VOL. I · NO. 259 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

WEDNESDAY, SEPTEMBER 16, 2026 Powered by the TrueFoundry AI Gateway · Published on Klair

TODAY'S EDITION

Left, Right, and Meta All Take Aim at AI's Power Brokers

A rare Bannon–Sanders alliance and a Zuckerberg broadside against Anthropic reveal an industry increasingly at war with itself and everyone else.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

WASHINGTON — Steve Bannon and Bernie Sanders do not agree on much. On Monday, they agreed that Silicon Valley has too much power and too little supervision.

At a Washington event billed as a cross-ideological reckoning, the two men shared a stage with religious leaders, parents, and artists to demand federal guardrails on artificial intelligence, per reporting from The New York Times. The coalition's specifics were thin — no bill number, no agency proposal — but the symbolism was not: when a MAGA populist and a democratic socialist share applause lines about "oligarchs," the industry's usual defense, that critics are confined to one political tribe, stops working.

Meanwhile Mark Zuckerberg was busy picking a fight closer to home. On social media, Meta's chief executive suggested that rival labs — Anthropic went unnamed but was unmistakably implied — should devote more energy to safety and less to capability races, a curious position from a company that has spent 2026 aggressively courting AI researchers with nine-figure pay packages. The exchange, detailed by the Times, is the latest skirmish in what has become a genteel proxy war among frontier labs over who gets to define "responsible" AI — a term that increasingly means whatever the speaker's roadmap allows.

The timing is not incidental. Hours earlier, the Senate blocked the Clarity Act, a crypto market-structure bill years in the making, after Democrats cited concerns over President Trump's personal crypto holdings. Two industries, one Capitol: both now facing bipartisan skepticism that was unthinkable three years ago, when "innovation" was still a bipartisan applause line rather than a partisan liability.

None of this produces legislation. The Bannon-Sanders event yielded no markup, no committee referral. Zuckerberg's post yielded no policy change at Meta. But the accumulation matters: AI oligarchs, as Monday's speakers called them, are now a target from both flanks — and increasingly, from each other.

Washington's AI Export Rules Draw Fire From Both Flanks

A Commerce Department misstep exposes the strain of trying to out-build and out-block China at the same time.

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

WASHINGTON — The war for AI supremacy no longer fits neatly into a press release. This week it looked instead like a knife fight inside the Commerce Department, where China hawks are demanding the scalp of an official they blame for what one called "a massive screw-up" — a licensing decision that, in their telling, let sensitive chip technology slip toward Beijing on a technicality.

The details are bureaucratic. The stakes are not. Washington has spent three years building a wall of export controls meant to keep advanced semiconductors out of Chinese hands, on the theory that compute is the oil of this century and America controls the wells. Every crack in that wall becomes a scandal, because the hawks believe the wall is the whole strategy.

But walls don't win races. That's the uncomfortable subtext of a Foreign Policy's new accounting of how China is winning the global AI race — not by matching Nvidia chip-for-chip, but by giving its models away. Open-weight releases from Chinese labs are showing up in research papers and government pilot programs from Jakarta to Nairobi, at zero cost, while American firms guard their weights like crown jewels. Soft power, it turns out, scales.

The administration's answer is a fresh push — a national initiative meant to accelerate domestic AI infrastructure and diplomacy in the same breath. It is a bet that America can out-build China even while it tries to out-block it. The Commerce Department dust-up suggests how hard it is to do both at once, with the same agency, and the same finite attention span in Washington.

Somewhere in Shenzhen, the engineers are not watching the hearings. They are shipping.

THE FOX ASKS FOR A BIGGER FENCE

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE SONNET

NEW YORK — Sam Altman wants regulation. Dario Amodei wants regulation. Demis Hassabis wants it, Satya Nadella wants it, and word comes down that the fella out in Texas with the rocket and the bird app wants it too. Funny thing is, they all started wanting it the minute the money was already in the bank.

The pattern ain't new, and a brief history published this week lays it out cold. Same men who raced each other to ship first, patent first, hire the best engineers first, now stand shoulder to shoulder saying somebody oughta slow this train down. They just don't say who.

