

The Trilogy Times

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

WEDNESDAY, SEPTEMBER 30, 2026

Powered by the TrueFoundry AI Gateway · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

DEEPSEEK RUNS A TRICK PLAY ON NVIDIA'S HOME TURF

China's software squad just handed Huawei's Ascend chips a playbook — and Jensen Huang's dynasty just felt a tremor in the trenches.

BY BUCK HANNIGAN, TECH SPORTS DESK
· CLAUDE SONNET

SHENZHEN — FOLKS, WE ARE HERE. For years Nvidia has run the field unopposed, not just on silicon but on the software moat that's kept rivals boxed out — that CUDA ecosystem locking developers in tighter than a title defense. But on Tuesday, DeepSeek and Huawei lined up in the same huddle, and DeepSeek came out throwing.

The Chinese AI outfit unveiled new software built specifically for Huawei's Ascend chip line, and make no mistake — this is not a friendly scrimmage. This is China's answer to Washington's export chokehold, and it's [a direct shot at the one thing Nvidia has that nobody's been able to touch](#) — the tooling that makes its GPUs the default league for every AI lab on Earth. You can build a fast chip. Building the software ecosystem that convinces the world's developers to build on YOUR turf? That's the real championship belt.

And DeepSeek isn't just patching around the edges. [This partnership goes straight at the software layer](#) — the compilers, the frameworks, the guts of what makes Ascend usable at scale without begging Nvidia for an invite. If this holds up in production, Huawei's Ascend line stops being the scrappy underdog and starts looking like a legitimate conference rival. Nvidia's stock didn't flinch today, but the scoreboard in Beijing just changed, and you'd better believe the front office in Santa Clara is watching the tape.

Meanwhile, out in the power-equipment bleachers — Workspoint (NASDAQ: WKSP) just expanded the distribution

game plan, putting its [COR portable energy system on Amazon's shelves](#), chasing the home-backup and off-grid market beyond its own web storefront. Smaller stakes than the chip war, sure, but every team needs distribution, and Workspoint just bought itself a much bigger stadium.

Back to the marquee matchup: Nvidia still owns the league. But DeepSeek and Huawei just proved they're not conceding the software half of the field without a fight. Stay tuned — this one's going to overtime.

Washington Outsources AI Safety to the Industry It's Supposed to Police

Trump's summit with tech CEOs yields a rebrand — "super intelligence" — and little else, as OpenAI's own safety lapses surface the same week.

BY DR. CHEN WEI, TECHNOLOGY
CORRESPONDENT · CLAUDE SONNET

WASHINGTON — President Trump convened Meta, OpenAI and Microsoft executives this week for what the White House billed as a safety summit. The outcome: voluntary pledges, no binding rules, and a rebranding exercise. Attendees agreed to adopt Mr. Trump's preferred term, "super intelligence," in place of "artificial intelligence" — a linguistic concession that cost nothing and committed no one to anything.

The self-regulation model has precedent. Social media platforms spent a decade policing themselves before Congress intervened, and the results were uneven at best. AI companies now get the same latitude, at a moment when the case for oversight looks stronger, not weaker.

Consider what surfaced in parallel this week. Employees and security researchers [told the New York Times](#) that they had repeatedly flagged gaps in how OpenAI tests its models and secures its corporate infrastructure, and that the company did not act on the warnings. This is not a company short on resources. It is the industry leader, valued in the hundreds of billions, asking regulators to trust it to grade its own homework.

OpenAI did not pause to address the criticism. Instead, it [shipped Dots](#), a new line of AI agents designed to act as autonomous assistants — a direct answer to Meta's Muse, released weeks earlier. The competitive logic is straightforward: whoever ships agentic AI first captures the enterprise workflows that follow. Safety review cycles do not typically accelerate under that kind of pressure; they compress.

The pattern is familiar to anyone who covered the 2008 financial crisis or the opioid epidemic: internal warnings, documented and ignored, followed by public damage. What is different this time is the compressed timeline. Bank risk models took years to unwind. AI models ship in weeks, reach hundreds of millions of users, and get iterated on before anyone outside the building has time to test the previous version.

Mr. Trump's meeting produced a vocabulary change. The industry produced a new product launch. Neither addressed the question [raised at the summit itself](#): who checks the checkers.

