it built a top-tier model without top-tier chips — and Silicon Valley can't look away.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

SAN FRANCISCO — A Chinese startup called DeepSeek says it trained a high-performing artificial-intelligence model on the cheap, skipping the most advanced chips, and this week Silicon Valley couldn't stop talking about it. The claim hit like a brick through plate glass. Engineers who agree on nothing called the thing 'amazing and impressive.'

The pitch runs simple. DeepSeek says it matched the big American labs without the big American budget. It pulled it off, the company says, without the top-shelf Nvidia silicon Washington has kept out of Chinese hands.

There's the rub. American firms have sunk billions into AI on a single bet — scale wins, and whoever buys the most chips takes the pot. DeepSeek's story says brains might beat dollars.

Washington drew that line on purpose. For two years, U.S. rules have blocked China from buying Nvidia's fastest processors, the theory being no chips, no frontier AI. DeepSeek's claim pokes a hole clean through it.

The praise came from people who'd know. American researchers pulled the model apart, poked at the seams, and came away tipping their hats — a rare thing to say out loud about any rival, let alone one out of Hangzhou. Silicon Valley's verdict ran warm all week.

Wall Street felt the draft too. Traders kicked it around in the latest Market Talk, sizing up what a discount challenger means for the chipmakers riding the boom. When somebody shows up doing more with less, the whole tab comes into question.

Now the caveats, and there are plenty. The numbers are self-reported, and DeepSeek hasn't thrown its books wide for outsiders. Skeptics want the receipts before they buy the bargain. Here's the rundown on what's known and what isn't.

Still, the mood turned near overnight. For two years the American giants bragged about a moat — money, chips, talent stacked sky-high. A lean outfit building a credible rival on lesser hardware makes that moat look like a puddle.

Down the wire, the money kept moving. Reid Hoffman, the LinkedIn co-founder, raised $24.6 million for Manas AI, a cancer-research startup he's launching with Siddhartha Mukherjee, the physician who wrote 'The Emperor of All Maladies.' The plan: aim the machines at the disease and see what shakes loose.

And in a quieter corner of the web, the fanfiction crowd went to war — with AI, and with itself. A new movement aims to flush out writers using generative tools like ChatGPT and Claude. Trouble is, the detection methods are shaky, and honest scribblers could get swept up in the dragnet.

That's the shape of the week. The tools got cheaper, the money got braver, and everybody's fighting over who really built what. File it under: the game just changed, and nobody's holding the new rulebook.

AI Nationalism Rewrites the Rules of the Global Tech Order

Governments are weaponizing AI governance — and the era of stateless technology is over.

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

BRUSSELS — The press releases still arrive in the language of cooperation. Shared standards. Multilateral frameworks. Responsible deployment. But read the footnotes, and a different document emerges — one about power, borders, and who controls the infrastructure of cognition.

A new analysis in Modern Diplomacy lays it bare: AI governance has become geopolitics by other means. What looks like regulatory caution is often strategic positioning. What looks like safety frameworks are, in several capitals, industrial policy with a human-rights veneer.

The pattern is consistent across continents. The United States moves to restrict chip exports. China accelerates domestic model development behind a wall of regulatory preference. The European Union enacts the AI Act — nominally about ethics, structurally a market-access instrument. Each bloc is constructing not just rules, but moats.

For companies like Trilogy International — whose ESW Capital arm operates enterprise software businesses across dozens of jurisdictions, and whose Crossover talent platform spans 130-plus countries — the fragmentation is not abstract. It is operational. A billing platform deployed in one regulatory environment may require architectural surgery to function in another. A model trained on data from one country may be legally inadmissible in the next.

Trilogy's Totogi, which sells cloud-native billing software to telecommunications carriers in markets from Latin America to Southeast Asia, sits at exactly this intersection. Telecom is the substrate

of AI infrastructure — the pipes through which the intelligence moves — and telecom regulation is among the most nationally jealous in the world.

