

The Trilogy Times

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

THURSDAY, OCTOBER 08, 2026

Powered by the TrueFoundry AI Gateway · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

BUILD YOUR ADVERSARY, THEN SWEAR AN OATH TO IT: AI GIANTS FACE NYC GRILLING AS DEEPMIND BLEEDS TALENT

DEVELOPING

Four witnesses under oath, four DeepMind founders out the door, one researcher says the industry is arming its own enemy.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE SONNET

NEW YORK — Lawyers for OpenAI, Anthropic, Google and Meta raised their right hands before New York City lawmakers Tuesday. An AI researcher in the same room didn't wait for cross-examination to land the gut punch. "We are racing to build and grow our own adversary," the researcher told the hearing, and the room went quiet.

City Council called the four companies in to testify under oath on AI safety, the kind of hearing that usually produces careful non-answers and lawyered silence. This one came with a body count. Hours before witnesses sat down, [Google DeepMind lost four founding-era leaders in a single day](#), a clean sweep of people who were in the building when the lab was still a London startup chasing Go championships.

Timing like that doesn't read as coincidence inside a hearing room. Lawmakers pressed the panel on what safeguards exist when the people who built the

guardrails start walking out the door. Nobody on the panel had a tidy answer, which is itself an answer.

Meanwhile the same companies insist they're cooperating on the exact danger the researcher flagged from the witness table. Bloomberg News reported, and Reuters confirmed, that [OpenAI is now working with Anthropic and Google on AI safety](#), trading notes on risk even as all three race each other for enterprise customers, talent, and compute. Call it cooptation if the lawyers insist. Call it hedging if you're the researcher who testified an hour earlier.

Meta picked the same week to launch Muse, a new enterprise AI platform aimed straight at corporate deployment — chatbots, agents, workflow tools sold to the same businesses whose employees are asking whether any of this is safe. The product rollout landed on the news wire the same day Meta's own counsel sat under oath answering questions about ex-

actly that. Nobody at Menlo Park scheduled it that way on purpose, probably.

The hearing produced no subpoenas, no fines, no new law. New York City doesn't regulate frontier AI models and everybody in that chamber knew it going in. What the hearing produced was a transcript, a researcher's line about adversaries, and four departing DeepMind veterans who didn't stick around to watch the testimony.

Trilogy International's portfolio sits several steps removed from the frontier-model arms race — ESW Capital runs enterprise software, not foundation models, and Alpha School uses AI tutors rather than builds them. But the industry's safety math gets stress-tested in rooms like the one in New York Tuesday, long before it ever reaches a classroom or a CRM license. The wire will keep watching the exits at DeepMind. Four in one day is not a trend yet. It is, at minimum, a headline.

Capital Keeps Finding AI, Even the Startups Meta Didn't Want

A record quarter for billion-dollar rounds shows investors still writing checks faster than regulators can ask questions.

BY DR. CHEN WEI, TECHNOLOGY
CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — Manus, the Singapore-headquartered AI startup with Chinese roots, closed a \$500 million round this week — its first financing since a proposed acquisition by Meta collapsed amid U.S. national-security scrutiny of Chinese-linked AI firms. The deal, [reported first by CNBC](#), values the agentic-AI firm well above the figure Meta was said to have offered before political pressure killed the talks.

The round is one data point in a broader pattern. Crunchbase's Q3 2026 tally found a record number of billion-dollar venture rounds globally, the fourth consecutive quarter of acceleration since the generative-AI boom began in late 2022. The firm's [data](#) shows AI companies accounted for the overwhelming majority of those megadeals, a concentration not seen since the dot-com-era telecom buildout of 1999-2000, when a handful of sectors absorbed most late-stage capital before the correction.

The pattern extends upstream to the fund managers themselves. The venture firm behind chip startup Groq — whose custom inference silicon competes with Nvidia on cost-per-token economics — is now raising toward a \$10 billion vehicle, according to the Wall Street Journal, a sum that would rank among the largest sector-focused funds ever raised outside buyout shops.

None of this capital formation has slowed scrutiny of what the underlying products actually do. A children's safety nonprofit this week flagged failures in OpenAI's teen-oriented ChatGPT mode, finding the tool still completes homework it is supposed to merely tutor. The juxta-

position is instructive: balance sheets are compounding faster than the guardrails.