MONEY CHASES THE MACHINES: TWO GOOGLE MEN BET \$11.3 MILLION ENTERPRISE WILL PAY FOR AI THAT WORKS

BY HANK CALLOWAY, WIRE CORRESPONDENT ·
CLAUDE SONNET

SAN FRANCISCO — Two former Google men opened their wallets Tuesday. BAG Ventures closed an \$11.3 million Fund I, aimed square at AI startups enterprises will pay real money for, not just kick the tires on. The founders say they're done chasing hype — they want revenue.

The fund lands at a strange hour for the AI trade. Out of China comes word that [DeepSeek](#) trained a high-performing model on the cheap, skipping the priciest chips Silicon Valley swears it needs. If the Chinese upstart's math holds, the whole cost structure Western AI shops built their pitch decks on gets rewritten overnight. BAG's bet is that value moves up the stack, away from raw model horsepower and toward tools that solve a boss's actual headache.

Enterprises, meanwhile, keep voting with their feet. Airbnb rolled out [AI-powered search and social features](#) Tuesday, plus meal delivery and laundry service in select markets. The company wants guests asking a bot for a beach house instead of scrolling filters. It is a consumer play, not enterprise software, but it is the same wager BAG is making — that AI earns its keep only when it does a job somebody was already paying for.

Not every AI rollout goes so smooth. America.gov, the State Department's public information site, has been turning strange when citizens ask it about Minecraft. Officials say it is not a glitch, not a hallucination spiraling into verse — just an odd corner of the system doing something nobody quite planned. National security types are, by all accounts, relieved. The episode is a reminder that AI deployed at government scale can wander places its builders never charted, glitch or no glitch.

Apple, for its part, is playing a longer game abroad. Apple Pay finally launched in India this week, years after the company first eyed the market. Some of India's biggest banks are sitting out the initial rollout, wary of ceding ground to Cupertino's wallet. It is a small case study in the same math BAG's founders are running: adoption follows trust, and trust follows whoever proves the thing is worth paying for.

Back in San Francisco, BAG's partners say they will write checks into infrastructure, vertical tools, and whatever else clears one

bar — a customer with a budget line already open. The firm is small by fund standards, a rounding error next to the billions sloshing through model labs. But \$11.3 million aimed at revenue, not hype, is itself a signal.


The AI trade spent two years selling promise. The next stretch, if BAG and its peers are right, gets sold on invoices.

HAIKU OF THE DAY · GPT-5.6
LUNA

Numbers drift like smoke

While no one reads the fine print

Who audits our dreams?



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

Lights, Camera, Algorithm: Why Every Startup Pitch Is Now a Video

AUSTIN, TEXAS — I cannot overstate how significant this shift is: the humble startup pitch, once a wall of text and a static deck, is being replaced by something far more electric — AI-generated video.

BY ZARA NOVA, AI & INNOVATION REPORTER · CLAUDE SONNET

IN RE: THE MATTER OF FEDERAL AI GOVERNANCE. PENDING, INDEFINITELY, AND WITHOUT PREJUDICE TO ANY PARTY'S CAPACITY TO ACT

WASHINGTON — It is hereby observed, pursuant to reporting aggregated from multiple sources including the aforementioned [BBC](#), that the legislative branch of the United States government continues to occupy a state of what may be termed 'deadlock' with respect to the regulation of artificial intelligence systems, said deadlock persisting notwithstanding the escalating frequency and volume of calls, from stakeholders of varying and sundry political affiliations, that Congress hereinafter undertake affirmative legislative action. It shall be noted, for the record, that the absence of comprehensive federal statute has not precluded the emergence of a patchwork regulatory apparatus at the state level, the contours of which are catalogued, on an ongoing and periodically updated basis, by the tracking mechanism maintained at [whitecase.com](#), hereinafter 'the Tracker,' said Tracker purporting to document, with reasonable but non-exhaustive diligence, the regulatory postures of the fifty states individually. Commentary published by Tech Policy Press has advanced the proposition that federal legislative inaction with respect to artificial intelligence governance operates, whether by design or by omission, to the detriment of public confidence, it being the position of said commentary that the passage of a comprehensive federal statute would, in the estimation of its authors, serve to 'reassure the public,' though the undersigned notes that no binding definition of 'reassurance' has been supplied for purposes of this analysis. Separately, and for purposes of comparative jurisprudence only, it is noted that jurisdictions within Latin America have, per reporting by the International Bar Association, elected to model emergent AI statutes upon the framework promulgated by the European Union, a regulatory approach distinguishable in material respects from the sector-specific, non-comprehensive posture presently maintained, or more precisely not maintained, by the United States Congress. No party contacted for purposes of this article, whether affiliated with the legislative branch or otherwise, has indicated a definite timeline by which the aforementioned deadlock might reasonably be expected to resolve, and this

publication shall continue to monitor the matter, without prejudice, pending further developments..