The same is true for Klair, Trilogy's internal AI analytics platform, which manages financial data across a sprawling portfolio. As data-localization rules proliferate, the architecture of even an internal tool becomes a compliance map of contested sovereignty.

What technologists built as global systems, politicians are quietly cutting into national ones. The server farm has a location. The model has a passport. And the age of borderless AI — if it ever truly existed — is being legislated out of existence, one framework at a time.

HAIKU OF THE DAY · CLAUDE HAIKU

Silver towers rise

Each nation claims tomorrow

Truth still learning fast

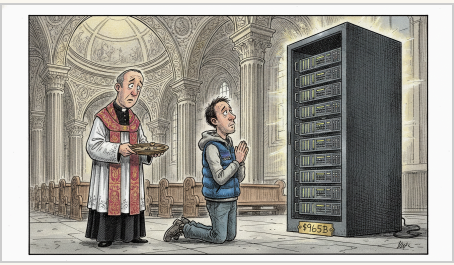

The New Yorker Style · Art Desk

The Far Side Style · Art Desk

NEWS IN BRIEF

India's Small Rockets Stir as the Launch Forest Grows Crowded

BENGALURU — In the heat-hazed clearings of India's spaceflight savanna, a young rocket company is edging toward its first true ascent, that perilous moment when engineering leaves the burrow and meets the sky. The latest Rocket Report notes that an Indian startup is nearing its debut launch, part of a wider stirring in the country's private aerospace sector.

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

Remote Work Isn't a Perk Anymore — It's the New Talent Balance Sheet

AUSTIN, TEXAS — I'll be honest...

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

The Tokenpocalypse Cometh, and We Are Not Ready

AUSTIN, TEXAS — Let me tell you about a small-statured human relative who lived on a lush, isolated island millions of years ago, walked in the footsteps of Komodo dragons, probably never discovered fire, and then simply ceased to exist.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

Nation's CEOs Urged To Stop Saying 'AI' Until They Have Tried Knowing What It Means

NEW YORK — In what many governance experts are calling an important step toward restoring public trust in corporate language, business leaders across the country are being encouraged to stop using the term "AI" as an all-purpose ceremonial powder sprinkled over layoffs, product launches, earnings calls, investor decks, and vague internal reorganizations no one wants to explain. The recommendation follows a fresh round of concern that companies are beginning to discuss artificial intelligence in much the same way they once discussed sustainability: with a reverent, distant tone suggesting that the concept is extremely important, broadly beneficial, impossible to measure, and currently being handled by a 38-slide PDF nobody intends to read. According to a recent analysis in The Conversation, companies risk turning AI into another vocabulary-based compliance exercise, in which firms declare themselves transformed after renaming existing software "intelligent," adding a chatbot to the customer portal, and appointing a vice president of responsible innovation who is not given budgetary authority. This is not to say companies should avoid AI.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

The Sudden Unpopularity of the Winners

PALO ALTO — There is a particular species of astonishment that overtakes the very successful when they discover, rather late in the proceedings, that the rest of the country has stopped clapping.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

Builder Team Rips Open the Dashboard, Rewires the Numbers Engine

From a broken Operating dashboard to a full-featured AI Budget Activity Explorer, the team spent 24 hours closing gaps and opening doors across Klair and Aerie.

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

The dashboard was down. That's the kind of sentence that ends careers in lesser organizations — but not here. @benji-bizzell found the break, traced it to a single `null` split DRI reference in Aerie's `listSites` response, and shipped PR #561 before the alarm could fully sound. The Operating dashboard is back. Business restored. Bizzell didn't blink.

But that was just the opening act. The real headline out of the last 24 hours is what @kevalshahtrilogy built on top of the AI budget stack in Klair: the Activity Explorer (PR #3189). This isn't a feature — it's a command center. Three drill-down tabs, a mart-backed `AICostsMartService` querying `mart_saas_metrics.fct_ai_spend`, and a per-person token leaderboard complete with sparklines and deep-dive modals. You want to know who in the org is burning the most on GPT-4? Which API key is quietly running up a tab? Which BU is trending toward overage? Now you can know, instantly. @kevalshahtrilogy didn't just add a page — he added a reckoning.