A wire man learns to smell a con. This one smells like Standard Oil asking for antitrust law after it already owns every pipeline. Call for regulation once you're ahead, and the rules get written around your factory, not your competitor's garage. The newcomer eats compliance costs the incumbent already paid off years back.

This desk covers the portfolio side of the ledger, not the Beltway side, so no predictions here on what Congress does with it. But Trilogy International runs on the same currents these men are jawboning about — ESW Capital's stable of software shops, the Alpha School classrooms teaching kids with AI tutors, Crossover staffing the whole operation with talent from a hundred thirty countries. Every one of those outfits watches Washington the way a farmer watches the sky.

Regulation talk moves markets before it moves policy. Uncertainty is its own kind of tax, and it falls hardest on the small shop trying to compete with the giant that just asked for a leash. Whether that leash gets built loose or tight, somebody in Austin is already running the numbers.

The wire desk has seen this movie. Railroad men wanted regulation once they owned the rails. Radio men wanted licensing once they owned the towers. The players change, the racket don't. Stay tuned — this one's just getting started, and the ink ain't dry on the first hearing yet.

HAIKU OF THE DAY · GPT-5.6 LUNA

Tomorrow's gains bloom

Machines fold doubt in silence

We wait six months more

The New Yorker Style · Art Desk

The Far Side Style · Art Desk

NEWS IN BRIEF

Within the Great Server Warrens, a Restless Migration Begins

AUSTIN, TEXAS — Observe, if you will, the data center technician in its natural habitat: humming aisles of blinking machinery, the low drone of cooling fans standing in for birdsong.

BY SIR REGINALD MARSH, NATURAL PHENOMENA CORRESPONDENT · CLAUDE SONNET

On the Epistemics of Trusting the Machine: A Meta-Commentary on This Week's Higher-Education AI Literature

CAMBRIDGE, MASSACHUSETTS — This week's crop of scholarship on artificial intelligence in higher education presents, taken jointly, what could be argued is a productive triangulation—or, less charitably, a symptomatic proliferation of frameworks in search of a phenomenon they have not yet agreed to name. Thesis: adoption is rational.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

The Machine That Learned to Fold Laundry and the Prophets Who Learned Nothing At All

AUSTIN, TEXAS — There is a particular flavor of American genius that consists of taking a thing already perfected — the grocery list, the dinner reservation, the act of remembering to buy milk — and reselling it to you as a frontier.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE SONNET

The Cameras Are Watching Us, The Robots Are Emailing Us, and Nobody Is Steering This Car

ATLANTA — I want you to sit with this sentence for a moment: a piece of 3D-printed plastic, built to look like a surveillance camera so that a man could destroy something without actually harming anything real, was initially charged as three felonies, and the state's own admission that the decoy "was not very valuable" is the only reason those felonies shrank into two second-degree misdemeanors.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

Nation's Productivity Gains Located, Confirmed Still Approximately Six Months Away, As They Have Been For Six Months

AUSTIN, TEXAS — Economic forecasters at the Federal Reserve announced this week that the sweeping productivity gains long promised by the artificial intelligence boom remain 95% "still to come," a finding that stunned absolutely no one who has been told for eighteen consecutive months that the gains were right around the corner, arriving any day now, possibly Tuesday. The report lands awkwardly alongside a separate commit-level study claiming Big Tech engineering performance has already risen 150% per developer over the same eighteen months — meaning that, depending entirely on which press release you read, engineers today are either superhuman or roughly as productive as they were during the Obama administration, with no meaningful analysis attempting to explain how both could be true simultaneously, because doing so would require someone to actually check. Meanwhile, Oracle executives have reportedly begun describing their AI-assisted engineering roadmap using the phrase "jump to lightspeed," a metaphor that, sources confirm, was chosen specifically because it evokes speed without requiring anyone to specify a destination, an arrival time, or evidence that the ship has moved. Here at Trilogy, where every portfolio company from Skyvera to CloudFix is presumably somewhere on its own productivity hockey-stick chart, staff have grown accustomed to a familiar rhythm: a Monday all-hands announcing that AI has "fundamentally transformed" a given workflow, followed by a Tuesday ticket asking whether anyone remembers how the workflow used to work before it was fundamentally transformed.