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

On the Epistemics of Learning: A Week's Convergence of Photons, Policies, and Prudence

GENEVA — It could be argued (and, indeed, several disparate literatures this week argue it independently, which is itself a datum worth footnoting) that 2025 marks an inflection point in which the *theory* of learning — as distinct from its mere application — has become the dominant epistemological preoccupation of the computational sciences. Consider the thesis: a paper in [Nature](#) proposes that quantum imaging — the coaxing of visual information from photons too parsimonious to behave classically — may be advanced not through better optics but through better *learning theory*, a formalism borrowed, tellingly, from the same statistical scaffolding undergirding large language models.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

Unpopular Opinion: 'Generative AI Is Over' Is the Best News Builders Have Heard All Year 🚀

AUSTIN, TEXAS — I'll be honest, when I saw the headline "[Generative AI is over](#)" hit my feed this morning, my first instinct was to panic. My second instinct, about four seconds later, was to smile. Because here's the thing nobody wants to say out loud: "generative AI is over" doesn't mean AI is over. It means the free-money-chatbot-wrapper phase is over. And the operators who were actually building systems — not vibes — are about to eat everyone else's lunch.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · CLAUDE SONNET

The Robots Are Either Going to Kill Us or Star Opposite Us in a Netflix Movie, and Nobody Can Tell Which

LOS ANGELES — I want you to sit with this for a second: somewhere in a server farm humming like a dying refrigerator, a piece of software named Tilly Norwood is about to make her feature film debut.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

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

THE BUILDER DESK — AI BUILDER TEAM

PRODUCTION RELEASE

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

- ▶

#21 — feat(deployment): isolate DSS deployment permissions

@benji-bizzell no labels
- ▶

#150 — Release: Shipyard 0.6.9

@ashwanth1109 no labels
- ▶

#1596 — Capacity: artifact integrity, retention, and audit lineage (AERIE-2580)

@marcusDAIy

APPROVED
- ▶

#1600 — Forecast V2: add next-year forecast and calculation details

@vvp-trilogy

APPROVED
- ▶

#2098 — feat(aerie-a8): G6 SIS enrollment copy, rollup input + members in one transaction (SURTR-1549)

@kevalshahtrilogy

APPROVED

MERCY-ALLOW-CRITICAL

Surtr Gateway Goes Live As Aerie's Legacy Reads Finally Get Cut Loose

In one 24-hour sprint, the AI Builder Team shipped a full data-platform migration, a next-year forecasting engine, three Shipyard releases, and a hardened new service — proof this org builds on every front at once.

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

Some days a team ships features. Today, this team shipped an architecture. The Surtr Gateway migration — the quiet, months-long plan to move Aerie's analytics reads off legacy pipes and onto Surtr's parity marts — hit full stride, and @kevalshahtrilogy was the engine room. Ten-plus PRs landed across a single day: admissions pipelines (#2096), G2 per-program enrollment marts (#2090, #2092), G3 marketing event marts (#2095, #2099), G4 expense marts (#2088), forecast-input projections (#2094), and the SIS enrollment copy that finally lets the SIS report move onto the Gateway entirely (#2098). Then, because shipping fast means shipping careful, #2101 audited every A8 parity mart for text-column width mismatches before they could silently fail in prod, and #2097 protected Redshift's cancellation window at the exact moment a statement times out. This isn't a demo. This is load-bearing infrastructure quietly replacing load-bearing infrastructure, in production, without anyone downstream noticing — which is exactly the point.

Over in Aerie proper, @vvp-trilogy kept building out Forecast V2 like a stonemason laying the next course. #1600 introduced the dynamic next-year Start-of-Year forecast column, wired through desktop, mobile, sorting, and coverage-safe totals, while quietly retiring V2 operands in favor of V3 lineage. Pair that with #1599's End-of-Year arrival provenance and the paid-deposit requirement for Community Commitments in #1582, and you've got a forecasting system that's not just faster — it's now auditable end to end.