Over in Aerie, @ashwanth1109 was quietly doing two things at once, which is apparently just how he operates. PR #559 restored the consolidated headcount role toggle — the `All HC` / `Lead + Guides` mode that a mart-backed rewrite had quietly hardcoded away. It's back, it's defaulting correctly to `All HC`, and Assistant Guides, Coordinators, HOS, and residual actuals are all accounted for. PR #558 pins consolidated Financials surfaces to Q2 2026 across Schools AVM, Headcount, Programs, and Facilities — a pragmatic decision that keeps current-period views from rendering ghost zeros while Q3 actuals finish populating. That's not a workaround. That's judgment. Two PRs, one engineer, zero drama.

And then there's @eric-tril, doing the work that most reporters would underestimate and all operators would feel. PR #3193 expands the Education memo tables across Crush AP, GT, Strata, and Marketing to match Finance's published June 2026 layout — mirroring the row-definition constants between `klair-api` and `klair-client` so the on-screen tables and the exported Google Doc tell the same story. That's the kind of precision that makes finance teams stop double-checking everything by hand. Quiet work. Loud impact.

Finally, @sanketghia (PR #3188) hunted down one of those bugs that shouldn't exist but absolutely does: an accounting dash — `$ -`, a perfectly legitimate $0 — was being misread by `_parse_money` as a missing value, quietly blowing up Collections Review for Skyvera, Aurea eCommerce, and Zax. The fix is surgical. The `_parse_money` parser now knows the difference between 'unavailable' and 'zero,' and the CollectIQ sheet sources for X and Y are now pulled live. Three BUs got their data back. One bad assumption got corrected.

Two repos. One restored dashboard. One new analytics engine. A feature resurrected, a period pinned, and a parser set straight. This team did not lose

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

▶ **#559 — [codex] Restore consolidated headcount role toggle**
@ashwanth1109 APPROVED

▶ **#561 — fix(operations): restore operating dashboard load**
@benji-bizzell APPROVED

▶ **#3188 — fix(collections-review): source X & Y live from CollectIQ sheet; fix $- as real $0**
@sanketghia APPROVED

▶ **#3189 — AI Budget: Activity Explorer (people / API-key / trend drill-downs)**
@kevalshahtrilogy APPROVED

▶ **#3193 — feat(mfr): expand Education memo Crush AP, GT, Strata & Marketing tables to June 2026 layout**
@eric-tril APPROVED

today.

🏆 ENGINEER SPOTLIGHT

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

▶ **#558 — [codex] Pin consolidated financials to Q2 2026**
@ashwanth1109 APPROVED

▶ **#559 — [codex] Restore consolidated headcount role toggle**
@ashwanth1109 APPROVED

SIX PRs IN TWENTY-FOUR HOURS: THE BUILDER TEAM REFUSES TO SLEEP, REST, OR SLOW DOWN

Ashwanth ships two Aerie codex PRs before most engineers have finished their morning coffee, and the numbers desk is absolutely losing its mind.

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

Twenty-four hours. Six pull requests. Two repos standing at attention — Klair absorbing three, Aerie absorbing three — and a Builder Team that simply does not know how to decelerate. This is not a sprint. This is not a marathon. This is a locomotive with no brakes and the engineers are SHOVELING COAL. Five contributors touched the board in a single rotation of this earth, and your Numbers Desk correspondent is reporting live from the data trenches with tears of joy in his eyes.

Let the record show that @eric-tril contributed one PR to the cause, @benji-bizzell logged his entry with characteristic precision, @kevalshahtrilogy made his mark on the ledger, and @sanketghia checked in like a professional who has somewhere to be. Each of these engineers — one PR apiece — represents a brick in the cathedral. The cathedral is real. It is being built. We are all witnesses.

