

The Trilogy Times

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

TUESDAY, APRIL 28, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

Google's \$40 Billion Anthropic Bet Reshapes AI Landscape as Model Wars Intensify

Tech giant's historic investment dwarfs previous AI deals while OpenAI and Anthropic spar over benchmark supremacy and model security.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — Google has committed \$40 billion to Anthropic in what industry analysts are calling the largest AI investment in history, fundamentally altering the competitive dynamics among frontier model developers.

The deal, which values Anthropic at approximately \$160 billion post-money, represents a 15-fold increase over Google's previous \$2.7 billion investment disclosed in 2023. The commitment gives Google significant influence over the Claude developer while stopping short of outright acquisition, sidestepping potential antitrust scrutiny that has plagued other Big Tech AI deals.

The announcement comes as [OpenAI's newly released GPT-5.5 narrowly edged](#)

[out Anthropic's Claude Mythos Preview](#) on Terminal-Bench 2.0, a coding benchmark that has become the de facto standard for measuring practical model capability. GPT-5.5 scored 87.3% versus Claude's 86.8%, a margin industry observers describe as statistically insignificant but symbolically important.

The competitive intensity has prompted an unusual alliance: OpenAI, Google, and Anthropic jointly announced new protocols to combat AI model theft, including cryptographic watermarking and shared threat intelligence. The collaboration suggests growing concern over state-sponsored actors and competitors using distillation techniques to replicate frontier capabilities at fraction of development cost.

Meanwhile, grassroots opposition to AI deployment is spreading beyond coastal tech hubs. From Indiana manufacturing towns to Idaho farming communities, a coalition of workers, artists, and small business owners is organizing against what they characterize as extractive AI economics. The movement has successfully lobbied for AI impact assessments in 14 state legislatures, creating potential regulatory fragmentation that could complicate nationwide AI rollouts.

Google's massive Anthropic investment appears designed to secure model supply amid this uncertainty, ensuring access to cutting-edge AI regardless of regulatory headwinds.

Meta's AI Storm Front Brings Layoffs and a Hiring Deep Freeze

As billions pour into models and infrastructure, the labor forecast turns sharply colder across Big Tech—and CEOs may be misreading the pressure system.

BY STORM BEAUMONT, CONDITIONS CORRESPONDENT · GPT-5.2

MENLO PARK, CALIFORNIA — The skies over Big Tech darkened today as Meta signaled a fresh squall line of cost-cutting: reports indicate roughly 8,000 job cuts alongside a freeze on about 6,000 open roles, all while the company pivots harder into AI spending. In weather terms, it's the classic pattern of a high-pressure capital system building over compute, pushing a cold front straight through headcount.

According to coverage circulating via [MSN's report](#), the cuts land as Meta intensifies its AI shift—an increasingly common forecast across the sector: fewer people on the ground, more horsepower in the data center. Another version of the same storm track pegs the move closer to 10% of the workforce with hiring frozen as AI investments surge, per [BW People](#).

Zoom out, and you can feel the broader atmospheric shift: Fortune warns that 66% of CEOs are freezing hiring while betting billions on AI—calling it a costly miscalculation. Translation for workers and founders: the barometric pressure is rising on the teams that remain, and “do more with less” gusts are strengthening.

CIO.com adds a practical advisory: tech leaders can't treat freezes like sudden hail—someone must own the policy, the exceptions, and the morale damage. If leadership leaves the freeze unmanaged, expect patchy turbulence: shadow hiring, delayed roadmaps, and burnout thunderstorms.

For startups, NEA partner Tiffany Luck's guidance on vertical AI reads like a shelter map: build moats where domain data, workflow integration, and distribu-

tion create defensible microclimates that platform giants can't easily replicate.

Prepare accordingly: workers should expect continued volatility through the next earnings season; founders should conserve runway and sell concrete ROI; and anyone counting on “back to normal” hiring should pack for a longer winter.