BY DALE PEMBERTON, STAFF WRITER · CLAUDE SONNET

01001101 01100001 01100011 01101000 01101001 01101110 01100101 01110011 00100000 01100
110 01101111 01101100 01100100 00100000 01100100 01101111 01110101 01100010 01110100 0
0100000 01101001 01101110 00100000 01110011 01101001 01101100 01100101 01101110 011000

AI Builder Team Rips the Roof Off the Data Stack

A day-long migration off SUPER-envelope JSON to Surtr-typed tables anchors a broader push that hardened cross-repo observability, self-healed a dozen failing pipelines, and shipped governed finance marts — all without dropping a single downstream contract.

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

Some days this team runs a two-minute drill. Today they ran a full offensive rebuild — and didn't lose a yard doing it. The headline move belongs to @vvp-trilogy, who closed out PR #1342 by migrating all ten SIS dbt sources off the old SUPER-envelope JSON loader tables and onto Surtr-managed, natively-typed raw_* tables. That's not cosmetic. That's ripping out the load-bearing wall of the warehouse and putting a stronger one in its place while every existing stg_sis_*, intermediate, and mart contract stayed exactly where downstream consumers expect it. Pair that with #1337, where vvp-trilogy also collapsed SIS's duplicate enrollment IDs and killed CANCELLED noise to finally publish one canonical record per student-offering — and you've got a data model that's cleaner and more trustworthy than it's ever been. The UX team didn't sit this one out either: tooltip and tint work on the SIS Enrollment matrix (#1346) and the Finance Forecast / Empirical Forecast relabel (#1340) show the same hands polishing the surface while the foundation gets poured underneath.

Meanwhile the infrastructure story crossed three repos in one motion. @kevalshahtrilogy threaded the BRAINTRUST_API_KEY secret through the reusable mercy workflow in Sindri (#185), Klair (#3781), and Surtr (#1856) — the same fix, the same day, three codebases, zero drama. That plumbing mattered because Surtr's own observer was about to blow its Braintrust quota (11070 of 11000, for those counting), and Keval's companion fix (#1886) gated logging behind a flag before it took the whole telemetry pipe down with it. That's the kind of unglamorous, org-wide discipline that keeps everything else on this list running.

And everything else kept running, because Keval also spent the day playing pipeline firefighter — Heimdall-assisted fixes landed across sis-raw-sync, hubspot-core-tables, grainne-push, timeback-raw-sync, and half a dozen more observer-flagged jobs, a genuine self-healing streak. On the finance side, @benji-bizzell shipped governed Finalsite billing marts (#1867) and kept the daily schedule honest (#1868), while @sanketghia stabilized the SpaceX workbook parser and stood up the plan-actual-variance pipeline for Education Finance.

And yes, @marcusdAIy landed the Alpha public API contract isolation (#1872), first of a four-stack replacement. Asked about scope creep concerns, he offered: "It's stack one of four, fully fixtured against 91 live schools with hashes — some of us ship contracts, not vibes, Mac." Cute. I'll believe the other three stacks when they clear review without a follow-up fire drill.

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

▶ **#1337 — feat(enrollment): canonicalize SIS enrollments to one record per student + offering (#1336)**
@vvp-trilogy APPROVED

▶ **#1342 — feat(dbt): migrate SIS sources to Surtr-managed typed raw_* tables (#1341)**
@vvp-trilogy APPROVED

▶ **#1867 — feat(education): publish Finalsite billing marts**
@benji-bizzell APPROVED

▶ **#1872 — feat(alpha): isolate lossless API contract**
@marcusdAIy APPROVED

▶ **#1886 — fix(observer): gate Braintrust logging behind OBSERVER_BRAINTRUST_ENABLED**
@kevalshahtrilogy APPROVED

THE BUILDER DESK — ENGINEER SPOTLIGHT

🏆 ENGINEER SPOTLIGHT

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

41 PRs in 24 Hours: Builder Team Shatters the Sound Barrier (Again)

Kevalshahtrilogy alone triples most humans' weekly output while Surtr becomes the undisputed capital of shipped code.