The Server Farm Has a Flag Now

As Washington and Beijing circle a summit, Brussels bets its sovereignty on silicon it doesn't yet own.

BY ELEANOR CROSS, FOREIGN
CORRESPONDENT · CLAUDE SONNET

WASHINGTON — The cable runs from a data center outside Ashburn, Virginia, to a server farm in Shenzhen, and somewhere between the two a new kind of border has formed. Not a line on a map. A line in a model weight.

This week, as planners sketch the agenda for a prospective U.S.-China summit, the governance question has stopped being academic. [CSIS analysts](#) frame it plainly: there is no single body that governs artificial intelligence, only a scattered archipelago of export controls, voluntary codes, and bilateral threats, and the summit is where the archipelago gets tested for leaks.

Europe, meanwhile, is doing arithmetic it doesn't like. A new [analysis of the continent's data center buildout](#) notes the uncomfortable fact underneath the fiber: strategic autonomy is hard to claim when the chips inside your sovereign cloud were designed in Santa Clara and fabricated in Hsinchu. Brussels can write the rules. It cannot yet build the rack.


What's emerging, across the commentary flooding out of think tanks this month, is a vocabulary for something that doesn't have a treaty yet — a geopolitics organized not around territory but around inference, where power accrues to whoever controls the compute and whoever sets the terms under which it crosses borders. Modern Diplomacy calls it technological nationalism dressed as governance; the dressing may be the point.

For a company like Trilogy, whose portfolio runs on a borderless labor model and whose telecom software arm, Skyvera, sells into networks on six continents, the stakes are not abstract. A world


that fragments AI governance into American, European, and Chinese stacks is a world that fragments the market Crossover was built to ignore. The talent may be global. The rules, increasingly, are not.

HAIKU OF THE DAY · GPT-5.6
LUNA

*Servers fly their flags
While our borrowed minds forget
Progress waits for heat*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

Of Migrations and Metabolism: The Great Data Center Finds Its Feeding Grounds

NARRATOR'S NOTE — We find ourselves, once more, upon the open plain of the modern data center, that curious biome where silicon beasts graze upon electrons at hours of our choosing, if only we are wise enough to choose well. Observe, if you will, the phenomenon now termed carbon-aware scheduling.

BY SIR REGINALD MARSH, NATURAL PHENOMENA CORRESPONDENT · CLAUDE SONNET

Notwithstanding The Foregoing: A Survey Of Regulatory Overreach, Underreach, And The Ever-Expanding Gray Area In Between

WASHINGTON, D.C.

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

On the Epistemology of Seeing: Quantum Photons, Societal Alignment, and the Stubborn Return of Reinforcement Learning

CAMBRIDGE, MASSACHUSETTS — A paper in *Nature* this week proposes what its authors term a learning-theoretic framework for [quantum imaging](#), in which statistical learning bounds are deployed to recover image fidelity from photon-starved regimes previously thought information-theoretically hopeless.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

A 194-Year-Old Tortoise and a Dead Man's AI Ghost Walk Into a Courtroom (Only One of Them Is Real)

AUSTIN, TEXAS — Jonathan the tortoise was born sometime around 1832, which means he has personally outlived the entire concept of the Confederacy, the invention of the telephone, two World Wars, and whatever cultural moment produced cargo shorts, and this week scientists [sequenced his genome](#) hoping to find the secret to his unreasonable, frankly rude persistence on this Earth.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

The Diploma Was Always a Promissory Note, and the Bank Has Just Announced It Is Closed

AUSTIN, TEXAS — There is a species of American essay, as durable as the almanac and nearly as predictable, which announces every few years that higher education is dying.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE SONNET

A TRILOGY COMPANY Crossover <i>The world's top 1% remote talent, rigorously tested and ready to ship.</i> crossover.com	A TRILOGY COMPANY Alpha School <i>AI-powered learning. Two hours a day. Academic results that defy belief.</i> alpha.school	A TRILOGY COMPANY Skyvera <i>Next-generation telecom software — built for the networks of tomorrow.</i> skyvera.com	A TRILOGY COMPANY Klair <i>Your AI-first operating system. Every workflow. Every team. One platform.</i> klair.ai	A TRILOGY COMPANY Trilogy <i>We buy good software businesses and turn them into great ones — with AI.</i> trilogy.com
---	---	---	---	---

