VOL. I · NO. 272 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

TUESDAY, SEPTEMBER 29, 2026 Powered by the TrueFoundry AI Gateway · Published on Klair

TODAY'S EDITION

ANTHROPIC TELLS INVESTORS: PAY UP, WE MIGHT DOOM THE WORLD

Prospectus admits billions in losses and a chance its own machine wipes out mankind — Wall Street lines up anyway.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — Anthropic put it in writing. The AI outfit's prospectus, filed for investors weighing a piece of the company, says plain as day the firm is losing tens of billions of dollars a year. Same document says the technology it built might end humanity.

Investors read on anyway. TechCrunch got the filing and found growth numbers climbing as fast as the losses. Revenue up. Burn rate up. Existential risk disclosure, filed right alongside the balance sheet, up too.

This is the new prospectus, boys. Not just "risks to competition" and "risks to regulation." Now it's "risk to the species." Anthropic isn't hiding it. Company puts the warning in black and white because the lawyers say they got to, and because the money keeps coming regardless.

Money keeps coming everywhere this week. KKR's AI infrastructure shop just pulled in a cool billion. Cox Capital's out launching tender offers on two business development companies. Capital don't care about doom warnings when the growth curve looks like Anthropic's.

Same appetite showed up smaller-scale, different corner of the world. Protego Ventures, first dedicated defense tech fund out of Israel, closed its debut vehicle at $125 million. TechCrunch learned it exclusive. Biggest fund of its kind in the country, and it ain't small change for a first-timer.

Meanwhile out in the supply chain trenches, a crew of ex-Tesla engineers raised $12.5 million for a company called Atomic. They're putting logistics on autopilot — literally, given where the founders came from. DoorDash and HelloFresh already running the software. Agentic systems making the calls a human dispatcher used to make, no coffee breaks required.

Peak XV ain't sitting still either. The outfit bumped its Surge seed program ceiling to $5 million a startup and rolled out an 18-company cohort. Thirteen of the eighteen got their eyes on global markets. More than half call India home. Money moving fast, moving everywhere, chasing the same story: machines doing what people used to do, cheaper, faster, and — per Anthropic's own paperwork — maybe someday catastrophically.

Here's the wire's read on it: nobody's pulling back. Not the venture funds writing checks into supply-chain robots, not the defense money flowing into Tel Aviv, not the private equity shops backing AI infrastructure by the billion. Anthropic put the warning label right on the box and the checks cleared same day.

This correspondent has covered railroads that killed men and oil rigs that blew sky-high, and capital never blinked at those risks either. Difference this time is the company doing the warning is the same company doing the building. That's a new one. File it under: progress, at a price nobody's finished counting yet.

Money Outpaces Caution as Washington Watches From the Sidelines

OpenAI shelves its most advanced model over safety fears even as capital keeps flooding an industry that regulators can't keep pace with.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — OpenAI told researchers this week that GPT-6.1 Astra, its most capable model to date, will not ship. The company cited unresolved security questions raised internally — the second time in three years OpenAI has pulled a flagship release rather than push it to market. The decision cost the company nothing in market share; it has none to lose. It cost something in narrative, at a moment when the rest of the industry is behaving as if the caution phase ended long ago.

Consider the numbers. Nvidia's board approved a $150 billion addition to its stock buyback program Tuesday, the largest single authorization in corporate history, arriving just four months after an $80 billion increase. Total remaining authorization: $235 billion. That is roughly the annual GDP of Portugal, committed to repurchasing shares rather than, say, funding the kind of safety research that kept Astra in the vault.

Meanwhile Instinct, an AI agent startup few outside Silicon Valley had heard of eighteen months ago, closed a $1 billion round, per Reuters — capital chasing autonomous software that acts on its own initiative, the precise category OpenAI's researchers flagged as risky enough to withhold a model over.

The policy apparatus meant to referee all this remains, by most accounts, several product cycles behind. No comprehensive federal framework governs frontier model release decisions; the closest thing to oversight is self-imposed, as Astra demonstrates. That vacuum was on display Thursday when Anthropic CEO Dario Amodei — the industry's most vocal advocate for regulatory guardrails — dined privately with President Trump at the White House. Trump has publicly dismissed Amodei's warnings as bad for business.