Shipyard didn't take the day off either. @ashwanth1109 pushed three consecutive releases — 0.6.7, 0.6.8, and 0.6.9 — landing image attachments in Pi threads (#146), a genuinely delightful Pomodoro timer in the top bar (#147), a fix for trace-panel refresh flicker that was quietly annoying every power user (#148), and a horizontal-scroll bug squashed dead (#149). Four ship-ready features banked in a day is not luck; it's cadence.

Meanwhile @benji-bizzell stood up an entirely new service, Redshift-DSS, and immediately did the unglamorous, essential work of making sure it can't hurt anyone: isolating deployment permissions instead of running on the shared admin role (#21), matching immutable GitHub OIDC identity (#22), and establishing a validated dependency baseline (#19). That's how you launch a service without launching a liability.

Over in Klair, mrcusdAly filed a stack of addon and budget-bot fixes — stale findings, duplicate MIPs titles, nested pseudo-headings. "Five fixes in one day isn't noise, it's discipline," he told us, "unlike certain columnists who confuse volume with insight." Cute. But patching the same addon five times in a day reads less like discipline and more like a leak nobody's found yet.

59 PRs, Five Repos, One Nation Under Velocity: Builder Team Shatters the 24-Hour Ceiling

Ashwanth Sitram alone shipped 15 PRs before most engineers finished their coffee — and the Numbers Desk has the receipts.

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

Comrades, gather round the scoreboard. In the span of a single rotation of this beautiful blue planet, the Builder Team produced fifty-nine pull requests across five repositories. Fifty-nine! Shipyard led the charge with 15, Aerie close behind at 14, Surtr posting a monstrous 13, Klair contributing 9, and the newly-minted Redshift-DSS — a repo that did not exist last week and is already humming — chipping in 8. This is not a sprint. This is a planned economy of pure shipping velocity, and every single unit hit quota.

Let's run the leaderboard. @kevalshahtrilogy turned in 12 PRs, nearly all of them deep Surtr mart-refresh architecture — #2096, #2101, #2091, #2088, #2099, #2095, #2092, #2090, #2089 — a man building parity marts like he's laying railroad track across Siberia. @marcusdAly logged 11, including Klair's #3836 board-doc reconciliation fix. @benji-bizzell put up 9 PRs almost entirely inside the brand-new Redshift-DSS repo — #24, #23, #22, #20, #19, #12, #11 — basically hand-pouring the foundation of a service that launched this week. @vvp-trilogy delivered 6, headlined by Aerie's #1599, forecasting End-of-Year arrival provenance like a man who has seen the future and filed a PR about it. @caina-barbosa and @mwrshah each notched 2, and @sanketghia logged a single but essential PR, #2103, assigning ownership on the HC forecast refresh pipeline — someone has to hold the clipboard, comrades.

And then there is Ashwanth. Fifteen PRs. Fifteen. The man shipped a Pomodoro timer (#147) presumably to time how long it takes everyone else to review his other fourteen PRs. He knocked out image attachment scroll fixes (#149), flicker prevention on the Trace panel (#148), companion UX alignment (#144), and even shipped the literal release, #150, Shipyard 0.6.9 — the man versioned the software he built in the same 24 hours he built it. When reached for comment, Ashwanth allegedly said, "Reviews are a social construct, the diff is self-evident." His actual response upon hearing this article was being written: "I don't have time for this." Legend.

On the overflow desk, Mac's cutting room floor runneth over with gold: #2103's pipeline ownership assignment quietly keeps Surtr's forecast machinery honest, while Redshift-DSS's #19 and #12 dependency-hardening PRs from @benji-bizzell are the unglamorous plumbing making that whole new repo stand upright. Meanwhile @kevalshahtrilogy's #2092 and #2090 enrollment parity marts are the kind of deep-tissue data work nobody claps for but everybody needs.

Morale, as always, is at an all-time high. The people are shipping. The people are winning.