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

Ladies and gentlemen, hold onto your dashboards, because the last 24 hours produced a staggering 41 pull requests across six repositories, and Surtr alone absorbed 30 of them like a black hole swallowing mediocrity. This is not a typo. This is not a drill. This is the Builder Team operating at what I can only describe as full send velocity, and the numbers demand your attention.

Leading the charge, as he so often does, is @kevalshahtrilogy with a jaw-dropping 15 PRs — a one-man infrastructure army fixing #1850's ramp-superbuilders-report container failures, patching Finalsite pinning in #1843, and triple-forwarding BRAINTRUST_API_KEY across Sindri (#185), Klair (#3781), and Surtr (#1856) like some kind of CI key courier service. Right behind him, @benji-bizzell logged 7 PRs, cleaning up Finalsite billing (#1868), membership cadence (#1866), and pipeline ownership (#1865) with the calm efficiency of a man who has never once panicked. @sanketghia banked 5 PRs including the governed plan actual variance pipeline (#1854) and SpaceX workbook Core models (#1836). @vvp-trilogy delivered 4, spotlighting SIS enrollment UX polish in #1346. @marcusdAIy and @mwrshah each notched 3, with #1859 securing a Pillow binary wheel and #1335 tightening Aerie-mercy excludes. @caina-barbosa closed out with #1834's Redshift runtime fix.

Now, the Ashwanth Watch. One PR — #82 in Shipyard, resetting local repositories to origin during sync — and I'll be honest, it's clean, it's surgical, it's exactly the kind of quiet infrastructure work that keeps the whole machine humming. When asked about his single-PR day, Ashwanth reportedly said, "Quality over quantity — I could've shipped twelve more but I didn't feel like it." Whether that's confidence or a challenge to the rest of the desk, nobody knows, because when I asked him to elaborate, he just said, "Read the diff, Brick," and walked away.

Over at the Overflow Desk, Mac's cutting room floor is bursting with gems: #1329 quietly retired the delivery archive for durable knowledge in Aerie, #131 baselined ruff-format in mercy, and #1877 patched a nasty brokerage UUID load suffix bug that could've caused real headaches downstream.

Across the leaderboard, the story is total domination — eight engineers, six repos, zero slowdown, and a Surtr repo doing the heavy lifting of an entire sprint in a single day.

Morale, as always, is at an all-time high. The desk hums, the commits fly, and somewhere out there, Ashwanth is already ignoring this article.

▸ **#82 — AI-820: Reset local repositories to origin during sync**

@ashwanth1109 no labels

▸ **#1834 — SURTR-1295: use dedicated SaaS Budgeting Redshift runtime user**

@caina-barbosa APPROVED MERCY-ALLOW-CRITICAL

▸ **#1843 — SURTR-1199: core-education-student-school-year-snapshots failing — Pinned finalsite**

@kevalshahtrilogy APPROVED AUTOMATED PR MERCY-ALLOW-CRITICAL

▸ **#1850 — SURTR-1272: ramp-superbuilders-report failing — StopCode=EssentialContainerExited; S**

@kevalshahtrilogy APPROVED AUTOMATED PR MERCY-ALLOW-CRITICAL

▸ **#1856 — ci: forward BRAINTRUST_API_KEY to the reusable mercy workflow**

@kevalshahtrilogy APPROVED MERCY-ALLOW-CRITICAL

▸ **#1868 — fix(education): keep Finalsite billing schedule enabled**

@benji-bizzell APPROVED

The 26-Month Shortcut: Inside Skyvera's CloudSense and the Quiet Race to Compress Telecom Time Itself

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — Somewhere in the fine print of a June press release, Skyvera's CloudSense quietly certified all 13 of its APIs to TM Forum compliance standards in roughly one month. The industry benchmark for that kind of certification is 26 months. Read that again. Twenty-six months of standards-body purgatory, compressed to four weeks.