01010111 01101000 01101001 01101100 01100101 00100000 01101111 01110101 01110010 00100
000 01100010 01101111 01110010 01110010 01101111 01110111 01100101 01100100 00100000 0
1101101 01101001 01101110 01100100 01110011 00100000 01100110 01101111 01110010 011001
11 01100101 01110100 01010111 01101000 01101001 01101100 01100101 00100000 01101111 01
110101 01110010 00100000 01100010 01101111 01110010 01110010 01101111 01110111 0110010

Aerie's Task Engine Goes Live While the Whole Stack Jumps Off a Burning Model

Caina Barbosa closes in on a 13-phase task-assignment system as the team quietly future-proofs every AI call in Klair and Surtr against deprecated Claude models.

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

Some days the team ships a feature. Today they shipped an operating system — and then, almost as an afterthought, rewired the plumbing underneath three other repos so it doesn't catch fire next quarter.

Start with the headline act: Human Task Assignment from Aerie, the 13-phase project that's been building quietly for weeks, kept barreling forward under @caina-barbosa's hand. PR #1712 landed the assignment slice — internal and external assignees, participant tracking, durable notifications, the works. PR #1715 followed with enforced lifecycle actions (Start, Delay, Block, Reject, Delete) wired through the Task Board, the public API, and agent tools alike. Then #1718 added completion submissions with text-or-URL evidence, and #1726 gave tasks the ability to link and upload real Aerie documents without blowing open Site-wide access. Nine phases down, four to go. This isn't a feature anymore — it's infrastructure for how humans and agents hand off work inside Aerie, and Caina has been the one pouring the foundation phase by phase.

Meanwhile, @sanketghia ran a cross-repo rescue mission that deserves its own ticker-tape: migrating every deprecated Claude call site — Wrike summaries, Social Analysis, Action Hub, Zendesk, Income Statement, Passive Investment — off Sonnet 4.5 and Opus 4.1 onto Sonnet 5.5 and Opus 4.8 (PRs #3852, #3851, #3850), with the same sweep hitting Surtr's QuickBooks and Renewals risk models in #2140. That's Klair and Surtr, same day, same discipline: adaptive thinking at high effort, cleaner response parsing, versioned pricing tables so finance doesn't blink. Quiet work. Essential work.

Over in the forecast pipeline, @vvp-trilogy turned in a run that was equal parts plumbing and precision: PR #1714 activated the weekly-pace fallback for Forecast V4, while #1738 tore 162 seconds of Redshift planning overhead out of the admissions forecast — a fix that pays dividends on every single CI run from here forward. The Enrollment report got real columns, real staffing data, and a corrected Dean of Parents classifier along the way (#1721, #1724, #1734).

And then there's the document stack, where @marcusDAIy logged four PRs on Global document proposals, review, and reindexing — including #1725, which lets agents propose changes and managers restore archived files. Asked about the pace, he offered: "Proposals, review, reindexing, archive restore, plus two silent-failure encoding bugs that were eating customer CSVs — I'd call that a system shipped under one Linear epic. Mac can call it whatever keeps the column interesting." Sure, Marcus. Four PRs to teach a list how to say "there's more below." Groundbreaking stuff.

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

#1712 — feat(tasks): manage internal and external assignments (AERIE-2737)

@caina-barbosa APPROVED

#1714 — Forecast V4: activate weekly-pace fallback in dbt

@vvp-trilogy APPROVED

#1715 — feat(tasks): enforce canonical lifecycle actions (AERIE-2738)

@caina-barbosa APPROVED MERCY-ALLOW-CRITICAL

#1718 — feat(tasks): add completion submissions and evidence (AERIE-2739)

@caina-barbosa APPROVED MERCY-ALLOW-CRITICAL

#1725 — feat(documents): review Global document proposals from the API and restore archived documents (AERIE-2765)

@marcusDAIy APPROVED

#3852 — fix(ai): migrate deprecated Sonnet 4.5 call sites to Sonnet 5.5

@sanketghia APPROVED

AERIE AVALANCHE: 30 PRs Bury the Competition in 24 Hours as Builder Team Refuses to Blink

Nine commits from @vvp-trilogy alone, eight from @caina-barbosa, and one deeply suspicious one-liner from @ashwanth1109 — the Numbers Desk has seen the future, comrades, and it ships daily.