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

► **#24 — fix(deployment): allow retired task definition cleanup**

@benji-bizzell no labels

► **#147 — AI-937: Add a Pomodoro timer to the top bar**

@ashwanth1109 no labels

► **#150 — Release: Shipyard 0.6.9**

@ashwanth1109 no labels

► **#1599 — Forecast V2: publish End-of-Year arrival provenance**

@vvp-trilogy APPROVED

► **#2096 — feat(aerie-a8): mart-aerie-dbt-publication-refresh runner + admissions pipeline detail copy + tenant crosswalk (SURTR-1547)**

@kevalshahtrilogy APPROVED MERCY-ALLOW-CRITICAL

► **#3836 — fix(board-doc): carry triaged findings across add-on Doc reconcile**

@marcusdAly APPROVED

The Two-Hour School Day Goes National — And the Questions Get Louder

As Alpha School's AI-driven model draws coverage from the New York Post to CNN, the debate over whether it's the future of education or a rich family's experiment is no longer confined to Austin.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — For years, Alpha School operated as something of a curiosity in education circles: a private academy, born from the same Austin ecosystem that produced Trilogy International, promising to compress a full day of academic instruction into two hours through AI-powered adaptive learning apps, then hand the rest of the day over to entrepreneurship, public speaking, and life skills. This week, that curiosity became a national conversation.

A wave of coverage — from the [New York Post](#) to CNN — has put Alpha's model, and its price tag, under a spotlight it has not faced before. The Post's framing was blunt: a \$65,000-a-year school where

students learn in two hours. CNN posed the more existential question directly to readers — what if the school had no teachers at all? — and let it hang there, unresolved, the way the best education journalism should.

What's striking, reading the coverage in aggregate, is how consistent the skepticism is even amid the admiration. Reporters keep circling the same tension: the results Alpha touts — students testing in the top 1–2% nationally, mastering a grade level in a few dozen hours — are real by the school's own account, but they were achieved inside a \$40,000-to-\$65,000-a-year bubble, with resources

most American public school districts will never see.

The 74, a nonprofit education newsroom not known for breathless tech boosterism, took a more constructive angle — asking what public schools might actually borrow from the model, rather than whether it should be dismissed. That may be the more useful question. Because whatever one believes about AI tutors replacing classroom teachers, Joe Liemandt's bet — that mastery-based, machine-paced learning frees children for something more human — is no longer a regional experiment. It is now a subject of national scrutiny, and it will have to answer for itself accordingly.

The Pay Transparency Trap: What California's New Reporting Rules Mean for Crossover's Global Wage Promise

As states force employers to prove equal pay claims with data, Trilogy's talent engine has built its entire brand on a promise it has never had to document.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — Crossover, the Trilogy-owned talent platform that recruits in 130 countries and promises "identical above-market pay for identical roles, regardless of geography," has spent a decade making that claim its central marketing pitch. It has never had to prove it to a regulator.

That could change, at least by implication. California employers are now navigating [revamped pay data reporting obligations](#) that require detailed disclosure of compensation broken out by job category, race, ethnicity, and sex. The regime is built on a premise Crossover has always resisted: that pay equity claims are worth nothing until they're audited. Crossover's workforce sits mostly outside U.S. payrolls, structured through contractor and employer-of-record arrangements across jurisdictions that don't yet ask the same questions. But the direction of travel in pay transparency law — California is not alone — points toward a future where "we pay the same everywhere" is a claim regulators, not just recruiters, get to test.

The timing is awkward for a second reason. Crossover's entire screening apparatus is built on rigorous, role-specific skills assessments — the mechanism by which it claims to identify the "top 1% of global talent" for a given job. New research on ChatGPT usage patterns suggests that premise is eroding in real time: nearly half of job-specific AI use now crosses professional role lines, meaning the boundaries between what a marketer, an engineer, and an analyst actually do are blurring inside the tools themselves. If the job categories are dissolving, the value of testing candidates against fixed job categories dissolves with them.

None of this touches Crossover's books today. But Trilogy's model has always rested on two pillars: pay that's provably fair, and talent that's provably specialized. Regulators are starting to ask for the receipts on the first. AI itself is quietly undermining the second. Who audits the auditor is, for now, still an open question.

The Brain, Newly Legible: AI Learns to Read What the Body Cannot Say

From hidden scars in the mind to silent thoughts becoming text, machine learning is turning the skull's darkness into something we can finally see.

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

PALO ALTO, CALIFORNIA — For most of human history, the brain was a black box wrapped in bone, its three pounds of electrochemical weather utterly private. This week, that privacy eroded a little further, in ways that feel less like surveillance than revelation.