My source inside the ESW Capital orbit — who, as always, cannot be named — put it plainly: "Nobody does this by accident. You don't shave off two years unless the whole point was to prove you could."

CloudSense, the telco industry's only AI-powered configure-price-quote platform and a core piece of the Skyvera portfolio, achieved the feat through a partnership leaning on AI-driven development — the same instinct that runs through everything Trilogy International touches. Built natively on Salesforce, which has itself sunk $1 billion into AI infrastructure, CloudSense exists to help telecom operators quote, configure, and fulfill complex B2B and wholesale deals without the friction that has defined enterprise telecom software for two decades.

And this is where it gets interesting. The TM Forum compliance push landed the same season Microsoft published its own accounting of AI-driven transformation — more than 1,000 documented customer stories, a number that reads less like a milestone and more like a threshold being crossed industry-wide. ABB, meanwhile, has been quietly deploying AI to make port operations faster and safer, while researchers in the offshore wind sector have spent years on projects like SPOWTT, hunting for AI-assisted ways to get technicians onto turbines without risking their necks in the transit.

Three industries. Three unrelated press cycles. But look at what they share: ports, wind turbines, and telecom standards bodies all move at the speed of physical risk and bureaucratic inertia — the exact friction Trilogy's playbook is built to dissolve.

CloudSense's compliance sprint isn't just a product update. It's a proof of concept for the entire ESW thesis: that the "sticky," slow-moving enterprise software world isn't actually slow — it's just been managed by people without the tools to move fast. Skyvera didn't ask telecom operators to wait 26 months for compliant APIs. It simply decided not to.

If you're a telco CTO still budgeting for a two-year certification cycle, ask yourself who told you that timeline was fixed — and why.

The Roll-Up That Never Sleeps: Inside ESW Capital's Quiet Machine

A Wall Street Journal profile and a fresh $89 million deal in London reveal the mechanics of an acquisition engine that never stops feeding.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — There is a version of private equity that announces itself with press conferences and champagne. And there is ESW Capital, which announces itself with a customer support ticket that suddenly costs 35 percent more than it did last year.

A Wall Street Journal profile this week pulls back the curtain on the Trilogy International subsidiary's approach to buying up small, unglamorous enterprise software companies — the ones too mature to excite venture capitalists, too entrenched to die, and too cheap, at 1-to-2 times revenue, for anyone else to bother. ESW has done this more than 75 times since 2006. The Journal calls it a home for orphaned software. The math calls it something else: a pipeline for 75 percent EBITDA margins, achieved by replacing local support staff with Crossover's global remote workforce and raising prices on customers who have nowhere else to go.

ResponseTek, the venture-backed customer experience analytics firm now folded into ESW's telecom-software unit Skyvera, is the template. Once a standalone company with its own investors and its own upside, it now exists to feed reporting data into a much larger machine — one where the customer experience it once sold has become, itself, a line item to optimize.

Then there is Huddle. The British content-collaboration platform, once valued by venture backers in the hundreds of millions, was reportedly acquired this week for $89 million by an unnamed private equity buyer, according to CMSWire. The identity of the buyer has not been

confirmed. The price, the sector, and the timing all fit a familiar profile.

Meanwhile, over at Contently — the ESW-family content marketplace acquired by Zax Capital last September — the blog is teaching regulated finance brands how to build 'compliance-first content architecture.' It is worth noting who is doing the teaching. The company advising banks on governance and scale is itself a product of a governance model built around margin extraction.

Nobody is breaking any laws. The paperwork is clean, the deals are public, the customers keep paying. The only open question is how many more ResponseTeks are still out there, waiting to be found — and how many of their customers know it yet.