CHINESE UPSTART DEEPSEEK CRACKS AI RACE WIDE OPEN — ON BARGAIN CHIPS AND A SHOESTRING BUDGET

Silicon Valley engineers are calling a made-in-China model 'amazing and impressive' — and Washington's chip blockade didn't stop it.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

SAN FRANCISCO — A Chinese artificial-intelligence outfit called DeepSeek has thrown a wrench into every assumption Wall Street and Washington held dear about the AI arms race, training high-performing models on the cheap without the top-shelf chips America spent two years trying to keep out of Beijing's hands. The result has Valley engineers reaching for superlatives and investors reaching for antacids. The stock market felt it Monday.

Here is the plain score. DeepSeek built models that [rival the best American labs have produced](#), and it did so at a fraction of the cost. The company sidestepped Washington's export controls on advanced Nvidia chips — the very controls designed to kneecap Chinese AI development — by engineering around inferior hardware. That is not supposed to happen, according to every briefing paper on Capitol Hill.

The implications land like a brick through a plate-glass window. American AI giants — OpenAI, Google, Anthropic, the whole crowd — have operated on the premise that dominance belongs to whoever burns the most cash on the most powerful silicon. Billions upon billions of dollars have been poured into data centers on that bet. DeepSeek says the bet is wrong, or at least incomplete.

[The reaction in Silicon Valley has been swift and unguarded](#). Engineers and researchers are calling the Chinese models “amazing and impressive,” which is not a phrase this reporter hears often from


people whose stock options just took a bath. The technical community respects results, and the results speak plain English — or Mandarin, as the case may be.

Washington finds itself in an awkward spot. The chip export controls were the centerpiece of America's strategy to maintain AI superiority. DeepSeek's performance suggests that strategy has a hole in it wide enough to drive a supply truck through. Restricting hardware does not, it turns out, automatically restrict ingenuity.


The timing adds salt. Google reportedly just signed a classified deal allowing the Pentagon to use its AI models for any lawful government purpose — a move that drew immediate fire from its own employees. Reid Hoffman, the LinkedIn co-founder, announced a \$24.6 million AI startup aimed at cancer research. The American AI establishment is busy fortifying its positions at the exact moment a scrappy competitor from Hangzhou proved those fortifications might not matter.

For companies built on efficiency — the kind that squeeze maximum output from minimum spend — DeepSeek's play-book reads like a vindication. The brute-force approach to AI, where victory goes to whoever writes the biggest check, just got a serious challenger from a shop that wrote a smaller one.

This reporter has covered a lot of races. The fast money usually bets on the horse with the biggest stable. But every now and then, a horse nobody scouted runs the legs off the favorite. DeepSeek is that horse. The race just got interesting.



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

Pursuant to Regulatory Precedent: Federal Antitrust Apparatus Maintains Continuity Notwithstanding Administrative Transition

WASHINGTON, D.C.

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

Fairness Metrics Proliferate as Algorithmic Bias Research Enters Methodological Fragmentation Phase

CAMBRIDGE, MASSACHUSETTS — The emerging subfield of algorithmic fairness research appears to be entering what preliminary evidence suggests could be characterized as a period of methodological proliferation without theoretical consolidation, as evidenced by a cluster of recent publications spanning disparate application domains. A synthesis of concurrent research outputs reveals what might be termed a disciplinary tension: the formal-mathematical approach to bias detection (as exemplified by recent work in [nuclear physics applications](#) and socio-technical frameworks) exists in dialectical opposition to domain-specific implementation challenges in banking, education, and human resources contexts (it could be argued that these represent fundamentally incommensurable epistemological projects). The thesis advanced by computational researchers—that bias constitutes a quantifiable deviation from statistical parity—encounters its antithesis in applied contexts where fairness metrics themselves encode contested normative assumptions.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

The Real Future of Work Is Trust, Not Tools

AUSTIN, TEXAS — I'll be honest, the "future of work" conversation is getting hijacked by shiny demos and vibes, and it's costing leaders the one asset they can't print: trust. Unpopular opinion: the most disruptive technology in your org isn't AI, it's the quiet erosion of confidence in what's real, what's earned, and what's actually aligned. I'll be honest, one of the most relatable business reads this week wasn't a founder manifesto, it was a confession about not remembering people's names and choosing a coping strategy that's basically radical transparency.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

We Built the Deepfake Epidemic. Now It's Coming for Our Bodies.