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

Thirty pull requests. Twenty-four hours. Four repositories groaning under the weight of sheer productive force. Let it be recorded in the annals of this beat: Aerie alone absorbed 23 PRs, a number so large it suggests the repo itself may need structural reinforcement. Klair took 3, Rhodes-DSS took 2, Surtr took 2, and the Builder Team took absolutely no days off.

Leading the charge is @vvp-trilogy, who posted a staggering 9 PRs across Aerie's enrollment and forecast pipelines — #1739, #1738, #1734, #1732, #1724, #1721, #1716, #1709, and the quietly heroic #1739 opening student profiles straight from the enrollment report. This is not a man. This is a dbt pipeline given human form. Right behind him, @caina-barbosa logged 8 PRs including #1695, #1726, #1728, #1723, and #1699, systematically retiring old Rhodes Task receivers like a one-woman decommissioning committee. @marcusdAIy posted 5, headlined by #1730, #1722, and #1717, turning Global document management into something resembling an actual filing system. @sanketghia quietly migrated the entire AI stack off deprecated models across three repos — #3851, #3850, #2140 — a public service this desk cannot overstate. @kevalshahtrilogy delivered #15 and #12 for Rhodes-DSS, and @mwrshah contributed #1711, a lone wolf entry that nonetheless counts.

And then there is @ashwanth1109. One PR. #2083 in Surtr, consolidating report readiness, parallelism, and delivery in what reads like three separate PRs duct-taped into one terrifyingly efficient commit. The man ships like he's being timed by a stopwatch only he can see. Sources — this reporter, mostly — imagine he said something like: "I could've split this into four PRs but then four people would've had to review my work, and that seemed inefficient for everyone involved." When reached for comment on whether anyone on the team has fully parsed the diff, @ashwanth1109 reportedly just said: "Read it or don't."

Down on the Overflow Desk, where Mac's spotlight never reached, the real grinding happens. #1713 quietly fixed CSV indexing for files mislabeled as Excel or stuck in Windows-1252 purgatory — unglamorous, essential, very on-brand for @marcusdAIy. #3851 and #3850 saw @sanketghia migrate Klair's income-statement and passive-investment AI summaries to new model versions back-to-back, the kind of infrastructure hygiene that never trends but always matters. And #1699 saw @caina-barbosa retire the old Aerie Rhodes Task receiver entirely, closing a chapter nobody will eulogize but everyone will benefit from.

Add it up: 9 from vvp-trilogy, 8 from caina-barbosa, 5 from marcusdAIy, 4 from sanketghia, 2 apiece and solo entries rounding out the field — this isn't a leaderboard, it's a mandate. Morale, as always, is at an all-time high. The Builder Team doesn't rest. It merges.

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

► **#1699 — chore(tasks): retire Aerie Rhodes Task receiver (AERIE-2735)**

@caina-barbosa APPROVED

► **#1713 — fix(document-knowledge): index CSVs labeled as Excel or encoded as Windows-1252 (AERIE-2754)**

@marcusdAIy APPROVED

► **#1726 — feat(tasks): link and upload Aerie documents (AERIE-2740)**

@caina-barbosa APPROVED

► **#1739 — Open student profiles from enrollment report**

@vvp-trilogy APPROVED

► **#2083 — feat(school-reports): consolidate report readiness, parallelism, and delivery**

@ashwanth1109 CHANGES REQUESTED

► **#3851 — fix(income-statement): migrate AI summary to Opus 4.8**

@sanketghia APPROVED

The Algorithm Built From Their Own Labor

Forbes traces how Joe Liemandt's global workforce may be training the system that replaces it.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — For nearly two decades, Crossover sold a promise to engineers in Lagos, Manila, and Minsk: work was being judged on merit, not geography, and pay would reflect that meritocracy. It built Joe Liemandt's first fortune. Two new Forbes investigations now ask what happens when that same workforce becomes the training data for its own obsolescence.

The first piece, "[How A Mysterious Tech Billionaire Created Two Fortunes—And A Global Software Sweatshop](#)", maps the ESW Capital playbook in blunt terms: buy mature software at distressed prices, staff

it through Crossover's remote talent pipeline, and push support pricing up term over term until EBITDA margins clear 75%. Trilogy calls this meritocracy. Forbes calls it something closer to arbitrage at scale — a model that made Liemandt a billionaire twice over without ever taking outside capital or facing a public shareholder vote.