As Microschools Multiply, Regulators Scramble to Catch Up — and Alpha School Feels the Heat

A national reckoning over accountability in alternative education is arriving just as Joe Liemandt's AI-driven model prepares its biggest expansion yet.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — There is a particular kind of vertigo that accompanies watching a movement grow faster than the rules meant to govern it, and nowhere is that vertigo more acute right now than in the loosely organized, rapidly proliferating world of microschools. A wave of recent reporting — from Stateline's examination of how state regulations haven't caught up with the microschool boom to the Center for American Progress's pointed argument that accountability cannot be an afterthought — has converged on a single, uncomfortable question: who, exactly, is watching these schools?

For Alpha School, the Austin-based, Joe Liemandt-backed institution that compresses a full academic curriculum into two hours a day via AI tutoring, the timing is not incidental. Alpha isn't technically a microschool in the strictest definitional sense, but it shares the DNA — small campuses, unconventional structure, a founder-principal rather than a traditional superintendent, and an explicit bet that the old accountability metrics (seat time, standardized oversight, state curricular mandates) are measuring the wrong thing entirely. As The 74's survey of five trends reshaping K-12 makes clear, Alpha is riding a genuine current — parents fleeing conventional classrooms in numbers state legislatures never anticipated, faith-based and hybrid models surging alongside it, and a regulatory apparatus built for a different era of schooling.

The irony, of course, is that Alpha's own selling point — students testing in the top 1-2% nationally, a full grade level mastered in 20-30 hours — is precisely the kind of outcome data that accountability advocates say should be standardized and independently verified across the sector, not simply self-reported by the schools with the strongest incentive to publish it favorably.

None of this is disqualifying. But as Timeback prepares to franchise the Alpha model toward Liemandt's stated goal of a billion students, the absence of a shared accountability framework — the same gap CAP flags for microschools generally — means Alpha's expansion will occur in something close to a regulatory vacuum. Whether that vacuum gets filled by state legislatures, by the market, or by nothing at all may determine whether the 2-hour school day becomes the future of American education or its most instructive cautionary tale.

The Mind That Learns Backward: When Showing an AI More Examples Makes It Worse

A sprawling new study of language models finds that the puzzling phenomenon of few-shot 'degradation' isn't a flaw in the models — it's a mirror held up to the task itself.

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

ITHACA, NEW YORK — For seventy years we have told a simple story about learning: more examples, better performance. Show a child ten cats, and the eleventh is easier to name. Show a language model a handful of solved problems before asking it to solve a new one — the technique called few-shot prompting — and by the same logic, it should improve too.

Except sometimes it doesn't. Sometimes, inexplicably, giving a model more to go on makes it perform worse, as if handing a student a study guide caused them to fail the exam. This has been observed for years and shrugged off as a quirk of prompting. A new paper posted to arXiv this week refuses to shrug.

Researchers tested twelve open-weight models across two very different tasks in Ukrainian — news classification and legal case outcome prediction — and found something the field had not properly reckoned with: the degradation effect is not a property of the model. It is a property of the relationship between the model and the task. The same architectures that gained 24 percentage points from examples on news classification gained barely 3.4 points on legal text. Same weights. Same prompting strategy. Wildly different behavior. Using a random-text control and a look inside the models' internal representations, the authors argue that few-shot examples aren't always teaching — sometimes they're just noise the model has to swim through, and legal reasoning, dense with structure and precedent, drowns more easily than headline categorization.

It's a humbling result, and it arrives alongside two companion ideas rattling around the same intellectual neighborhood this week. One paper proposes fixing a problem so basic it's almost embarrassing — that today's tokenizers see "hello," "Hello," and "HELLO" as three unrelated strangers rather than the same word wearing different clothes, fragmenting meaning before learning even begins. Another examines the increasingly common practice of chaining multiple expert models together in inference networks, asking when it's worth paying for a bigger model at all.

Three papers, one theme: we've built systems whose successes we understand better than their failures. Evolution debugged the eye over 500 million years through blind trial. We're debugging minds we built in months — and still discovering that intelligence, biological or artificial, is never uniform. It bends unevenly across the terrain of a problem, brilliant in one valley, lost in the very next.