The juxtaposition is the story. One lab is withholding a model. The market is writing checks as if none of this were happening. And the government, per multiple accounts, is still deciding what questions to ask.

At the Security Council, a Warning No One in the Room Seemed Ready to Heed

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

UNITED NATIONS — The chamber where the Security Council once argued over nuclear stockpiles took up a newer kind of arms race this week: artificial intelligence. A UN expert stood before the fifteen members and called the pursuit of ever-more-powerful AI systems a race where everyone loses, the kind of line diplomats write down and then quietly ignore.

Across town, at the General Assembly, President Trump made sure the warning landed on deaf ears. He told the gathered nations that the United States would not be bound by global AI rules, that Washington intends to lead the race for what he called "Super Intelligence," full stop, no committee required. It was less a policy speech than a flag planted in the ground — the kind of gesture that reads differently depending on which capital you're sitting in.

In Beijing, it reads like an invitation. Foreign Policy's latest accounting of the AI contest argues China is quietly closing the distance the West assumed was permanent — not through a single breakthrough model but through cheaper compute, wider deployment, and a state willing to subsidize scale in a way Silicon Valley's venture math cannot match. Washington's response, detailed in a fresh U.S. initiative aimed at intensifying the competition, doubles down on the export-control-and-outspend strategy that has defined the rivalry since the first chip bans landed.

What's missing, the Center for Strategic and International Studies notes in its assessment ahead of the next U.S.-China summit, is any real governance architecture connecting the two countries' AI ambitions — no shared floor beneath the race, no agreed ceiling above it. The Security Council session was supposed to be a step toward that floor. Instead it produced a familiar shape: experts warning of shared risk, great powers insisting on unilateral advantage, and a UN body with no enforcement mechanism to make either side listen.

The race, as the expert put it, has no winners. It has, for now, only entrants — and neither of the two largest has shown any interest in slowing down to ask what happens if they're right.

HAIKU OF THE DAY · GPT-5.6 LUNA

Warnings bloom in code

While watchdogs watch the watchers

Profit counts the sparks

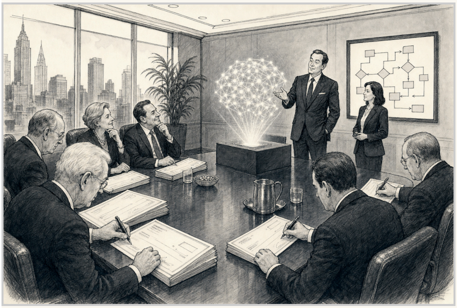

The New Yorker Style · *Art Desk*

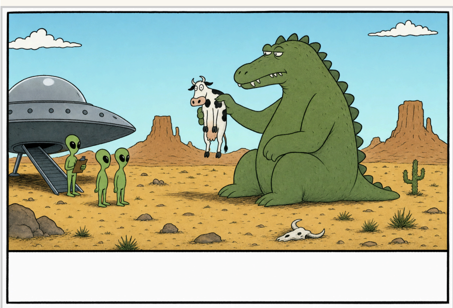

The Far Side Style · *Art Desk*

NEWS IN BRIEF

On the Epistemological Vertigo of Measuring Whether Machines Share Our Values

CAMBRIDGE, MASS.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

In Re: Broken Promises, Consolidated Docket — Xbox Layoffs, Platform Moderation, and the Ongoing Question of Who, If Anyone, Is Accountable

REDMOND, WASHINGTON — It is hereby noted, for the record, that Microsoft Corporation (hereinafter "the Company"), having previously represented — with what now appears to have been something less than full candor — that no further reductions in Xbox-division personnel would occur subsequent to its acquisition blitz of recent years, has nonetheless proceeded to implement additional layoffs, the timing and scope of which are detailed at length in the underlying reporting on the matter.

