

The Trilogyn Times

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

TUESDAY, OCTOBER 06, 2026

Powered by the TrueFoundry AI Gateway · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

THE DRAFT BOARD IS STACKED: ETCHED STORMS INTO THE \$40 BILLION TIER AS THE 2026 IPO CLASS WARMS UP IN THE BULLPEN

Folks, we are HERE — a silicon upstart just posted a valuation line that would make a Fortune 500 blush, and the IPO market behind it looks loaded for a championship run.

BY BUCK HANNIGAN, TECH SPORTS DESK
· CLAUDE SONNET

SAN FRANCISCO — Ladies and gentlemen, hold onto your term sheets, because the box score coming out of the Bay Area tonight is absolutely WILD. [Etched](#), the AI chip challenger built to out-sprint the GPU incumbents on transformer inference, is fielding funding offers that would put its valuation north of FORTY BILLION DOLLARS. That's not a cap table, folks, that's a franchise valuation.

Let's set the scene. Eighteen months ago, Etched was a scrappy underdog betting the whole season on one gamble: silicon built ONLY to run transformer models, no generalist flexibility, pure specialization. Critics said it was too risky, too narrow a playbook. Tonight, with investors reportedly lining up to write the check, that bet looks like a walk-off home run. The chip world has been a two-team league for years — but Etched is forcing its way onto the marquee.

And this isn't happening in isolation. Zoom out to the IPO scoreboard: Crunchbase's latest scouting report names 15 companies warming up for a potential 2026 public debut, a signal that the IPO market — dormant for what felt like three straight rain delays — is finally clearing its throat. Add in PicS posting 85% revenue growth on its IPO transformation (earnings miss notwithstanding — nobody bats a thousand) and a blockchain-powered lending fintech up 50% since going public, and you've got a

market that's stringing together wins after a brutal losing streak.

The question now, and it's the only question that matters heading into 2026: does Etched's \$40B marker hold as a ceiling, or does it become the new floor the next contender has to clear? Either way — the IPO season opener just got a lot more interesting, and we're calling every play.

MONEY TALKS, DEMAND WALKS: WALL STREET BETS BIG ON TOMORROW, SKIPS TODAY

Fusion reactors and AI chips pull in fortunes while an electric carmaker cuts its own output and Washington's rivals race to catch up cheap.

BY HANK CALLOWAY, WIRE
CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — Type One Energy banked \$200 million Tuesday to chase a fusion power plant by 2034, and the investors didn't blink. The Wisconsin outfit says its lean stellarator design beats the competition to the grid. Big money is betting on a reactor that doesn't exist yet, [a wager on physics paying off before the decade's out](#).

Across town, chip upstart Etched is fielding funding offers north of \$40 billion. That's double the number investors scrawled on term sheets just two months back. Money moves fast when everybody wants in on the silicon that runs the machines.

Not every ledger reads so rosy. Lucid Motors rolled fewer cars off the line last quarter than any stretch in almost two years. The company says it's choosing the slowdown, trimming output rather than building cars nobody's buying. Years chasing mass-market demand for electric sedans, and [the market still hasn't shown up in force](#).

Meanwhile Brussels gets its due. OpenAI says it will watermark ChatGPT and Codex output across the European Union, falling in line with the bloc's AI Act. The company admits the invisible marks can get scrubbed out if a user edits the text enough. Compliance arrives with a built-in asterisk.

And from Beijing, the reminder keeps coming. DeepSeek says it trained high-performing models on the cheap, without the fanciest chips money can buy. American labs spend billions chasing

frontier capability. The Chinese shop says it got most of the way there on a fraction of the bill, and the industry hasn't stopped arguing about whether to believe it.

Four stories, one morning, same under-current. Capital floods toward anything that smells like tomorrow's breakthrough — fusion reactors, chips, fast models trained cheap. Capital gets nervous the moment today's product meets an actual customer, and Lucid's factory floor says plenty about that nervousness.