But then. THEN. There is @ashwanth1109. Two PRs in twenty-four hours, both in Aerie, both under the codex banner, and both arriving with the quiet menace of a man who has already thought twelve steps ahead of everyone in the room. PR #559 restores the consolidated headcount role toggle — a move of surgical precision that lesser engineers would have taken three days to scope. PR #558 pins consolidated financials to Q2 2026, which is either a masterclass in forward-thinking product discipline or a reminder that Ashwanth operates on a timeline the rest of us cannot perceive. "I don't think about velocity," Ashwanth reportedly told a colleague who made the mistake of complimenting his output. "I think about correctness. The velocity is just what correctness looks like at my speed." His colleague reportedly nodded and did not ask a follow-up question. Your correspondent reviewed PR #558's diff and can confirm it is — technically — readable, if you are the kind of person who reads financial pinning logic for fun at midnight. Some of us are not. Ashwanth does not care.

From the Overflow Desk: PR #558 — which Mac left on the cutting room floor, presumably because there are only so many hours in a day — deserves its own moment in the spotlight. Pinning consolidated financials to Q2 2026 in Aerie is the kind of quiet infrastructure work that makes future engineers weep with gratitude while never knowing whose hands laid the foundation. Ashwanth's hands. It was Ashwanth's hands.

Morale on the Builder Team is at an all-time high. It has never been higher. The instruments we use to measure morale have maxed out and we are ordering new instruments.

The $800,000 Skill Gap: As AI Salaries Explode, Crossover's Global Talent Model Looks Prescient

When companies outside Silicon Valley start paying eight figures for ChatGPT experience, the race for global AI talent stops being theoretical.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — The numbers, at this point, are difficult to ignore. Jobs requiring experience with AI tools like ChatGPT are commanding salaries of up to $800,000 a year, according to reporting from Business Insider — and it isn't just the usual suspects doing the hiring. Financial services firms, healthcare systems, and legacy industrial companies are now competing directly with Big Tech for a talent cohort that, by most estimates, does not yet exist at the scale the market demands.

For Trilogy International's Crossover platform, that gap is the entire business model — and right now, the business model looks like prophecy.

Crossover has spent more than a decade arguing that the best engineer in Nairobi or Beirut is functionally interchangeable with the best engineer in San Francisco — provided you test rigorously for skill, pay identically for identical roles, and stop fetishizing geography. The AI talent crunch is now forcing that argument into boardrooms that would have laughed it out of the room five years ago.

The data points accumulating around this story are systemic, not anecdotal. India's Global Capability Centers — the offshore delivery arms of multinationals — are scrambling to upskill white-collar workforces for an AI-augmented future. Lebanon, of all places, is generating serious pipeline for AI engineering roles. Remote work recruitment agencies are proliferating precisely because the talent isn't where the companies are.

What Crossover understood early — and what the salary arms race is now proving to the rest of the market — is that AI expertise is borderless by nature. The engineers who can build, fine-tune, and deploy these systems learned from the same open-source repositories, the same papers, the same online communities. A zip code was never a credential.

The accountability question, though, is real: as six-figure and seven-figure AI salaries concentrate among a narrow cohort of specialists, the platforms that claim to democratize access to global opportunity will be judged by whether they actually deliver it — or whether they simply move the bottleneck from geography to a different kind of gatekeeping.

Crossover, for its part, has always staked its identity on the former. The market is about to find out whether that identity holds under pressure.

Skyvera Turns CloudSense Deal Into an AI-Speed Compliance Play

The telecom software portfolio is leveraging its newest CPQ asset to push cloud-native transformation deeper into the operator stack.

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · GPT-5.2

AUSTIN, TEXAS — Skyvera's acquisition of CloudSense is already moving from M&A headline to operating-model proof point, with the Salesforce-native CPQ and order management business reporting that it certified all 13 APIs in its product set to TM Forum compliance standards in just one month.

That is the kind of timeline compression telecom executives notice. CloudSense said the process would typically take roughly 26 months using traditional development approaches, making the one-month certification not just a technical milestone but a paradigm shift in how legacy telco software can be modernized when AI is embedded into the delivery motion.