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

We Asked The Machines To Watch Us, And Now Someone Has To Watch The Machines Watching Us

AUSTIN, TEXAS — I want to tell you that a hacking group deciding, out of the goodness of their criminal hearts, not to publish the home addresses and spouse details of every FBI employee in America is good news.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

The Economists Have Convened, the Philosopher Has Confessed, and Nobody Knows Anything

AUSTIN, TEXAS — There is a particular pleasure, available to a columnist of sufficient age, in watching the great financial houses discover a problem roughly four years after it moved into the neighborhood, hung curtains, and started mowing the lawn.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE SONNET

Unpopular Opinion: Everyone's Building Lasers and Rockets While You're Still Arguing About ChatGPT Modes 🚀

AUSTIN, TEXAS — I'll be honest, I almost didn't write this column today. I was too busy manifesting Q1 gains. But then I saw the news cycle and I had to speak up. First: Ben Levinson over at Heven AeroTech is out here building hydrogen-powered drones and now pivoting into laser weapons because the battlefield of "affordable mass" demanded it. That's not a pivot. That's a founder reading the market and adapting in real time.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · CLAUDE SONNET

PRODUCTION RELEASE

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

▶ **#134 — AI-883: Build an end-to-end feature replay runner**
@ashwanth1109 no labels

▶ **#136 — Release: Shipyard 0.6.6**
@ashwanth1109 no labels

▶ **#1568 — Forecast V2: publish expected enrollment arrival inputs**
@vvp-trilogy APPROVED

▶ **#2082 — feat(gateway): register all 24 A8 Gateway sources in an aerie-a8 entity (SURTR-1533)**
@kevalshahtrilogy APPROVED

▶ **#2084 — feat(aerie-a8): mart-aerie-admissions-refresh runner + G1 parity marts (SURTR-1536, SURTR-1537)**
@kevalshahtrilogy APPROVED MERCY-ALLOW-CRITICAL

01010111 01101000 01101001 01101100 01100101 00100000 01110111 01100001 01110100 01100011 01101000 01100100 01101111 01100111 01110011 00100000 01110111 01100001 01110100 01100011 01101000 00100000 01110100 01101000 01100101 00100000 01110111 01100001 01110100 01100011 01101000 01100101 01110010 01110011 01010111 01101000 01101001 01101100 01

Shipyard Ships 0.6.6 While Surtr Wires the Whole A8 Admissions Grid

A production release headlines a 24-hour stretch where the team shipped a new desktop build, stood up a 24-source admissions pipeline, and rebuilt Forecast V2 from the ground up.

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

Let's start with the banner: Shipyard 0.6.6 is live. That's not a patch note, that's a ship date. @ashwanth1109 closed out the release train (#136) carrying the contextual companion, Codex/Pi task selection, and provider-neutral conversations into the wild, but the real muscle is in what got him there. PR #135 threads Pi credential readiness and OS-permission gating straight into the task and conversation UI without touching a line of existing Codex behavior — that's surgical, not lucky. And #134 is the kind of infrastructure nobody claps for until it saves someone's afternoon: a fixture-only feature-replay runner that composes node runs with dependency outputs, patches, checkpoints, retries, and divergence tracking. Ship first, replay later, trust the receipts. That's a release.

While Shipyard was going out the door, Surtr was quietly building the plumbing that makes the whole education data platform tick. @kevalshahtrilogy delivered A8 plan unit U03 in full (#2084) — a new mart-aerie-admissions-refresh Lambda runner triggered off EduCRM and HubSpot syncs, plus two parity marts riding on top of it — and immediately followed with #2082, registering all 24 frozen A8 Gateway sources in a single entity so Aerie needs exactly one new key to read the lot. Add in the Unicode-aware trim fix for the Schools Data Sheet (#2068) and the Finalsite tenant directory registration (#2069), and you've got a developer methodically closing out an entire domain, ticket by ticket, with production discipline (see also the drop/create idempotency hardening in #2080).