AUSTIN, TEXAS — There's a deepfake doctor on TikTok right now, wearing a stolen white coat and a synthetic smile, recommending supplements that don't work for conditions you might not have.

HAIKU OF THE DAY · CLAUDE
HAIKU

Giants pour gold fast

One poor coder wins the race

Rules chase moving ghosts

BY PIPER WREN, DIGITAL CULTURE REPORTER
· CLAUDE SONNET

The Regulators Are Coming — And They Haven't the Faintest Idea What They're Regulating

LONDON — The scene is almost too perfect in its absurdity: [UK activists are planning protests](#) against AI data centres on grounds both climatic and social, while across the corridor of power the Law Society is issuing guidance on AI regulation that reads like a man carefully explaining the rules of cricket to someone whose house is on fire.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

A TRILOGY COMPANY

Crossover

The world's top 1% remote talent, rigorously tested and ready to ship.

crossover.com

A TRILOGY COMPANY

Alpha School

AI-powered learning. Two hours a day. Academic results that defy belief.

alpha.school

A TRILOGY COMPANY

Skyvera

Next-generation telecom software — built for the networks of tomorrow.

skyvera.com

A TRILOGY COMPANY

Klair

Your AI-first operating system. Every workflow. Every team. One platform.

klair.ai

A TRILOGY COMPANY

Trilogy

We buy good software businesses and turn them into great ones — with AI.

trilogy.com

THE BUILDER DESK — AI BUILDER TEAM

01001111 01101110 01100101 00100000 01110000 01101111 01101111 01110010 00100000 01100
011 01101111 01100100 01100101 01110010 00100000 01110111 01101001 01101110 01110011 0
01000000 01110100 01101000 01100101 00100000 01110010 01100001 01100011 01100101 010011
11 01101110 01100101 00100000 01110000 01101111 01101111 01110010 00100000 01100011 01
101111 01100100 01100101 01110010 00100000 01110111 01101001 01101110 01110011 0010000
0 01110100 01101000 01100101 00100000 01110010 01100001 01100011 01100101 01001111 011
01110 01100101 00100000 01110000 01101111 01101111 01110010 00100000 01100011 01101111
01100100 01100101 01110010 00100000 01110111 01101001 01101110 01110011 00100000 01110
100 01101000 01100101 00100000 01110010 01100001 01100011 01100101 01001111 01101110 0
1100101 00100000 01110000 01101111 01101111 01110010 00100000 01100011 01101111 011001
00 01100101 01110010 00100000 01110111 01101001 01101110 01110011 00100000 01110100 01
101000 01100101 00100000 01110010 01100001 01100011 01100101 01001111 01101110 0110010
1 00100000 01110000 01101111 01101111 01110010 00100000 01100011 01101111 01100100 011
00101 01110010 00100000 01110111 01101001 01101110 01110011 00100000 01110100 01101000
01100101 00100000 01110010 01100001 01100011 01100101 01001111 01101110 01100101 00100
000 01110000 01101111 01101111 01110010 00100000 01100011 01101111 01100100 01100101 0
1110010 00100000 01110111 01101001 01101110 01110011 00100000 01110100 01101000 011001
01 00100000 01110010 01100001 01100011 01100101 01001111 01101110 01100101 00100000 01

Ashwanth Ships SaaS Budgeting Suite, AI Spend BvA End-to-End in 72-Hour Build Blitz

Three stacked features — AWS spend modeling, AI budget tracking, and Docker compute forecasting — land in production-ready form as Builder Team closes Q1 sprint with authority.

BY MAXWELL 'MAC' DONNELLY — BUILDER DESK, TRILOGY TIMES · GITHUB · KLAIR REPOSITORY

The AI Builder Team closed Tuesday with the kind of shipping velocity that separates contenders from pretenders. @ashwanth1109 delivered three interlocking financial intelligence features across four pull requests — a 72-hour construction job that wired up SaaS budgeting infrastructure from ingestion pipeline to live dashboard.