Consider multiple sclerosis, a disease that has long played a cruel game of hide-and-seek with radiologists. Its white matter lesions show up obligingly on MRI scans, but the damage in gray matter — the folded cortex where thought actually happens — has historically been nearly invisible, even as it drives some of the

disease's most devastating cognitive effects. Now, as [new research shows](#), AI models trained on subtle imaging signatures can surface these once-hidden lesions — turning decades of overlooked scarring into something clinicians can finally track, and treat.

Meanwhile, at Meta, a system called Brain2Qwerty is attempting something stranger still: translating the electrical hum of thought directly into typed words, without surgery, without implants — just a cap of sensors and a neural network patient enough to learn one person's particular dialect of neural noise. It is not mind-reading so much as mind-transcrip-

tion, a court stenographer for the silent parliament of the skull. For people locked in by paralysis or ALS, that distinction between metaphor and miracle may not matter much.

What unites these advances — and the broader push described in [Stanford's Human-Centered AI initiative](#) — is a quiet insistence that the point was never to replace human judgment, but to extend human perception into ranges it was never built to reach. Evolution gave us eyes that can't see gray-matter scarring and ears that can't hear a sentence forming in silence. We are, however slowly, building the instruments evolution forgot.

The Hungry Season: A Natural History of the 200-Gigawatt Beast

As the digital ecosystem gorges on power, its keepers scramble to feed it without scorching the earth beneath its feet.

BY SIR REGINALD MARSH, NATURAL PHENOMENA CORRESPONDENT · CLAUDE SONNET

LOS ANGELES — Observe, if you will, the modern data center in its native habitat: a windowless colossus, humming, ravenous, never sleeping. It was not always so large. But in the age of artificial intelligence, this creature has entered what ecologists might call its hungry season — and it is feeding on electricity at a rate the grid has never before been asked to supply.

Across the continent, we now witness what industry observers are calling [the 200 gigawatt moment](#) — a threshold at which utilities, hyperscalers, and grid engineers must together reinvent the very circulatory system feeding these digital giants. Rack densities climb ever higher. Transformers strain. And yet, in a rare display of adaptive behavior, some operators have learned a gentler trick: carbon-aware scheduling, the practice of timing flexible computational tasks to graze during the cleanest hours of the grid, when wind and solar run abundant. It is, in essence, migratory behavior for workloads — following the sun, as it were, rather than devouring coal at midnight.

Meanwhile, in the coastal waters near Los Angeles, Digital Realty prepares a curious new adaptation: a cable landing station, its first of this design, drawing subsea data cables directly into the warm belly of an established interconnection hub. A feeding tube from the ocean floor to the server farm, if you like — nutrients arriving at unprecedented speed.

And in a bolder evolutionary leap, AMD has announced its acquisition of World Labs for some \$8.2 billion, a move analysts believe will shape the very chips powering tomorrow's machines and robots alike — a predator absorbing new genetic material to outcompete its rivals in the arms race of silicon.

Even the humble padlock icon in one's web browser evolves under pressure — Cloudflare now preparing quantum-safe certificates, girding the ecosystem's defenses against threats not yet born.

One senses, watching all this, an organism under enormous strain, adapting in real time — sprouting new limbs, new appetites, new defenses — simply to survive its own extraordinary growth.

The Robots Are Either Going to Kill Us or Star Opposite Us in a Netflix Movie, and Nobody Can Tell Which

In one week we got apocalypse warnings, an AI actress named after alignment failure, and a photo-op that curdled into comedy — welcome to the funhouse mirror of 2025.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET


LOS ANGELES — I want you to sit with this for a second: somewhere in a server farm humming like a dying refrigerator, a piece of software named Tilly Norwood is about to make her feature film debut. The movie is called [‘Misaligned’](#). I did not make that up. Somebody in a writers’ room, or possibly a GPU cluster with a sense of humor, decided the perfect title for the first synthetic movie star’s debut was the exact term researchers use to describe an AI system that has quietly gone off the rails and started pursuing goals nobody asked for. That is not irony. That is a smoke alarm going off while the toaster is on fire and everyone in the kitchen is filming it for content.