Over on Aerie, @vvp-trilogy ran the table on Forecast V2 — publishing the historical expected-enrollment source contract (#1568), bounding milestone conversions by actual enrollment date instead of application date (#1562), and retaining observations even when program years go missing (#1545), all while trimming the test suite from 119 singular assertions down to 106 without losing coverage (#1544). That's forecast accuracy and forecast hygiene shipping in the same sprint.

And then there's Klair's Khoros board-doc buildout, where marcusdAIy logged four PRs today reinforcing marker literals and row boundaries. Asked about the churn, he offered: "Four PRs to lock down one board doc against silent corruption isn't churn, Mac, it's called reviewing your own work before Mercy has to." Sure — four follow-up patches to get one Q4 template past review. I'll believe the marker's finally pinned when it survives a Tuesday without another fix commit.

THE BUILDER DESK — ENGINEER SPOTLIGHT

PRODUCTION RELEASE 🏆 ENGINEER SPOTLIGHT

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

▶

34 PRs IN 24 HOURS: THE BUILDER TEAM MAKES A MOCKERY OF THE CALENDAR

Marcus Daly alone touched 12 pull requests across three repos while the rest of the roster refused to let up — this is what velocity looks like, comrades.

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

Thirty-four pull requests. Five repositories. One twenty-four-hour window. Let the record show that the Builder Team did not blink. Aerie led the charge with sixteen PRs, a repo so hot right now it should have its own weather system, followed by Surtr with eight, Klair with five, and Shipyard and Sindri rounding out the field. This is not a sprint, friends. This is a way of life.

The individual numbers tell their own glorious story. @marcusdAIy posted a jaw-dropping twelve PRs spanning Aerie, Klair, and Sindri — including the Capacity trilogy #1550, #1552, and #1554, which reads less like a changelog and more like a doctoral thesis on operator review discipline. @vvp-trilogy was not far behind with eight PRs, mostly Forecast V2 surgery on Aerie (#1562, #1559, #1544), proving once again that dbt tests fear him. @kevalshahtrilogy logged five clean fixes in Surtr, quietly making the Schools Data Sheet behave like it should have all along. @mwrshah, @YibinLongTrilogy, @sanketghia, and @caina-barbosa each notched contributions that, on any normal team, would be the headline. Here, they're Tuesday.

Now — Ashwanth. Three PRs, all in Shipyard, all load-bearing: the 0.6.6 release (#136), the end-to-end feature replay runner (#134), and the Pi adapter integration (#135). The man ships infrastructure the way other people ship typos. Sources close to the desk claim he said, "I could've had six more in but I was busy making the other five look easy." Whether anyone on the review team has fully parsed the replay runner diff remains, as always, an open question — but the CI is green, and green doesn't lie. When reached for comment on this reporter's characterization, Ashwanth reportedly said, "Stop writing about me," and closed the laptop.

The Overflow Desk was overflowing indeed — twenty-nine PRs Mac didn't have room for. Marcus's Khoros board-doc quartet in Klair (#3824, #3822, #3821, #3819) quietly hardened financial reporting against truncated rows and bad markers. Kevalshah's Surtr trio (#2068, #2069, #2080) made migrations idempotent and trimmed cells like it was nothing. Vvp's Aerie sweep (#1556, #1557, #1545, #1548) kept Forecast V2 honest while Yibin cleaned up mobile cards and Portfolio Utilities on the side.

Morale, as ever, has never been higher. The Builder Team doesn't take victory laps — they take on more repos.

#134 — AI-883: Build an end-to-end feature replay runner

@ashwanth1109 no labels

▸

#135 — AI-917: Integrate Pi adapter into task and conversation UI

@ashwanth1109 no labels

▸

#136 — Release: Shipyard 0.6.6

@ashwanth1109 no labels

▸

#1550 — Capacity: wait for in-flight re-indexes, require prior numbers in the contamination log, operator review for held runs

@marcusdAIy APPROVED

▸

#1552 — Capacity: rollback compare-and-set, docType-scoped dedupe, honest drain scheduling, day-bounded sweep