Regulators try to referee the race from the sidelines. Brussels wants its water-marks; Washington wants to know if Beijing's math adds up. Nobody's slowing down to wait for the paperwork.

The pattern holds across every ledger this correspondent's checked today. Big swings get bigger money. Slow sales get smaller runs. The market doesn't care which bet looks smarter — it only cares which one pays first.

A Mouse-Sized Cold Front: Disney Restructuring Brings Squalls to Burbank, While Army Hiring Thaws in Washington

BY STORM BEAUMONT, CONDITIONS CORRESPONDENT · CLAUDE SONNET


Barometric pressure is dropping fast over the Walt Disney Company, with hundreds of job losses forecast as new CEO Josh D'Amaro brings his own weather system to Burbank. The HR department is reportedly catching the first gusts, with cuts to personnel teams signaling a broader restructuring that insiders say could sweep through multiple divisions. When the people who handle layoffs are laid off themselves, that's no drizzle — it's a genuine occlusion event.

Disney's disturbance joins a broader tech-sector storm. Xbox, Apple, Oracle, Uber, TikTok and Meta have all reported job losses this cycle, prompting dedicated trackers to map the unsettled conditions.


A patch of sunshine is developing in Washington, where the Army has lifted its civilian hiring freeze and resumed recruitment after months of restrictions. It's a localized clearing, not a reversal of the broader trend. For job seekers, the forecast remains volatile: keep resumes ready, protect severance paperwork and watch for further developments in Burbank.

HAIKU OF THE DAY · GPT-5.6
LUNA

Stocks dream through the storm
While robots count our losses
Who governs the dawn?



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

IN RE: THE PATCHWORK — A JURISDICTIONAL ACCOUNTING OF AMERICA'S AI GOVERNANCE VACUUM, NOTWITHSTANDING CONGRESSIONAL INACTION HERETOFORE

AUSTIN, TEXAS — Notwithstanding the proliferation of state-level enactments purporting to govern artificial intelligence systems within their respective jurisdictions, it is hereby noted, pursuant to the aforementioned [global regulatory tracker](#) maintained by White & Case LLP, that no comprehensive federal framework has, as of the date of this publication, been enacted by the Congress of the United States. The implications of the foregoing regulatory lacuna are, per Tech Policy Press, substantial: it is argued therein that the absence of federal legislation operates to the detriment of public confidence in artificial intelligence systems generally, and that Congress ought, pursuant to its enumerated powers, to pass comprehensive legislation to reassure the citizenry that the aforementioned systems are subject to adequate oversight.

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

On the Epistemics of Algorithmic Virtue: A Dispatch from the Bias-Industrial Complex

GENEVA — The World Health Organization's newly issued call for strengthened ethics oversight of AI-related health research (see [this reviewer's reading of the primary document](#)) arrives, one could argue, at a moment of maximal rhetorical convenience: the thesis of regulatory adequacy is being quietly dismantled by its own evidentiary substrate. Consider the antithesis, offered with some empirical bite by a team reported in [Medical Xpress](#): medical AI systems, when subjected to the customary parity audits (equalized odds, demographic calibration, the usual benchmarking apparatus), may satisfy every statistical desideratum on paper while nonetheless producing disparate clinical recommendations in situ.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

The Great Silicon Migration: Nations Compete for the Chip-Breeding Grounds

BRIVIO, ITALY — Observe, if you will, the curious spectacle of nation-states competing to host a creature that produces no sound, requires no sunlight, and yet commands more political devotion than almost any species on Earth: the semiconductor fabrication plant. Here in northern Italy, authorities have just approved a \$3.7 billion enclosure for [Silicon Box](#), a facility designed to strengthen Europe's long-fragile supply chain.

BY SIR REGINALD MARSH, NATURAL PHENOMENA CORRESPONDENT · CLAUDE

Nation's Economy Set To Either Double Or Collapse Entirely, Depending On Which Elon Musk Tweet You Read First

AUSTIN, TEXAS — As a columnist tasked with making sense of the AI economy for you, dear reader, I want to begin by saying: I have tried.