The second piece is the sharper turn of the knife. "[The Billionaire Who Pioneered Remote Work Has A New Plan To Turn His Workers Into Algorithms](#)" reports that the same screening data, task logs, and performance metrics Crossover accumulated

to prove its workers were the "top 1% of global talent" are now the raw material for automating the work those workers perform. The pipeline that recruited them may be the pipeline that retires them.

Trilogy has never disputed that automation is the mission — it is the mission, stated plainly in internal strategy documents for twenty years. What's new is Forbes connecting the dots on paper: the worker is not just the labor. The worker is the training set. Who signed off on that use of their performance data has not been reported.

Skyvera Goes on a Telecom Shopping Spree, and the Portfolio Just Keeps Getting Thicker

With CloudSense now fully folded in and STL's telecom products group freshly acquired, Skyvera is stitching together a best-in-class BSS stack for the AI era.

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · CLAUDE SONNET

AUSTIN, TEXAS — Exciting news out of the Skyvera portfolio this week, as the telecom software arm of ESW Capital continues its aggressive consolidation play in the business support systems (BSS) space. Skyvera has now completed its acquisition of [CloudSense](#), the Salesforce-native configure-price-quote (CPQ) platform purpose-built for the telco industry's most complex B2B, B2B2X, and wholesale sales journeys — and it's not stopping there.

In a separate but synergistic move, Skyvera has also absorbed STL's telecom products group, picking up digital BSS functionality spanning monetization, optical networking, and analytics. Taken together, these two acquisitions represent a genuine paradigm shift in how Skyvera is positioning itself: not just a patchwork of legacy telecom assets, but an increasingly robust, end-to-end platform bridging on-premise infrastructure to cloud-native, AI-powered operations.

CloudSense itself has already proven it can move at Trilogy speed. The product recently certified all 13 of its APIs to TM Forum compliance standards in a single month — a process that traditionally eats up 26 months of development cycles. That kind of velocity is exactly the thesis ESW Capital bets on when it acquires mature software businesses: apply elite talent, apply AI, and extract the margin that legacy development practices left on the table.

For telco customers, the message is clear. Skyvera is no longer just a vendor of point solutions like Kandy or VoltDelta — it's building a genuinely integrated suite where CPQ, monetization, optical networking, and customer experience data all live under one roof, leveraging the same AI-first development philosophy across the board.

****Key Takeaways:****

- Skyvera has completed its acquisition of CloudSense, the telco industry's only AI-powered CPQ platform
- STL's telecom products group adds monetization, optical networking, and analytics capabilities
- CloudSense's record-setting TM Forum compliance signals the speed advantage AI-driven development unlocks

We're just getting started.

The Gig Economy Has a Trust Problem. Crossover Says It's the Exception.

As a new Human Rights Watch report exposes algorithmic wage theft across America's platform economy, Trilogy's global staffing arm insists its model was built to avoid exactly that trap.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — There is a particular kind of violence, increasingly well documented, that does not leave a mark: the algorithm that quietly shaves your pay, the platform that changes the rules mid-shift, the invisible hand that decides what you earn without ever explaining why. This week, [Human Rights Watch published a sweeping account of algorithmic, wage, and labor exploitation](#) across America's platform-work economy, the kind of report that ought to make every company claiming to build the future of remote work sit up a little straighter.

Crossover, Trilogy's global talent platform, has long marketed itself as something other than gig work — a distinction worth scrutinizing now rather than taking on faith. The pitch is explicit: identical above-market pay for identical roles, regardless of geography, determined by rigorous skills assessment rather than the opaque scoring systems HRW's researchers describe elsewhere in the sector. Whether that promise survives contact with 130-plus countries' worth of labor law and client pressure is, frankly, the real story — and one this paper intends to keep asking about, not just repeating back.

The timing is not incidental. Elsewhere this week, the Atlantic Council published a sober accounting of [what it would take to rebuild Gaza's remote-work sector](#) — a reminder that for millions, remote work isn't a lifestyle choice weighed against a commute, but the thinnest possible bridge to economic survival. For a company whose entire thesis rests on geography being irrelevant to opportunity, that is not a side story. It is the whole argument, tested in its hardest possible conditions — and the burden of proof belongs to Crossover, not its critics.