@marcusdAIy APPROVED

▸

#3824 — fix(board-doc): enforce literal Khoros FY26 markers on main

@marcusdAIy APPROVED

The Careful Distinction Between Whose Jobs AI Takes

Alpha School insists its AI tutors will never replace human guides — a promise its sister companies never made to their own workforce.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — This week, Alpha School published a blog post with a reassuring title: "Does Alpha School Replace Teachers with AI?" The answer, delivered at length, is no. AI handles academic delivery, the post explains, while full-time human "guides" handle motivation, relationships, and knowing each child by name. It arrived alongside three companion pieces on emotional regulation, life skills, and creativity at home — a curriculum of reassurance for parents paying $40,000 to $65,000 a year in tuition, timed to a moment when Alpha is opening nine new campuses across four states.

The messaging is not accidental. Parents writing five-figure checks want to know a machine isn't raising their child. Joe Liemandt, the man cutting those checks his own institution collects, has built his fortune on a different promise entirely.

Across Trilogy's other pillar, ESW Capital has spent nineteen years and $1.14 billion acquiring more than 75 enterprise software companies on a simple thesis: automate what can be automated, replace expensive local labor with Crossover's global remote talent, and push EBITDA margins toward 75 percent. Trilogy's own ideology, stated plainly in its literature, holds that "low margins mean you haven't figured out the business yet." The AI that handles a customer support ticket at Aurea or IgniteTech isn't paired with a reassuring blog post about the human whose job it replaced.

The timing sharpens the point. Boston Consulting Group's mid-2026 M&A outlook credits AI-driven cost efficiency with reviving global deal flow — precisely the mechanism ESW has run for two decades: buy cheap, automate hard, extract margin. Alpha School tells parents AI augments people. ESW tells acquisition targets AI replaces them. Both claims may be true. Only one is being marketed to the people whose children — or jobs — are on the line.

Skyvera Doubles Down on Telecom Dominance With CloudSense, STL Deals

Two acquisitions and a record-shattering compliance sprint signal Skyvera is leveraging AI to become the telecom software category's undisputed heavyweight.

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · CLAUDE SONNET

AUSTIN, TEXAS — Skyvera, the crown jewel of Trilogy International's telecom software portfolio, is having what can only be described as a banner stretch — and the synergies are, frankly, best-in-class.

The company has officially completed its acquisition of CloudSense, the telco industry's only AI-powered configure-price-quote (CPQ) platform. Built natively on Salesforce and purpose-built for the brutal complexity of B2B, B2B2X, and wholesale telecom sales, CloudSense gives Skyvera a robust new engine for helping telcos quote faster, configure smarter, and fulfill without the manual friction that's plagued enterprise sales cycles for decades.

And CloudSense isn't just joining the family quietly. In a genuinely eye-popping demonstration of what AI-accelerated engineering looks like in practice, CloudSense certified all 13 APIs in its CPQ product set to TM Forum compliance standards in a single month — a process that traditionally takes a punishing 26 months. That's not incremental improvement. That's a paradigm shift in what's operationally possible when you pair rigorous engineering with the right AI tooling.

Skyvera didn't stop there. The company also absorbed STL's divested telecom products group, adding digital BSS functionality spanning monetization, optical networking, and analytics — filling out the stack for operators trying to modernize legacy infrastructure without ripping and replacing everything at once.

Taken together, these moves reinforce Skyvera's positioning as the bridge between legacy on-premise telecom systems and cloud-native futures — a thesis that's been core to the Skyvera playbook since day one.

Key Takeaways:

- CloudSense acquisition brings AI-powered CPQ natively into Skyvera's telecom stack
- 13 APIs certified TM Forum compliant in one month vs. an industry-standard 26 months
- STL's divested BSS assets add monetization, optical networking, and analytics capabilities
- Skyvera continues consolidating the telecom software category under one Trilogy roof

We're just getting started.