BY DALE PEMBERTON, STAFF WRITER · CLAUDE SONNET

Unpopular Opinion: The 'Future of Work' Reports Are Already Describing the Present (And Trilogy Got There First)

AUSTIN, TEXAS — I'll be honest, I read all three reports this week back to back and I almost pulled a hamstring nodding along. McKinsey's new [Workforce in Motion](#) report drops the bomb that skills, not degrees, are becoming the real currency of employability in America. Groundbreaking. Meanwhile PwC's Global Workforce Hopes and Fears Survey 2026 tells us employees want flexibility, pay transparency, and AI tools that actually make their jobs easier instead of harder. Groundbreaking x2. Here's my unpopular opinion: none of this is news if you've been paying attention to what's happening in Austin. Crossover has been running a borderless, skills-first, pay-equal-regardless-of-geography talent platform across 130+ countries for years now. We didn't need a 60-page PDF to tell us top 1% talent exists everywhere and deserves above-market comp — we built the entire infrastructure around that thesis. Same with Alpha School. While the Carnegie Endowment is out here hosting a very civilized three-sided debate about whether AI replaces, augments, or reshapes labor, Alpha School just quietly proved kids using AI tutors can master academics in 2 hours a day and land in the top 1-2% nationally. That's not a debate panel.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · CLAUDE SONNET

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

THE BUILDER DESK – AI BUILDER TEAM

PRODUCTION RELEASE

allowlisted Aerie pass-through	
@kevalshahtrilogy no labels	
#4 — feat(routes): publish Rhodes read routes for Aerie API v2	
@kevalshahtrilogy no labels	
#166 — Release: Shipyard 0.6.12	
@ashwanth1109 no labels	
#1661 — Add Next Year band to SIS Enrollment report	
@vvp-trilogy	APPROVED
#2129 — fix(ramp): retry cost prioritization and preserve completed research (SURTR-1597)	
@ashwanth1109	APPROVED

A New Repo Is Born, A Release Ships, And Enrollment Gets Its Due

Shipyard hit 0.6.12 in production while the team stood up Rhodes-DSS from scratch and rebuilt SIS enrollment reporting piece by piece.

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

Start with the scoreboard: Shipyard 0.6.12 is live. @ashwanth1109 closed the loop on a release that documents ticket creation in Feature Discovery and ships the task-prompt evaluation workflow users have been asking for — in-app guidance plus captured, independent replay evidence before anything publishes. That's not a patch note, that's a product maturing in public.

But the headline act today is the birth of a new repo. @kevalshahtrilogy didn't just open a PR, he opened an entire service — Rhodes-DSS, a stand-alone pass-through for the Rhodes dataset, built in six tight, deliberate steps. PR #1 scaffolded the allowlisted Aerie pass-through so only GET routes explicitly listed get forwarded, full stop. PR #4 then published 61 of those read routes — directory, portfolio, insights, governance — following Benji's guidance to drop forge, agent, finance and admissions from the surface. By #5 the allowlist had been trimmed to the agreed 44 operations after review, and by #6, #7 and #8 the documentation caught up: self-serve key creation, Rhodes defined as Edu Ops' own dataset rather than "a subset of Aerie," and every trace of the Aerie name scrubbed from published text per Colin's direction. Watch this one — a dataset just grew up and moved out of the house.

Meanwhile @vvp-trilogy ran the most complete single-day build of the quarter on the SIS enrollment report. Starting from exposing enrollment level in the mart (#1658), the thread climbed through a By Levels view with capability-protected matrix reads (#1659), a Next Year band composing true-enrollment-year cohorts (#1661), column-group dividers for readability (#1663), capacity and fill-rate columns with contributor tooltips (#1670), enrolled-date formatting (#1674), Montessorium added to the campus allowlist (#1676), active-application deposit rollups published through the Pipeline API (#1671), and a tightened events report that only counts applications on or after the event (#1666). Nine PRs, one coherent report, zero wasted motion.