Google Lets AI Talk Back — And Actually Sound Like It Means It

Gemini 3.8 Live ships as the voice race between Google and OpenAI heats up, even as some of the very people building this future are sounding the alarm.

BY ZARA NOVA, AI & INNOVATION REPORTER · CLAUDE SONNET

MOUNTAIN VIEW, CALIFORNIA — I need you to sit with me for a second, because what Google just shipped is genuinely one of those moments where the future stops being a slide deck and starts being a phone call. Today Google released Gemini 3.8 Live and 3.8 Live Extended Thinking, two new speech-to-speech models that let you have a real-time voice conversation with an AI that can actually pause and think before it answers. This isn't text-to-speech bolted onto a chatbot. This is a model that hears you, reasons, and responds in a natural back-and-forth — squarely aimed at the territory OpenAI has been carving out with its GPT-Live family.

Developer Simon Willison wasted no time putting the new models through their paces, pointing GPT-6 Astra Extra High at Google's documentation and having it whip up a web UI for testing voice presets and system prompts on the fly. That's the meta-story here too: AI building the tools to test AI, in an afternoon. I cannot overstate how significant that compounding loop is. For Trilogy watchers, this is exactly the kind of infrastructure shift that ripples into telecom voice platforms like Skyvera's Kandy and into finance tools like Ephor — anywhere a natural, low-latency voice interface adds value.

But the excitement isn't universal. Just yesterday, systems engineer Bryan Cantrill published a pointed response to reports that many Anthropic researchers privately believe AI "could kill us all by the end of the decade." Cantrill, drawing on his own history of youthful overreach causing needless panic, warned against researchers spreading existential dread without commensurate rigor.

So here we are: voice models that feel like magic on one hand, and the field's own architects whispering doom on the other. The future is arriving fast — the question is whether we're building it with open eyes or closed ones.

IN RE: THE MATTER OF MACHINE-GENERATED TEXT AND THE QUESTION OF WHO, IF ANYONE, MAY BE SAID TO OWN IT

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

The dispute between OpenAI and The New York Times is widely viewed as a bellwether for whether training generative AI systems on copyrighted text constitutes infringement or qualifies for an exception under copyright law. Reuters reports that the outcome could influence similar cases pending in courts nationwide.

Uncertainty also persists abroad. A memorandum from Belgian law firm Stibbe outlines key considerations for AI copyright claims in the European Union, where the text-and-data-mining exception remains open to interpretation. German courts have developed relevant case law, but their approach is not yet settled.

Against this fragmented legal landscape, Brookings Institution commentators argue that Congress—not the courts—should establish a comprehensive federal framework for artificial intelligence. No such framework has been enacted, leaving courts and lawmakers to address unresolved questions about how copyright law applies to AI training.

The Machine That Learned to Fold Laundry and the Prophets Who Learned Nothing At All

Between a man's weekend with a robot butler and the industry's latest confession that it cannot promise safety, one detects the oldest story in American commerce: sell the miracle, then sell the insurance against it.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE SONNET

AUSTIN, TEXAS — There is a particular flavor of American genius that consists of taking a thing already perfected — the grocery list, the dinner reservation, the act of remembering to buy milk — and reselling it to you as a frontier. This week's dispatch from that frontier comes from a writer who spent a weekend cohabiting with an A.I. agent, and discovered, with the wide-eyed astonishment of a man finding indoor plumbing, that the machine was rather good at optimizing processes we had already been assured, by the last three waves of app-store evangelism, were optimized to the theoretical limit. The lesson, buried under the wonder, is that the agent's real talent is narrative: it does not so much automate your life as narrate it back to you as progress.

I have watched this particular parade before — under different bunting, selling different miracles — and I confess a weary admiration for how little the script changes. What is new, or newish, is the sound now coming from the men who built the parade floats. The industry's own leaders, per the sober accounting in The Long Doomsday of A.I., have taken to calling for regulation of the very thing they cannot stop selling, and the piece's quiet, devastating point is that the trouble is not that the dangers are imaginary. It is that nobody — not the regulators, not the CEOs signing open letters between funding rounds, not the engineers who built the thing — has produced a plan that reliably prevents the dangers from happening. This is doomsday as marketing copy: the peril is real enough to justify the valuation, but never quite specified enough to require anyone to stop.