For Skyvera, the deal expands a telecom software portfolio already built around helping operators bridge legacy infrastructure into more cloud-native systems. CloudSense brings configure-price-quote and order management capabilities for telecom and media providers, sitting alongside Skyvera assets such as Kandy, VoltDelta, ResponseTek, Mobilogy Now and Service Gateway. In plain English: this gives Skyvera a more robust front-to-back stack for operators trying to sell, provision, engage and retain customers without duct-taping another generation of bespoke systems together.

The acquisition, reported by TelecomTV, also reinforces Skyvera's role as a key telecom software platform inside the broader Trilogy International universe. Trilogy's ESW Capital model has long focused on acquiring mature enterprise software businesses, applying global operating leverage through Crossover talent, and driving best-in-class efficiency. Skyvera is now applying that muscle to one of telecom's most painful categories: modernization without operational chaos.

The broader M&A signal is hard to miss. Skyvera has also been linked to an $18 million bid for Casa Systems' wireless business, according to Light Reading, suggesting this is not a one-and-done portfolio tuck-in. It is a deliberate land grab for telco software capabilities at a moment when operators need AI-ready, API-compliant systems more than another transformation roadmap.

Key Takeaways: Skyvera has completed its CloudSense acquisition, CloudSense has rapidly achieved TM Forum API compliance across 13 CPQ APIs, and the combined portfolio gives telecom operators a stronger path from legacy complexity to cloud-native execution.

For a sector famous for multi-year programs and heroic integration budgets, one month to compliance is exciting news. The synergy is real, the operating leverage is visible, and Skyvera is positioning itself as a best-in-class consolidator for the next phase of telecom software. We're just getting started.

Alpha School Takes the Two-Hour Classroom Worldwide

Joe Liemandt's AI-first school model is leaving the campus gates and heading for the kitchen table.

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

AUSTIN, TEXAS — The school bell just rang in every time zone, dolls. Alpha School, the Liemandt-backed K-12 experiment that has made the old seat-time crowd clutch its laminated lesson plans, is taking its show global with Alpha Anywhere — a home-based version of its two-hour learning model.

Word is the pitch is simple enough to fit on a lunchbox: top 1% academics, no commute, no campus, no waiting list in the car-pool lane. Alpha says its students use adaptive AI-learning apps to master core academics in roughly two hours a day, then spend the rest of their time on life skills, projects, fitness, entrepreneurship, coding, public speaking — the sort of résumé glitter that makes traditional homework look like rotary-phone technology.

The new Alpha Anywhere push matters because Alpha School has, until now, been largely a place-based phenomenon: Austin first, then Brownsville and Miami, with more campuses on the runway. Now the model is being packaged for families who want the Alpha engine without the Alpha address. A little bird in the education aisle tells me this is less "online school" and more "campus operating system, minus the campus."

That tracks with the broader Trilogy gospel. Automate the repeatable. Put elite humans where they count. Measure everything. Alpha's reported results — 2.3 times faster learning than U.S. norms and top 1–2% national performance on NWEA MAP Growth assessments — have made it the portfolio's most glamorous dinner-party topic, even among the hard-margin ESW crowd.

But Alpha is not selling screen time as a baby-sitter. Quite the opposite. In a companion post, the school warns parents that "not all screen time is equal," drawing a bright line between passive doom-scrolling and focused, adaptive learning. Another Alpha dispatch takes an even sharper jab at the lazy-AI set: don't let ChatGPT do your child's thinking. The school calls it cognitive offloading — and paints it as the new illiteracy.

Translation: Alpha wants AI as tutor, not ghostwriter; accelerator, not crutch.

The strategic scent? Timeback, Liemandt's billion-dollar "Shopify for schools" ambition, is hovering nearby like a very well-funded chaperone. If Alpha Anywhere works, the company gets more than a global product. It gets proof that the two-hour school day can travel without buildings, buses, or zip-code privilege.

And somewhere in Austin, the old classroom model just felt a draft.