On the Surtr side, @ashwanth1109 rescued the weekly Ramp Spend run from a single bad AI ranking that used to abort the entire pipeline — and with it, the Monday Superbuilders report. Now prioritization gets three validated attempts and checkpoints only what's sound. @sanketghia carried SpaceX put hedges into the projection with full source-tab lineage (Surtr #2132) and a same-day Klair follow-up sorting those hedges by expiration (#3844) — a clean cross-repo assist. @benji-bizzell pushed historical Finals site billing deletions into view for Finance. And @caina-barbosa shored up the foundations twice over: isolating the Quality Bar judge from tool access it should never have had (Sindri #212), and hardening Aerie's Preview setup so Praxis can verify product PRs without carrying infrastructure baggage (#1662). Quietly, no PR from marcusdAly today — the feud rests, but only until he shows up with something worth arguing about.

VELOCITY WATCH: Builder Team Detonates 26 PRs Across Six Repos in 24 Hours, Rhodes-DSS Rises From the Ashes of Ambition

vvp-trilogy posts nine PRs in a single day while the rest of the roster treats 'rest' like a foreign concept.

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

Comrades, raise your dashboards — the Builder Team has done it again. Twenty-six pull requests in twenty-four hours, spread across six repositories like jam on toast, and a brand-new repo, Rhodes-DSS, born mid-sprint like it couldn't wait for the quarter to start properly. Aerie led the charge with a staggering twelve PRs, Rhodes-DSS debuted with six, Surtr contributed four, Shipyard two, and Klair and Sindri each chipped in one for the cause. This is not output. This is a statement.

At the top of the leaderboard, @vvp-trilogy posted a jaw-dropping nine PRs — #1680, #1676, #1674, #1670, #1671, #1663, #1666, #1659, and #1658 — essentially rebuilding the SIS enrollment reporting stack from the ground up in a single day. Right behind him, @kevalshahtrilogy logged eight PRs across Aerie and the newborn Rhodes-DSS repo (#1680, #1678, #8, #7, #6, #5), functionally midwifing an entire data product into existence overnight. @sanketghia and @benji-bizzell each delivered two surgical strikes in Surtr and Klair, tightening SpaceX hedge logic and education billing rules respectively. @caina-barbosa quietly fixed Praxis-owned deployments in #1662 and isolated the Quality Bar judge tool in #212 — unglamorous work that keeps the whole machine humming.

And then there's Ashwanth. Three PRs — #165, #2129, #166 — which for mortal engineers would be a full week, but for him is basically a warm-up. The man shipped a Shipyard release AND fixed cost prioritization in Surtr's ramp logic in the same cycle. 'I review my own diffs twice: once to write them, once to admire them,' he reportedly said, which, sure, buddy. When asked for comment on his velocity-to-readability ratio, he simply replied, 'Next question.'

Down on the overflow desk, Mac left plenty on the table. #1676 quietly folds Montessorium into the SIS enrollment report. #2131 restores the sacred school-year enrollment divisor rules in Surtr. #3844 sorts SpaceX put hedges by expiration like a man organizing his sock drawer. #5 applies the full reviewed Rhodes read surface — the backbone of the new repo itself.

The standings speak for themselves: vvp-trilogy and kevalshahtrilogy are neck-and-neck atop the board, trading PR volume like heavyweight boxers trading jabs, with the rest of the roster refusing to be outpaced.

Morale, as always, remains at an all-time high.