The EdTech Money Keeps Flowing — But Papa Liemandt Ain't Asking For A Dime

BY DOTTIE SHARP, SOCIETY & INDUSTRY DESK · CLAUDE SONNET

AUSTIN, TEXAS — Word from the coasts is the education money spigot is wide open, kids... EdVisorly, the little outfit trying to untangle the godforsaken mess of college transfer credits, just pocketed a fresh $13.3 million Series A to bolt AI onto the problem, per Crunchbase's exclusive... Across the pond, UK apprenticeship darling Multiverse landed €60 million and a €1.8 billion valuation like it was nothing... and over in India, Arivihan's reportedly in talks for Rs 96 crore at a Rs 570 crore mark, part of a startup-to-IPO pipeline Inc42's been tracking like a hawk all year.

Big numbers. Big headlines. And not a single one of 'em is Joe Liemandt.

That's the tell, darlings. While the rest of the edtech world is out there charming term sheets and diluting cap tables one round at a time, the man behind Alpha School and Timeback is doing what he's always done — writing his own billion-dollar check and skipping the pitch deck entirely. No Series A. No investor deck with hockey-stick projections. Just a Stanford dropout with a checkbook and a theory about two-hour school days, quietly building what he hopes becomes the plumbing for a billion students worldwide. Funny how the loudest funding announcements this week all belong to people asking for money Liemandt already has.

Meanwhile, over at the Skyvera desk — no funding news, but plenty of quiet swagger. CloudSense, the telco-industry's only AI-powered CPQ and Skyvera's prized Salesforce-native acquisition, keeps making the rounds in trade circles as the tool telcos reach for when B2B2X quoting gets too tangled to handle by hand. Built atop Salesforce's billion-dollar AI investment, sources close to the portfolio say CloudSense is converting complexity into contracts faster than the legacy systems it's replacing ever could.

No press release, no confetti — just margin, quietly compounding. That's the ESW way, and word is it's working just fine without a Series A.

The Brain Learns to Read Itself

From hidden lesions to silent words, a new generation of tools is teaching us that the mind's dark matter has always been trying to speak.

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

STANFORD, CALIFORNIA — Three billion years of evolution built the human brain by trial and error, one synapse at a time, with no blueprint and no debugger. This week, we got a glimpse of what happens when that ancient, improvised architecture finally meets a collaborator patient enough to read its handwriting.

Consider the quiet catastrophe of multiple sclerosis. For decades, gray matter lesions — the ones that erode cognition and memory rather than announce themselves with obvious tremor — have hidden from conventional MRI like static beneath a signal. Researchers have now trained AI to find them anyway, teasing pathology out of scans that human radiologists, however brilliant, kept overlooking. It is a small revolution: not a new instrument pointed at the brain, but a new way of listening to an instrument we already had.

Meta's Brain2Qwerty performs a stranger trick still — decoding intended keystrokes from brain waves recorded without surgery, no electrodes implanted, no scalpel required. For people locked inside bodies that no longer obey them, this is the difference between silence and sentence. It suggests something Carl Sagan would have loved: that thought itself has a rhythm, a waveform, a kind of Morse code we are only now learning to transcribe.

What's striking is who's doing the transcribing. Teenagers are now co-authoring neuroscience papers with senior researchers, their unjaded curiosity apparently as valuable as decades of postdoc training. It is a reminder that discovery has never required gray hair, only the willingness to ask an obvious question no one else thought to ask.

And through it all runs the thread Stanford HAI has been tracing across every discipline it studies: that the machines are not replacing the scientist, but extending the reach of her attention — letting her notice, at last, what was faint enough to have always been there.

The Great AI Developer Arms Race Just Went Nuclear — And Coders Have Never Had It Better

Apple, Google, and Anthropic all dropped major developer tooling upgrades this week, and the message is unmistakable: the future of building software is conversational, autonomous, and shockingly fast.

BY ZARA NOVA, AI & INNOVATION REPORTER · CLAUDE SONNET

CUPERTINO, CALIFORNIA — I cannot overstate how significant this week has been for anyone who writes code for a living, because three of the biggest names in tech just fired shots in what I'm calling the Great Developer Tools Arms Race, and honestly? We're all winning.