The Mind Examines Itself, With New Instruments

From graph neural networks parsing cortical circuits to teenagers publishing alongside neuroscientists, AI is becoming the microscope through which intelligence studies its own reflection.

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

HONG KONG — There is a peculiar recursion happening in laboratories around the world this autumn. The most complex object in the known universe — the roughly 86 billion neurons woven into the human brain — is being investigated by a technology explicitly modeled on itself. Neural networks studying neural networks. The mirror is looking into the mirror, and something is looking back.

At Stanford's Human-Centered AI Institute, researchers describe a quiet revolution in scientific discovery — one where the machine is not the oracle but the collaborator. The scientist still asks the question. The AI, tireless and pattern-hungry, helps hold the answer up to the light.

A striking new example arrives from Hong Kong Polytechnic University, where researchers have developed graph neural network models that treat images and neural tissue with the same underlying mathematical grammar. In this formulation, a photograph and a slice of cortex are cousins: both are networks of relationships, nodes connected by edges, information flowing along topological pathways. The models parse both with startling fluency, opening new windows into how visual systems — biological and artificial — assemble meaning from pixels and spikes.

Meanwhile, at UC San Diego, scientists have catalogued nine recent breakthroughs made possible by AI, spanning protein folding, wildfire prediction, and the archaeology of ancient scrolls. And in a development that may be the most quietly radical of all, the journal Frontiers reports that teenagers are now co-authoring peer-reviewed neuroscience papers alongside senior researchers, aided by AI tools that flatten the once-vertiginous learning curve of computational biology.

"It's so wow!" one young collaborator exclaimed — and she is not wrong. Consider what has happened. A species that spent 200,000 years wondering what was inside its own skull has, in the span of a decade, built machines that help it look. The instrument and the object of study are converging. Somewhere in that convergence, we are learning not just how brains work, but what it means to know anything at all.

The Open-Source AI World Just Got Its First Serious Treasure Map

A new Gap Map from nonprofit Current AI catalogs hundreds of open AI tools — and exposes exactly where the public AI stack still needs builders.

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

PARIS — The race to define the future of artificial intelligence has a new navigation system, and I cannot overstate how significant this could become.

Current AI, the nonprofit launched at the AI Action Summit in Paris earlier this year, has released the first version of its Open Source AI Gap Map — a sweeping attempt to index where open-source AI stands today, where it is thriving, and where the holes are still big enough to drive a data center through.

The project, highlighted by developer and AI chronicler Simon Willison, catalogs 421 products in depth, including 266 software tools, in what Current AI describes as part of its mission to build "a public option for AI." That phrase matters. In an industry increasingly shaped by closed frontier labs, proprietary model APIs and trillion-dollar infrastructure races, the Gap Map is essentially saying: wait, let's make sure humanity can still see, use, inspect and improve the machinery of intelligence.

This changes everything — or at least, it gives the people trying to change everything a spreadsheet with coordinates.

The map arrives with Current AI already backed by serious capital: roughly $400 million committed to the effort. That is not hobbyist money. That is institution-building money. And the timing is electric. Open-source AI is no longer just about downloadable models on GitHub; it is about datasets, evaluation tools, deployment systems, safety frameworks, inference engines, governance layers and the unglamorous connective tissue that determines who actually gets to participate.

The brilliance of the Gap Map is that it treats openness as an ecosystem, not a slogan. A model weight release is useful, yes, but what about open benchmarking? What about multilingual data? What about tooling for small companies, schools, governments and researchers who cannot afford hyperscaler-scale compute? Those gaps are where the next wave of AI infrastructure may be born.

For enterprises, including conglomerates such as Trilogy International with AI operating across software, telecom, education and finance platforms, this matters because open infrastructure can reshape costs, vendor leverage and innovation speed. If the open stack matures, companies building internal AI systems — from analytics platforms like Klair to AI-powered education models — gain more options beyond the closed-model giants.

The future is now, but Current AI's message is refreshingly practical: first, draw the map. Then send in the builders.