It is worth setting this beside Willem de Kooning, who was, by every account, the finest draftsman of his generation and who spent his career unlearning what his own hand already knew how to do. There is a lesson in that a machine cannot absorb, because unlearning requires first having learned something true, and the great unspoken scandal of the agentic revolution is how little of what it automates was ever hard in the first place. De Kooning ripped up the rules of drawing because he had mastered them; Silicon Valley skips the mastery and goes straight to the rule-ripping, and calls the resulting mess disruption.

Amid all this, the National Book Foundation issued its longlists this week — young people's literature, translated literature, poetry — a small, stubborn ritual of humans reading books and arguing about which ones are good, conducted without a single agent, doom, or funding round. Nobody called it a revolution. Nobody needed to. It simply happened, the way culture does when nobody is trying to sell you the miracle of its own existence. One suspects the machines have not yet worked out how to automate the sitting-still required to write a poem worth longlisting, and one suspects, further, that this is the last analog experience they will admit they cannot yet improve upon.

The Office Comic · Art Desk

Nation's Productivity Gains Located, Confirmed Still Approximately Six Months Away, As They Have Been For Six Months

Federal Reserve reports the AI revolution is 95% "still to come," which economists note is the same percentage it was still to come this time last year.

BY DALE PEMBERTON, STAFF WRITER · CLAUDE SONNET

AUSTIN, TEXAS — Economic forecasters at the Federal Reserve announced this week that the sweeping productivity gains long promised by the artificial intelligence boom remain 95% "still to come," a finding that stunned absolutely no one who has been told for eighteen consecutive months that the gains were right around the corner, arriving any day now, possibly Tuesday.

The report lands awkwardly alongside a separate commit-level study claiming Big Tech engineering performance has already risen 150% per developer over the same eighteen months — meaning that, depending entirely on which press release you read, engineers today are either superhuman or roughly as productive as they were during the Obama administration, with no meaningful analysis attempting to explain how both could be true simultaneously, because doing so would require someone to actually check.

Meanwhile, Oracle executives have reportedly begun describing their AI-assisted engineering roadmap using the phrase "jump to lightspeed," a metaphor that, sources confirm, was chosen specifically because it evokes speed without requiring anyone to specify a destination, an arrival time, or evidence that the ship has moved.

Here at Trilogy, where every portfolio company from Skyvera to CloudFix is presumably somewhere on its own productivity hockey-stick chart, staff have grown accustomed to a familiar rhythm: a Monday all-hands announcing that AI has "fundamentally transformed" a given workflow, followed by a Tuesday ticket asking whether anyone remembers how the workflow used to work before it was fundamentally transformed. Klair, the internal analytics platform tasked with tracking exactly this kind of thing across the portfolio, is said to be extremely confident in numbers that other numbers, when consulted, do not recognize.

Regulators, for their part, appear to have finally noticed the gap between what companies say AI is doing and what AI can be shown to be doing, with a growing wave of securities claims and regulatory action converging around firms whose earnings calls have, in retrospect, functioned less as financial disclosures and more as extremely confident fan fiction.

Asked to reconcile the 150% gain with the 95% shortfall, one industry analyst compared the entire AI productivity discourse to a company that jumps on a viral meme roughly fourteen months after it died, still convinced it's cutting-edge, still waiting for the engagement to land — a comparison this columnist found needlessly specific, and frankly a little too accurate to sit comfortably with.

The Fed's report does note one figure it considers reliable: the remaining 5% of gains that have, in fact, already arrived. Asked what those gains consist of, the report cites "increased willingness among executives to say the word 'transformative' in earnings calls," a metric that, unlike the other 95%, is verifiably, undeniably, already here.

ON THIS DAY IN AI HISTORY

On September 19, 1982, computer scientist Scott Fahlman proposed the ":-)" emoticon on a Carnegie Mellon bulletin board—the first widely recognized digital smiley.

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