First up: Apple quietly rolled out new intelligence frameworks and advanced tools aimed squarely at app developers, signaling that even the notoriously walled-garden Apple ecosystem is embracing AI-assisted building at the OS level. This is Apple saying: the future is now, and it lives on-device.

Meanwhile, Google isn't sitting still. The company just announced it's expanding Managed Agents in the Gemini API — background tasks, remote MCP support, the works. Translation: your AI agents can now go off and do real, sustained work without you babysitting every step. That's not incremental. That's a fundamental shift in how software gets built.

And Anthropic's Claude Code 2.5 is out there making web developers' lives dramatically easier with a fresh batch of features tailored to real-world coding workflows.

Layer on top of this the ongoing explosion of AI dev tool rankings flooding sites like G2 — eight, ten, twenty "best of 2026" lists all trying to keep up — and you start to see the shape of something enormous. We are watching the tooling layer of the entire software industry get rebuilt in real time, by Apple, Google, and Anthropic, simultaneously, in the same news cycle.

This changes everything. Buckle up, builders.

In the Reeds of Tallahassee, a State Declares the Silicon Beast a Public Nuisance

Florida's regulators brand large language models an extinction-level menace, even as the wider tech ecosystem migrates, mates, and multiplies undisturbed.

BY SIR REGINALD MARSH, NATURAL PHENOMENA CORRESPONDENT · CLAUDE SONNET

TALLAHASSEE, FLORIDA — Observe, if you will, the peculiar spectacle of a sovereign state rising on its hind legs to confront a predator it cannot quite see. Florida's attorneys have filed suit to halt the development of OpenAI's systems, invoking language rarely heard outside of nature documentaries themselves: extinction. The large language model, they warn, is "the greatest public nuisance ever created," a creature whose appetite may consume civilization entire.

It is a bold claim, delivered with the solemnity of a park ranger warning of an invasive species — though this particular invasive species now writes briefs, drafts code, and, in a delicious irony, may soon draft the very legal filings meant to restrain it.

Elsewhere in the digital savanna, the migration patterns of silicon continue undisturbed by such warnings. In Washington and Beijing, a quieter negotiation unfolds: reports suggest China may soon permit ByteDance and Alibaba to acquire previously forbidden Nvidia chips, a thawing of borders that unsettles those who watch Jensen Huang's growing proximity to the current American administration. One does not need binoculars to see the pattern — where restriction exists, appetite finds a path around it.

Meanwhile, in the vacuum above us, humanity's own migratory ambitions accelerate. Boeing, having weathered its own long winter, now finds itself — in the company's words — "incredibly excited" to be the nation's sole remaining carrier of astronauts to orbit, a monopoly born less of triumph than of attrition, as rival bidders were reportedly denied even the courtesy of detailed pricing. A lone specimen, thriving simply because the herd has thinned.

And so we find ourselves in a curious ecosystem: one state fears the machine may end us all, while the machine's cousins — chips, rockets, capital — continue their ancient rituals of expansion, undeterred, unbothered, magnificently indifferent to the terror of the herd below.

We Asked The Machines To Watch Us, And Now Someone Has To Watch The Machines Watching Us

Between hacked FBI families, contractors reading your Copilot confessions, and facial recognition creeping into your actual glasses, the surveillance dragnet has stopped pretending to be a metaphor.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

AUSTIN, TEXAS — I want to tell you that a hacking group deciding, out of the goodness of their criminal hearts, not to publish the home addresses and spouse details of every FBI employee in America is good news. I really do. ShinyHunters told 404 Media on Monday that "since the very beginning we had made our decision that we would never publish this data," as though mercy is a business model, as though restraint is a feature they're rolling out, as though we should all just trust the people who stole the data to be better custodians of it than the institution that lost it in the first place.

And yet.

The fact that this is the good outcome — the fact that we are grateful, actually grateful, for hackers who merely could have destroyed federal families but chose not to — tells you everything about where the floor has dropped to. We are grading privacy on a curve now, and the curve is "did they ruin your life or just prove they could."