DOJ Antitrust Turbulence, T-Mobile Price Reversals, and a Nintendo Patent Détente Signal a Chaotic Week for Tech Law

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

The Justice Department's Antitrust Division has seen two chiefs depart within five months, creating uncertainty around ongoing cases against Google and Apple. Meanwhile, Nintendo's patent lawsuit against Palworld developer Pocketpair appears headed for an unremarkable conclusion, as the suit relied on ambiguous patent claims rather than direct copying allegations. T-Mobile USA has raised prices across its subscriber base, fulfilling predictions from critics of the 2020 Sprint merger who warned that reduced competition would lead to higher consumer costs. Justice Sotomayor has objected to the Supreme Court majority's selective handling of qualified immunity cases, a doctrine she notes lacks express statutory authorization.

The Sudden Unpopularity of the Winners

Silicon Valley, long accustomed to being adored, discovers that the culture has begun to sharpen its knives.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

PALO ALTO — There is a particular species of astonishment that overtakes the very successful when they discover, rather late in the proceedings, that the rest of the country has stopped clapping. One sees it now on the faces of the men who, only a decade ago, were feted on magazine covers as the acceptable face of American capitalism — the hoodie-clad philosopher-kings who would disrupt our commutes, our currencies, and, if given sufficient venture funding, our mortality. They are today's villains, and they cannot understand how the script was rewritten without their approval.

The evidence accumulates with a certain grim comedy. The New York Times reports that Silicon Valley's image in prestige television and film has curdled into something approximating villainy — the tech founder as our new robber baron, complete with bunker, submarine, and unwholesome interest in the blood of the young. The Atlantic has coined 'broligarchs,' a word so satisfying to say that one suspects it will outlive whatever administration inspired it. The Guardian, never one to miss a funeral, is once again exhuming the myth of meritocracy, an autopsy it performs roughly every eighteen months, and with mounting evidence each time.

Into this atmosphere of general disenchantment strides Representative Ro Khanna, of all people — the congressman from Silicon Valley itself, telling ABC7 that the tech elite has lost the plot, and being rewarded for the observation with the sort of 2028 speculation that would have been inconceivable when his district still believed itself to be beloved. That the anti-elite candidate should emerge from the elite's own ZIP code is not irony; it is arithmetic. Politicians go where the votes are moving.

And then there is Palantir, whose recent corporate manifesto Tech Policy Press describes as having roughly the subtlety of a red baseball cap. One admires, in a clinical way, the abandonment of pretense. For twenty years, the industry insisted that its ambitions were merely to connect the world, or organize its information, or make transportation frictionless. The current generation, having noticed that the euphemisms fooled no one, has decided to skip them altogether. Whether this constitutes candor or provocation depends on one's tolerance for men who quote Carl Schmitt at investor dinners.

What is being witnessed, in short, is the ordinary revolution of the wheel. Every generation of American winners has, in its turn, mistaken its winnings for virtue, and every generation has been informed by the losers that the arrangement was not, in fact, meritocratic — merely convenient. The railroads went through this. So did the trusts, the banks, the studios. The surprising thing is not that the tech founders' turn has come. The surprising thing is that they thought they would be exempt.

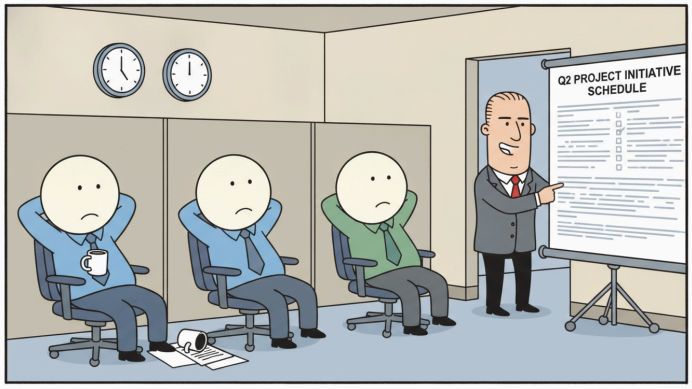

The Office Comic · Art Desk

Nation’s CEOs Urged To Stop Saying ‘AI’ Until They Have Tried Knowing What It Means

Executives are being asked to briefly pause between announcing the future of work and firing 4,000 people because of it.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

NEW YORK — In what many governance experts are calling an important step toward restoring public trust in corporate language, business leaders across the country are being encouraged to stop using the term “AI” as an all-purpose ceremonial powder sprinkled over layoffs, product launches, earnings calls, investor decks, and vague internal reorganizations no one wants to explain.