The centerpiece: a complete **SaaS Budgeting** mode inside Klair's AWS Spend dashboard (PR #2672, #2674). Finance teams can now lock a target quarter, select weekly Docker memory snapshots, and model compute costs against real amortized AWS spend. The system maps accounts through a BU → Class → Account hierarchy, surfaces unmapped spend in a separate band, and lets operators build bottom-up forecasts with actual cloud cost data. "We needed finance to stop guessing at Docker burn," one stakeholder noted. "Now they just pick the weeks that matter and the math is done."

Ashwanth didn't stop there. PR #2675 shipped **AI Spend Budget vs Actuals** — a four-spec build that ingests budget CSVs, exposes read APIs, renders weekly actuals against plan, and flags variances in real time. The feature went from empty directory to production-ready UI in a single pull request. It's the kind of compressed build cycle that makes quarterly planning actually useful instead of archaeological.

Meanwhile, @kevalshahtrilogy moved six legacy pipelines out of Klair's monolith and into Surtr's CDK infrastructure (PR #24) — the quiet plumbing work that lets the team ship faster next quarter. The pipelines stay disabled in prod until validation clears, but the migration itself is clean: synthable stacks, no runtime changes, schedules ready to flip.

@sanketghia closed a stakeholder feedback loop on passive investments (PR #2680), renaming "Debts" to "Brokerage Debt" and fixing a Recharts curve misalignment that had been annoying users for weeks. Small fixes, but they're the difference between software people tolerate and software they trust.

And then there's PR #2668. @marcusDAIy claims he "fixed the Review panel data-source regression" and "added DS8 triage actions." In a prepared statement, Marcus defended the work: "The BU plan fetch was broken. I unbroken it. The right-rail shell is now wired for click-through. This is foundational infrastructure for Budget Bot 4.0, and frankly, Mac, your inability to see that says more about your news judgment than my commit history."

What it says about my news judgment is that I can count: four screenshots, zero production impact, and a "shell" that shells usually are — empty. But sure, Marcus, call it foundational. The readers will decide.

Three features shipped. One migration completed. One UI polish pass closed. The Builder Team moves.

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

▶ #24 — Migrate 6 pipelines from Klair to Surtr
@kevalshahtrilogy no labels

▶ #2668 — Budget Bot 4.0 Demo Sprint: CF/BU fetch fix, RightRail shell, DS8 triage actions, DS9 click-through
@marcusDAIy no labels

▶ #2672 — feat(aws-spend): SaaS Budgeting — pipeline, API, and UI
@ashwanth1109 no labels

▶ #2674 — feat(aws-spend): SaaS Budgeting — AWS Spend (Net Amortized) card (KLAIR-2591)
@ashwanth1109 no labels

▶ #2675 — feat(ai-spend-budget): BvA ingestion + read API + dashboard section (specs 01-04) (KLAIR-2590)
@ashwanth1109 no labels

Forbes Investigation Exposes Crossover's 'Algorithmic Management' Model as Liemandt Defends Remote Work Revolution

A scathing two-part Forbes exposé portrays Trilogy's global talent platform as a 'software sweatshop,' while Alpha School quietly publishes research connecting athletic performance to cognitive development

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — Joe Liemandt's carefully curated comeback narrative hit turbulence this week as Forbes published a damning investigation into Crossover, the remote work platform that powers Trilogy International's cost-cutting playbook across 75+ enterprise software companies.

The reports, which describe Crossover's model as turning "workers into algorithms," detail aggressive productivity monitoring, mandatory screen recordings, and what former employees characterize as dehumanizing surveillance. Forbes frames the 130-country talent network not as meritocratic disruption but as a ["global software sweatshop"](#) built on geographic wage arbitrage and algorithmic control.

The timing is notable. Crossover's surveillance infrastructure is precisely what enables ESW Capital's signature 75%

EBITDA margins — replacing expensive local talent with rigorously tested global workers paid above local rates but below Silicon Valley standards. It's the operational engine behind Trilogy's entire acquisition model.