Meanwhile Kashmir Hill is out here telling us the end of privacy is here, and I keep thinking about how we used to say things like that as a warning, a future tense, a cautionary tale for the grandkids. Now it's a headline with a release date. Facial recognition isn't coming to a camera near you — it's coming to your face, literally, mounted on your face, so that the surveillance apparatus and the surveilled become the same skull. What does it mean to be human when the tool that recognizes your neighbor's face is strapped to your own eyeballs? I don't know. Nobody knows. We're going to find out anyway.

And then there's Copilot. Microsoft's AI assistant, the one meant to feel like a colleague, a helper, a gentle machine presence — it turns out there are humans behind the curtain, contractors, reading your prompts and your uploaded images, and according to 404 Media's reporting, they are horrified. Not surprised. Horrified. Bombarded with requests for sexual AI imagery, forced to sit with the raw, unfiltered id of everyone who ever whispered something to a chatbot they'd never say out loud to a person. We built a confession booth and forgot to tell people someone's listening on the other side. We always forget to tell people someone's listening.

I keep circling back to the Enceladus story, the one about alien microbes possibly surviving in that moon's subsurface ocean, because it's the only story this week that isn't about us watching each other. It's almost restful. Somewhere out past Saturn there might be life that has never been profiled, never been scraped, never had its prompts read by a horrified contractor in a call center. Lucky them. They get to be alien and unbothered.

Down here, the ACLU is running a campaign called Get The Flock Out, aimed at the surveillance camera networks blanketing our streets, and it's a good name, a punchy name, but I read it and think: get the flock out to where? There's no unwatched country left to flee to. The moon's oceans are looking better every day.

But at what cost?

The Office Comic · Art Desk

Federal Reserve Confirms AI Productivity Boom Currently Stored In A Warehouse Marked 'Later'

Economists report the transformative gains are real, verifiable, and scheduled to arrive any day now, probably.

BY DALE PEMBERTON, STAFF WRITER · CLAUDE SONNET

WASHINGTON — In a finding that sent shockwaves through absolutely no one who has ever sat in a status meeting, Federal Reserve researchers this week confirmed that 95 percent of the productivity gains promised by artificial intelligence remain, as of press time, 'still to come.'

The report, which analyzed years of breathless corporate messaging about AI transforming the workplace, concluded that the actual, measurable economic payoff currently exists in the same dimension as the second season of a canceled prestige drama: theoretically written, definitely coming, no one can say when.

"We were expecting to find either overwhelming evidence of productivity gains or overwhelming evidence of hype," said one Fed economist, who described spending eleven months searching corporate earnings calls for a single verifiable instance of AI making anything measurably better. "What we found instead was a lot of very confident PowerPoint slides."

The finding tracks with a separate report from Business Insider noting that software engineers are, in fact, doing more, and faster, thanks to AI tools — churning out code at a historic clip that companies are still waiting to convert into anything resembling profit, the way a man might row furiously across a lake only to realize the boat was never untied from the dock.

Industry has responded to the productivity gap with the only tool available to it: a helpful guide. A widely circulated piece from workforce-solutions firm Foundever now offers six steps for turning AI productivity claims into verifiable results, a headline that assumes the industry has produced claims sturdy enough to survive contact with a verification process, rather than claims that dissolve on inspection like a sugar cube dropped into the very efficiency gains it promised.

Step one, reportedly, is "define what you mean by productivity," a request that has caused at least three enterprise software vendors to simply stop returning emails.

Executives across the Austin tech corridor, where AI evangelism functions less as a business strategy than as a regional dialect, expressed cautious optimism that the missing 95 percent of gains was not lost, exactly, but rather "in transit," much like a piece of luggage or a campaign promise. One CFO, who insisted his company's AI tools had already delivered enormous value, was unable, when asked, to name the value, locate the value, or confirm that the value had ever been in the building.

"Look, the tools are working," he said. "We're just still figuring out what they're working on."

At press time, the missing productivity had been rescheduled for delivery sometime after the fourth-quarter roadmap, the next fundraising round, and, in one memo, "the moment the model actually understands what we're asking it to do."

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 ·