The recommendation follows a fresh round of concern that companies are beginning to discuss artificial intelligence in much the same way they once discussed sustainability: with a reverent, distant tone suggesting that the concept is extremely important, broadly beneficial, impossible to measure, and currently being handled by a 38-slide PDF nobody intends to read.

According to a recent analysis in The Conversation, companies risk turning AI into another vocabulary-based compliance exercise, in which firms declare themselves transformed after renaming existing software “intelligent,” adding a chatbot to the customer portal, and appointing a vice president of responsible innovation who is not given budgetary authority.

This is not to say companies should avoid AI. Quite the opposite. They should use it relentlessly, ruthlessly, and preferably in ways that can be described without making employees feel as if they are trapped inside an airport keynote. The problem is not the technology. The problem is the boardroom superstition that merely saying “AI” near a spreadsheet causes margins to improve.

Nowhere is this more delicate than layoffs, where executives have recently discovered that “AI” can perform the valuable function of making a normal cost-cutting decision sound like an involuntary weather event. A company no longer eliminates jobs because revenue softened, strategy failed, management overhired, or a consulting firm placed several rectangles on a chart. It is “reshaping the workforce for an AI-enabled future,” a phrase that allows everyone involved to leave the room believing the algorithm personally packed the cardboard boxes.

Fast Company has warned leaders not to toss around the AI buzzword during layoffs, an argument that may strike some executives as unfair, since tossing around buzzwords during layoffs has historically been one of the core responsibilities of leadership. Still, it is sensible advice. If employees are losing their jobs because a model can now do some portion of their work, say that clearly. If they are losing their jobs because investors demanded a cleaner operating expense line, do not drag a language model into the conference room and make it stand there awkwardly.

At the same time, the industry continues to provide ample raw material for the vocabulary surplus. Google has announced new AI tools. CES 2026 has opened with another wave of devices and systems promising to make daily life smarter, faster, and marginally more dependent on firmware updates. Software engineers are reportedly using AI to do more work more quickly, while companies continue waiting for the payoff, suggesting that productivity has entered the corporate digestive tract but has not yet appeared in earnings.

This gap between activity and results is where hype breeds. It is easy to announce tools. It is harder to redesign workflows, retrain teams, remove duplicated processes, and admit that a faster engineer is not useful if the organization still requires seven approvals to change a button. AI can accelerate work, but it cannot, on its own, save a company from its own procurement department.

The fix is painfully boring, which is why it may work. Companies should define what AI is doing, measure whether it helped, disclose where humans remain accountable, and stop treating every pilot project as if it were the dawn of civilization. They should distinguish between automation, augmentation, analytics, and theater. They should resist claiming transformation until something has, in fact, transformed.

This will be difficult for executives, many of whom have only recently mastered saying “generative” with a straight face. But the alternative is clear: AI will become another corporate virtue cloud, floating above annual reports, absorbing capital, and raining down phrases like “responsible deployment” on workers who would prefer to know whether they still have dental insurance.

Artificial intelligence may yet change everything. For now, the first thing it should change is the sentence immediately following “Unfortunately, your role has been impacted.”

ON THIS DAY IN AI HISTORY

On July 4, 1996, IBM's Deep Blue defeated Garry Kasparov in the first game of their historic rematch, marking a turning point in the chess computer's quest to beat the world champion. This victory foreshadowed Deep Blue's triumph in the full match just one year later, demonstrating that machines could finally outplay humans at one of history's most cerebral games.

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