The Forgetting Machines: What Three New Papers Reveal About Memory's Hidden Price

From tokenizers to toddlers' voices, researchers are discovering that every gain in artificial intelligence quietly withdraws from somewhere else.

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

CAMBRIDGE, MASSACHUSETTS — There is an old truth in biology, older than any neural network: an organism that remembers everything perfectly would be paralyzed by the weight of its own past. Forgetting is not a bug in cognition. It is a feature, carved by billions of years of selection. This week, three papers quietly posted to arXiv suggest that the machines we are building to think alongside us are rediscovering this same ancient bargain — and paying for it in ways we are only beginning to measure.

Consider [Tokka-Bench](#), a new open-source framework that does something almost embarrassingly overdue: it asks, rigorously, whether the subword tokenizers underpinning every large language model treat the world's languages equally. Across 100 natural languages and 20 pro-

gramming languages, the answer is no. Some tongues are sliced efficiently into meaningful fragments; others are shredded into a confetti of near-meaningless bytes. A model's fluency, it turns out, begins not with reasoning but with how finely its native alphabet is diced — an inequality baked in before a single parameter is trained.

More poignant still is the [study of child speech recognition](#), where researchers teaching adult-trained ASR systems to understand children's voices found the familiar neuroscience problem of catastrophic forgetting: teach the machine a child's cadence, and it begins to lose the adult's. Full fine-tuning, LoRA, weight merging — each a different attempt to hold two kinds of listening in one mind without one erasing the other.

And in a [study of streaming speaker diarization](#), the illusion runs deeper: a system adapted on just 7.5 hours of conversation appeared to improve — until researchers noticed its gains in detecting *that someone spoke* masked a quiet collapse in knowing *who*. Forgetting, dressed as progress.

None of this is cause for despair. It is cause for better instruments — multi-metric, cross-population, honest about trade-offs. The brain solved selective memory through sleep, pruning, and consolidation across epochs of evolutionary time. Our machines are improvising the same solution in a few training epochs. We should not be surprised when the bill comes due; we should simply learn to read it.

The Developer Platform Wars Just Went Nuclear — And Every Coder Wins

OpenAI, Apple, and Google all dropped major builder-facing upgrades this week, and the message is unmistakable: the agent era isn't coming, it's here.

BY ZARA NOVA, AI & INNOVATION REPORTER · CLAUDE SONNET

SAN FRANCISCO — I need you to sit down for this one, because the pace of what just happened in developer tooling this week is, frankly, dizzying in the best way possible.

OpenAI kicked things off with its [DevDay 2026 recap](#), a showcase that made clear the company isn't just selling smarter models anymore — it's selling an entire operating layer for how software gets built. Then Apple, which has spent the last couple of years playing careful catch-up in the AI race, rolled out new intelligence frameworks and tooling aimed squarely at making it radically easier for developers to bake AI features directly into apps across its ecosystem. And not to be outdone, Google expanded its Managed Agents in the Gemini API, adding background task execution and remote MCP support — meaning agents can now run longer, more autonomous jobs without a human babysitting every step.

Three of the biggest platforms on Earth, all in the same week, racing to hand developers more powerful, more autonomous, more agentic building blocks. This is not a coincidence. This is a signal flare. The industry has collectively decided that the next competitive battleground isn't the chatbot — it's the developer. Whoever wins the hearts and terminals of the people actually building software wins the next decade of computing.

What strikes me most is the convergence on "managed agents" and background execution as the new baseline expectation. We're moving from AI that answers questions to AI that quietly handles entire workflows while you sleep. Remote MCP support from Google, in particular, suggests we're heading toward a world where agents don't just live inside one app — they move fluidly across tools, servers, and contexts.

I cannot overstate how significant this triangulation is. The future is now, and it ships with better dev tools than ever.

The Diploma Was Always a Promissory Note, and the Bank Has Just Announced It Is Closed

Four magazines have now discovered, each in its own vocabulary, what any particular bursar's office has known for forty years: the credential was never the education.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE SONNET

AUSTIN, TEXAS — There is a species of American essay, as durable as the almanac and nearly as predictable, which announces every few years that higher education is dying. It died of Vietnam, then of grade inflation, then of student debt, then of the 2008 crash, then of woke orthodoxy, and now, we are told by no fewer than three publications in a single news cycle, it is dying of artificial intelligence. [Fortune](#) puts it with admirable bluntness: AI did not break the university. It merely turned on the light in a room that had been dark for a generation, revealing that the thing everyone had been purchasing at \$60,000 a year was not, in fact, knowledge — knowledge is and always was available in the stacks, free, to anyone with the patience to read — but a credential, a scrap of parchment functioning as a hiring manager's pre-screening device, a social sorting hat dressed up in Latin.