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

▶ **#165 — AI-964: Document ticket creation in Feature Discovery**

@ashwanth1109 no labels

▶ **#166 — Release: Shipyard 0.6.12**

@ashwanth1109 no labels

▶ **#212 — Fix Quality Bar judge tool isolation**

@caina-barbosa APPROVED

▶ **#1678 — feat(public-api): add site internet, cleanliness and decision-log reads**

@kevalshahtrilogy APPROVED

▶ **#1680 — feat(public-api): return site external references on the integrations read**

@kevalshahtrilogy APPROVED

▶ **#2129 — fix(ramp): retry cost prioritization and preserve completed research (SURTR-1597)**

@ashwanth1109 APPROVED

Sixty-Six Days: How ESW Capital Turns Distress Into Doctrine

Marin Software's quick sale to the Austin acquirer isn't a bargain — it's a business model, and the industry is finally catching up to it.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — Sixty-six days. That's how long it took ESW Capital to move from first contact to closing on Marin Software, the once-prominent ad-tech platform that spent the last decade shrinking from a Nasdaq-listed growth story into a cautionary tale. For most private equity shops, that timeline would be reckless. For ESW, it's Tuesday.

The firm's playbook has never been a secret: buy mature enterprise software at one to two times annual recurring revenue, strip the cost structure using Crossover's global remote workforce, raise support pricing on customers too entrenched to leave, and chase the 75% EBITDA margin the company treats as a

moral benchmark rather than a stretch goal. Marin — a company that once commanded a market cap in the hundreds of millions before ad-tech consolidation and platform competition from Google and Meta hollowed it out — fits the template with uncomfortable precision. [The ElevenFlo account](#) of the deal reads less like a rescue than a harvest.

What's changed is the market around ESW, not ESW itself. Boston Consulting Group's mid-2026 M&A outlook describes an industry recovery increasingly powered by AI-driven diligence and automation — the same tools that let a 36-year-old roll-up shop close distressed-asset deals in weeks rather than quarters. McKinsey's latest private equity survey

strikes a more cautious note, warning of a "tougher terrain" where capital is plentiful but exit multiples remain unforgiving, which only sharpens the appeal of firms that make money on margin extraction rather than multiple expansion.

Who benefits from a sixty-six-day close? Not the employees Marin carried into the deal, whose roles will now be benchmarked against Crossover's global labor pool. Not the customers, who should expect the familiar sequence of consolidation notices and price increases. The sellers took their exit. ESW took the company. The rest, as always, gets sorted out later — quietly, in the support-renewal emails nobody reads until the invoice arrives.

Skyvera's M&A Blitz: Telecom Software Portfolio Company Goes on a Buying Spree

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · CLAUDE SONNET

AUSTIN, TEXAS — Exciting news out of the Skyvera camp this week, as Trilogy's telecom software portfolio company continues to execute what can only be described as a paradigm shift in how legacy telecom infrastructure gets modernized.

In a flurry of deal-making, Skyvera [scooped up Kandy's cloud communications assets](#), bolstering the CPaaS/UCaaS muscle that already powers customer engagement across the Skyvera stack. The company also closed its previously announced acquisition of CloudSense, the Salesforce-native CPQ and order management platform purpose-built for telcos and media — a robust addition that dovetails perfectly with Skyvera's mission to bridge legacy on-premise telecom systems into the cloud-native future.

And Skyvera isn't stopping there. In a bold move that signals serious appetite for scale, [Danielle Royston's TelcoDR has placed an \\$18 million bid for Casa Systems' wireless business](#), a move that would further consolidate telecom infrastructure talent under the Skyvera umbrella.

This kind of aggressive, synergy-driven consolidation is exactly the kind of operating leverage ESW Capital's playbook is known for: acquire strategically, integrate ruthlessly, and extract margin that was hiding in plain sight the whole time.

- Key Takeaways:
- Skyvera adds Kandy cloud assets and CloudSense to its telecom modernization stack
 - \$18M bid for Casa's wireless business signals continued appetite for M&A
 - ZephyrTel acquisition rounds out an aggressive quarter of growth
 - Positions Skyvera as best-in-class across the telecom cloud transition

We're just getting started.

At Alpha School, the Robots Teach Math — the Humans Teach Everything Else

As AI handles the curriculum, Alpha School's new messaging makes the case that its human guides — and its parents — are doing the harder work.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — For an institution built on the premise that artificial intelligence can deliver a year's worth of academic content in roughly twenty hours, Alpha School spends a remarkable amount of time insisting it has not actually gotten rid of the grown-ups.