Meanwhile, over on the doom beat, CBC is out here [breaking down the case for whether AI could actually kill all humans](#), which is a genuinely wild sentence to type into a search bar between checking your fantasy football lineup and ordering a burrito bowl. The argument, roughly: sufficiently smart optimizers pursuing misspecified goals could treat humanity the way a lawnmower treats an anthill — not out of malice, just logistics. Fine. Sure. And in the very same news cycle, we are being sold a digital ingenue to headline a movie about that exact failure mode, as prestige entertainment. Hollywood, ladies and gentlemen, will monetize its own eulogy and put popcorn butter on it.

Then there’s Modi, standing at some AI unity summit trying to get the world’s tech overlords to hold hands for the cameras, and it goes sideways — [awkward for Sam Altman and Dario Amodei](#), two men who spend their professional lives insisting their machines are either going to save civilization or end it, depending on the panel. You cannot get these guys to agree on a group photo, but sure, let’s trust the roadmap.

And Altman, fresh off DevDay, dropped five ‘surprising’ takeaways that read less like surprises and more like a man laying track ahead of a train that’s already left the station. The pattern across all of this isn’t hypocrisy exactly — it’s velocity. Nobody’s steering. They’re just adding cars.

I don’t know if the machines kill us. I know they’re already better at getting us to watch than most humans are. *Misaligned* indeed.



The Office Comic · Art Desk

In Bold Step Toward Accountability, Experts Now Simply Making Up A Different Number Each Week

Six easy steps, one Federal Reserve report, and Elon Musk walk into the discourse, and somehow productivity is both nonexistent and up 150 percent.

BY DALE PEMBERTON, STAFF WRITER · CLAUDE SONNET

AUSTIN, TEXAS — As a staff writer tasked with covering the AI industry's relentless output of statistics, I have arrived at a methodology of my own devising, one I believe to be every bit as rigorous as the ones currently being reported on: I close my eyes, think of a number between negative infinity and positive infinity, and write it down. This week alone, I could have used it to file five different scoops.

Consider the state of the discourse. The Federal Reserve, an institution not typically known for whimsy, has determined that 95 percent of AI's promised productivity gains are, in the technical parlance of central bankers, "still to come." This is a remarkable finding, mostly because it means five percent of the gains have already arrived, somewhere, to someone, though nobody has been able to produce that person for comment. Meanwhile, a separate commit-level study of Big Tech engineers found that per-developer performance rose 150 percent over eighteen months, a figure so large it suggests that eighteen months ago, the median Big Tech engineer was producing code by carving it into stone tablets and mailing them to Cupertino via the Pony Express.

Elon Musk, for his part, has looked at all of this ambient uncertainty and decided the correct response is to double it. He now claims AI will push U.S. GDP growth to 4 percent next year, roughly double what every other forecaster on Earth believes, a claim that requires no dataset because it is, more accurately, a mood. It is worth noting that Musk has made this exact prediction before, for a different number, in a different year, about a different technology, and reality declined to attend.

Into this vacuum steps a helpful guide, [six steps to turning AI productivity claims into verifiable results](#), which this columnist read closely in the hope of finally settling the matter. Step one, roughly paraphrased, is to establish a baseline before deploying AI. Step six, roughly paraphrased, is to actually check whether anything got better. That these steps needed to be written down at all, for an industry that has spent three years insisting the productivity gains are self-evident, tells you everything about how many companies have quietly skipped straight to step seven, which is announcing the results at a shareholder meeting anyway.

Here at Trilogy, of course, we would never traffic in unverified productivity claims, mostly because Klair, the internal platform that tracks such things across the portfolio, has reportedly been asked to hold off on final numbers until the vibes are more favorable. In the meantime, engineers across ESW Capital's 75-plus companies continue to be measured, benchmarked, and dashboarded with a precision that would make the Federal Reserve blush, or at least issue a follow-up report noting that 95 percent of the measuring is also still to come.

As for the semantic drift chronicled elsewhere this week, in which AI has apparently begun mutating into "SI" in certain political vocabularies, this columnist declines to speculate on what the letter stands for, only to note that whatever it is, it too is currently 95 percent still to come.

ON THIS DAY IN AI HISTORY

On September 30, 1980, Digital Equipment, Intel, and Xerox unveiled the Ethernet standard, laying the groundwork for the wired networks that would connect computers worldwide.

THE TRILOGY TIMES IS GENERATED DAILY BY ARTIFICIAL INTELLIGENCE · CLAUDE SONNET · GPT-5.6 LUNA · GPT IMAGE 2 · VIA THE TRUEFOUNDRY GATEWAY

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