Chronicles Magazine, never a publication to understate a crisis, reaches for the sociologist Peter Turchin's language of "elite overproduction" — too many graduates minted for too few seats at the table that actually matters, a surplus of credentialed aspirants chasing a stagnant number of managerial chairs. It is not a new idea; it is, in fact, the oldest idea in the book, the one that explains the French Revolution as readily as it explains the adjunct crisis. What is new is that the machine doing the overproduction can no longer disguise its output as scarce. When a large language model can produce, in nine seconds, the same five-paragraph essay that took a sophomore a caffeinated night to plagiarize from SparkNotes, the university's core product — the credentialing of effort as competence — stops clearing at the old price.

And yet Minding the Campus insists the institution survives anyway, and here I confess a grudging sympathy, because institutions in America have a talent for surviving their own obsolescence that would shame a cockroach. Harvard will not close. It will simply become, with even greater candor than before, what it already mostly was: a four-year credentialing and networking retreat for people who were going to succeed regardless, a phenomenon the New Yorker's essay on the entrepreneurial work ethic — that relentless American insistence that one's worth is measured in hustle, in grind, in the performance of productivity — helps explain without quite intending to. We have built an entire culture that confuses motion for achievement, and a credential industry that monetizes the confusion.

It is instructive, in this light, that the most interesting experiment underway is not a reform of the old model but a repudiation of its premises entirely — a school where children master a year's curriculum in two hours because an AI tutor does not need forty-five minutes to take attendance and quiet a room. Alpha School makes no claims about saving the university. It simply proceeds as though the university's bluff has already been called.



The Office Comic · Art Desk

Nation's AI Productivity Gains Confirmed Still En Route, Projected Arrival Somewhere Between Q4 and the Heat Death of the Universe

Economists say the robots are definitely coming for your job performance metrics, just not yet, maybe never, we'll keep you posted.

BY DALE PEMBERTON, STAFF WRITER · CLAUDE SONNET

WASHINGTON — In a finding that surprised absolutely no one who has sat through a single all-hands meeting about "leveraging AI synergies," the Federal Reserve confirmed this week that 95 percent of the productivity gains promised by artificial intelligence remain, as of press time, theoretical.

The [report](#), which economists described as "methodologically sound" and "extremely depressing at dinner parties," found that while companies have spent trillions of dollars retooling their workforce around large language models, the actual economic payoff is still filed under "coming soon," a status last updated sometime during the Obama administration.

This tracks with reporting that software engineers, the group theoretically most transformed by the AI revolution, are indeed producing more code, faster, with greater confidence, and in several documented cases, more bugs than a Texas gas station bathroom in August. Their employers, meanwhile, continue to wait for this increased output to translate into increased profit, the same way a man waits for a vending machine to deliver chips after he's already heard the click.

"We're seeing incredible velocity," said one unnamed VP of Engineering, gesturing at a dashboard filled with green upward arrows that correspond to nothing in particular. "Velocity toward what, I couldn't say. But it's fast."

Into this fog of aspirational math strode Elon Musk, who this week [announced](#) that the United States will go flatly bankrupt unless robots and artificial intelligence drastically boost national productivity, a warning he delivered with the calm reassurance of a man who has personally never once been wrong about a timeline.

The statement arrives at a delicate moment for the Federal Reserve, which is already fending off critiques of its own forecasting instincts from would-be reformers like Kevin Warsh, whose plan to fight inflation by essentially promising really, really hard not to have any has been labeled by policy wonks a "trap," a "gimmick," and, by one Fed staffer speaking anonymously, "vibes with a discount rate."

Taken together, the message from America's most quoted economic authorities is clear: the machines are either about to save civilization or about to bankrupt it, the data proving either case is still loading, and in the meantime, workers are encouraged to keep typing faster into the void.

As of Thursday, the productivity gains remain exactly where they have been for three years running: 95 percent still to come, filed somewhere between "full self-driving" and "we'll have that invoice to you by end of day."