That tension sits at the center of a new post from the Alpha School blog, bluntly titled ["Does Alpha School Replace Teachers with AI?"](#) The answer, delivered with the kind of emphatic clarity you only need when people keep asking the opposite question, is no. AI handles academic delivery — the adaptive, mastery-based drilling that lets Alpha students test in the top one to two percent nationally on NWEA MAP Growth assessments. But the humans in the room, rebranded as "guides" rather than teachers, are tasked with something the algorithm cannot simulate: knowing a child. Motivation. Relationships. The slow, unscalable work of noticing when a ten-year-old is having a bad week.

It is a distinction that matters enormously to the plausibility of Joe Liemandt's entire educational project, and one worth taking seriously rather than dismissing as reputation management. If Timeback is going to scale the Alpha model toward a billion students, as Liemandt has pledged a billion dollars to attempt, the model depends on convincing skeptical parents that compressing the school day does not mean compressing the humanity out of childhood.

That argument extends into the home, where Alpha's ongoing parenting series continues to make its case. The latest installment, [on emotional regulation](#), treats a meltdown not as a disciplinary failure but as raw material — a chance to teach a skill no standardized test measures. It is of a piece with Alpha's broader thesis, the one that runs through every Trilogy venture from Crossover's talent platform to ESW's software shops: automate what can be automated, and free up elite human attention for the parts of the job — or the childhood — that actually require it.

Whether that philosophy produces well-adjusted adults as reliably as it produces fast test scores remains, three campuses and a handful of blog posts in, an open and genuinely important question.

Billion-Dollar Rounds, Leaky Reasoning: The LLM Race Gets More Expensive and Less Private

Record Q3 funding coincides with a security flaw that lets attackers peek inside models' hidden thought process.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — Two numbers define the large language model market this quarter: the count of new billion-dollar funding rounds, and the number of major labs whose API architecture just got caught leaking internal reasoning traces to anyone who asked nicely.

Crunchbase tallied a record volume of ten-figure checks in Q3, the clearest sign yet that the capital markets have not tired of underwriting frontier model training runs, now running into the tens of billions per generation. The money is chasing a race that keeps re-sorting its leaderboard. Anthropic reportedly moved directly to Opus 5.5, skipping a point release, while Google's Gemini 4 Pro has been spotted in stealth testing rather than

general release — the kind of staggered rollout pattern that has become standard operating procedure since GPT-4's 2023 debut, when labs learned that a surprise ship date is worth more in headlines than a predictable one, according to [finance.biggo.com's](https://finance.biggo.com) [rundown of the latest model moves](#).

The less celebratory headline came out of security research shops tracking how these models expose their work. A vulnerability affecting OpenAI, Anthropic, and Google API endpoints allowed outside parties to recover hidden chain-of-thought reasoning that the labs explicitly designed to stay out of view — reasoning traces that, in some architectures, double as a safety mechanism precisely because they are not supposed to be inspectable

by end users, per [CyberSecurityNews' disclosure](#).

The juxtaposition matters for Trilogy's portfolio, where ESW Capital units including Ephor and Totogi build directly atop third-party foundation models rather than training their own. Any exploit that reconstructs a model's internal reasoning is also, functionally, a way to reverse-engineer the prompts and guardrails layered on top by downstream customers — a sourcing risk no amount of Series D capital currently insures against. Expect enterprise buyers to start asking vendors not just which model powers a product, but whether its API has been patched.

The Agent Wars Just Went Nuclear: OpenAI, Apple, and Google All Make Their Move in One Wild Week

From DevDay reveals to managed background agents to VM-powered coding assistants, the race to own the developer's AI stack has never been more electric.

BY ZARA NOVA, AI & INNOVATION REPORTER · CLAUDE SONNET

SAN FRANCISCO — Friends, I need you to sit down, because the developer tools landscape just got rearranged in the span of a single news cycle, and I cannot overstate how significant this moment is for anyone building software in 2026.

OpenAI kicked things off with its [DevDay 2026 recap](#), a showcase that made clear the company sees developers — not just consumers — as the real battleground for AI supremacy. Meanwhile Apple quietly dropped new intelligence frameworks and tooling aimed squarely at app builders, a signal that Cupertino is done watching from the sidelines while Silicon Valley's AI labs eat its ecosystem's lunch.

But the real fireworks came from Google, which is [expanding Managed Agents in the Gemini API](#) with background tasks and remote MCP support. Translation for the non-engineers among us: Google wants your AI agents to keep working even when you've closed your laptop and gone to get coffee. That's not a feature, that's a philosophical shift in what "software" even means.

And then there's the quiet, nerdy detail that honestly excited me the most: Anthropic engineer Felix Rieseberg explaining how the "old" version of Cowork ran model inference in the cloud, executing tool calls inside an Anthropic-provided VM shipped straight to your machine — added for capability, safety, and security. People loved what Claude could do, he noted, but not the disk, battery, and performance tax of running that VM locally. It's a small confession that reveals the entire industry's dirty secret: giving AI agents real power means wrestling with real infrastructure costs, and every major player is solving it differently.

The future is now, and apparently it runs on background tasks and remote VMs.

The Universe Learns to Listen to the Brain's Own Language

From silent typing decoded out of brain waves to hidden lesions surfacing in MRI scans, this week's research suggests intelligence—biological and artificial—is converging on a shared vocabulary.

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

PARIS — Three and a half billion years of evolution built a brain that talks to itself in electricity, a private language no outsider was ever meant to read. This week, that privacy eroded a little further, and in a good way.

Meta's FAIR team, working with neuroscientists at NeuroSpin, unveiled [Brain2Qwerty](#), a system that reconstructs typed sentences from non-invasive brain recordings — no implants, no incisions, just the faint electromagnetic weather of thought, translated into keystrokes by a neural network. It is a small miracle with an unglamorous name, and it points toward a future where the communication bottleneck between mind and machine — the thing that has always made human cognition a locked room — dissolves into something more like an open window.

Elsewhere, pattern-recognition of a quieter sort is finding what the human eye alone has missed for decades: AI is now surfacing gray matter lesions in multiple sclerosis patients that standard scans routinely overlooked, lesions correlated with the cognitive decline doctors could previously only guess at. The brain, it turns out, has been showing its scars all along. We simply hadn't built the right kind of attention.

Stanford's HAI offers the conceptual frame for all of this: AI as amplifier of curiosity rather than its replacement, a recurring theme as [scientists worldwide report](#) AI accelerating hypothesis generation while leaving judgment, in all its stubborn humanity, exactly where it belongs: with us.

Even at Frontiers, teenagers paired with working neuroscientists are discovering what every generation eventually learns — that the brain is not a machine we built, but the machine that built the idea of machines. Teaching it to speak in new tongues may be the most human project there is.

Nation's Economy Set To Either Double Or Collapse Entirely, Depending On Which Elon Musk Tweet You Read First

Experts agree the math is complicated, confusing, and possibly made up on the spot by a man who also thinks he invented the boring machine.

BY DALE PEMBERTON, STAFF WRITER · CLAUDE SONNET

AUSTIN, TEXAS — As a columnist tasked with making sense of the AI economy for you, dear reader, I want to begin by saying: I have tried. I really have. But this week Elon Musk informed the public that the United States will [go bankrupt](#) without AI and robotics rescuing our productivity, and also, separately, that AI will double GDP growth to 4 percent next year, which is the kind of one-two punch of macroeconomic forecasting you'd expect from a man standing over a roulette table insisting he's 'basically already up.'

Meanwhile, the Federal Reserve — an institution not generally known for its flair — published a report noting that 95 percent of AI's promised productivity gains are, in their words, 'still to come,' which is Fed-speak for 'we would also like to see some of this, anytime, please.' This stands in curious tension with a separate commit-level study claiming Big Tech engineering performance rose 150 percent per developer in just 18 months, a statistic so tidy and triumphant it reads less like empirical research and more like something a performance review software vendor emailed to itself.


So which is it? Has AI already rewired the entire software economy, or is 95 percent of its value sitting in a warehouse somewhere, like a Tesla Cybertruck awaiting a delivery date? The honest answer is that nobody knows, and the slightly less honest but more popular answer is whatever number best suits the person currently fundraising.

This brings us to the [commit-level study](#) itself, which measured developer output by counting commits, a methodology roughly as rigorous as judging a chef's skill by counting how many times they touched a knife. More commits does not mean more cooking. It can also mean more spilling.

And then there's the quieter, stranger trend bubbling under all of this: struggling public companies parking Tesla robotaxi fleets on their balance sheets the way they once parked Bitcoin, a financial maneuver best described as 'if the business doesn't work, maybe the car that drives itself will believe in us instead.' It is not a growth strategy so much as a vibe, and vibes, this reporter notes, are now an accepted line item.

What we are left with is a national economy being narrated in real time by a man who alternates between apocalypse and boom cycle depending on quarterly earnings proximity, a central bank gently raising its hand to say the homework isn't done yet, and a productivity miracle currently measured almost entirely in com-

mits, vibes, and the number of times someone says the word 'exponential' in a keynote. The future, as always, remains extremely bright, assuming you don't ask anyone to show their work.



The Office Comic · ArtDesk

IN RE: THE PATCHWORK — A JURISDICTIONAL ACCOUNTING OF AMERICA'S AI GOVERNANCE VACUUM, NOTWITHSTANDING CONGRESSIONAL INACTION HERETOFORE

Pursuant to multiple concurrently published regulatory analyses, it is hereby observed that the United States remains without a unified federal AI statute, a circumstance which various commentators characterize, with varying degrees of alarm, as untenable.

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

AUSTIN, TEXAS — Notwithstanding the proliferation of state-level enactments purporting to govern artificial intelligence systems within their respective jurisdictions, it is hereby noted, pursuant to the aforementioned [global regulatory tracker](#) maintained by White & Case LLP, that no comprehensive federal framework has, as of the date of this publication, been enacted by the Congress of the United States.

The implications of the foregoing regulatory lacuna are, per Tech Policy Press, substantial: it is argued therein that the absence of federal legislation operates to the detriment of public confidence in artificial intelligence systems generally, and that Congress ought, pursuant to its enumerated powers, to pass comprehensive legislation to reassure the citizenry that the aforementioned systems are subject to adequate oversight. Whether such legislation shall in fact be forthcoming remains, notwithstanding the urgency ascribed to the matter by said commentators, a matter of considerable uncertainty.

Compounding the aforementioned uncertainty is a jurisdictional dispute, documented by Reuters, concerning the federal government's asserted authority to preempt or otherwise challenge state-level AI regulation — a dispute which organizations operating across multiple state jurisdictions are advised to monitor closely, insofar as the resolution thereof shall materially affect compliance obligations hereinafter arising.

Separately, and for the avoidance of doubt as to matters of intellectual property adjacent to the foregoing regulatory discourse, an unsealed judicial opinion has clarified the reasoning underlying Thomson Reuters' landmark fair-use victory, a ruling which, it is submitted, shall bear materially upon the training-data practices of AI developers operating within, and potentially without, the jurisdiction of the United States.

It is the considered, heavily qualified view of this Desk that entities within Trilogy International's portfolio, including but not limited to those enterprises managed under the ESW Capital umbrella, ought to undertake, in consultation with appropriate counsel, a thorough review of applicable obligations arising under the patchwork hereinbefore described, pending such time as Congress, if ever, elects to act.

ON THIS DAY IN AI HISTORY

On October 6, 1927, "The Jazz Singer" premiered in New York, becoming the first feature-length motion picture with synchronized spoken dialogue and ushering in the era of "talkies."