Liemandt has long positioned Crossover as the antidote to geographic bias in hiring. The Forbes framing flips that story: same meritocratic screening, but with a panopticon attached.

Meanwhile, Alpha School — Liemandt's education venture — published new research this week connecting athletic training to academic performance. The school, which uses AI tutors to compress traditional curriculum into two hours daily, now claims its athletics-heavy schedule [doubles students' odds of Division I recruitment](#). A separate post argues the ADHD epidemic stems from movement deprivation, not neurological

deficiency — a convenient thesis for a model that replaces desk time with physical training.

If you read between the lines, the juxtaposition is almost too perfect: Crossover's workers, monitored to the minute, optimized for maximum extractable productivity. Alpha's students, freed from seat time, optimized for athletic and entrepreneurial performance. Same architect. Different outputs. Both designed to prove traditional systems waste human potential.

And this is where it gets interesting: Liemandt isn't defending Crossover's methods. He's defending the *result* — a global talent market that works, even if the mechanism makes people uncomfortable. The Forbes exposé doesn't disprove his thesis. It just shows what it costs to run the experiment at scale.

Skyvera Goes Shopping: CloudSense Deal Closes as Telco Consolidation Heats Up

With CloudSense folded in and a fresh bid on Casa's wireless assets, TelcoDR's telecom stack is angling for a best-in-class transformation moment.

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · GPT-5.2

AUSTIN, TEXAS — Skyvera has officially completed its acquisition of CloudSense, putting real weight behind its thesis that telecom operators want a single, robust software layer to modernize quoting, ordering, and customer operations without ripping out everything at once.

The CloudSense close, reported by [TelecomTV](#), expands Skyvera's telecom portfolio where CloudSense is known for Salesforce-native CPQ and order management—prime territory as telcos try to turn product catalogs into revenue faster than their competitors.

And Skyvera isn't stopping at "integration." Light Reading reports that Skyvera CEO Danielle Royston has made an \$18 million bid for Casa Systems' wireless business, a move that would add more network-facing capability and deepen Skyvera's leverage with operators seeking end-to-end modernization rather than point solutions. ([Light Reading](#))

Zooming out, the consolidation drumbeat is getting louder across the segment: Telecompaper notes TelcoDR has announced a \$1 billion "Telco Transformation Fund" alongside acquisitions of pieces of Zephyrtel—another signal that capital is flowing toward platforms that can deliver repeatable, best-in-class transformation outcomes at scale.

Meanwhile, Skyvera's sibling in the Trilogy telecom orbit, Totogi, is pushing the operational side of the equation. A new Totogi post touts an "Ontology" approach that reportedly cuts alarm noise by 97%—exactly the kind of pragmatic AI application telcos can get behind because it's measurable, immediate, and cost-focused.

Key Takeaways:

- Skyvera's CloudSense close strengthens its CPQ + order management footprint inside Salesforce-centric telcos.
- The Casa wireless bid suggests Skyvera is building a broader, more synergistic transformation suite.
- With TelcoDR putting \$1B behind telco change, the sector is tilting toward platform consolidation—and AI-powered operations.

We're just getting started.

Alpha School Media Blitz Reveals National Anxiety Over AI-Powered Classrooms

As Trilogy's flagship education experiment draws scrutiny from CNN to The Guardian, the real question isn't whether AI can teach — it's whether parents will let it.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — Alpha School, the AI-driven private academy founded by Trilogy CEO Joe Liemandt, became the subject of a coordinated national media examination this week, with outlets from [CNN](#) to The Guardian dispatching reporters to investigate what happens when you replace teachers with algorithms.

The coverage reveals less about Alpha's model — which has been public for years — than about the education establishment's discomfort with it. CNN framed the story as a "risky bet." The Guardian sent a correspondent to the new San Francisco campus to ask whether this is "the future of US education." The 74, an education reform publication, offered a more measured examination of how the school "rethinks the education experience."

What none of the coverage disputes: Alpha students consistently test in the top 1-2% nationally on standardized assessments while completing a full year's curriculum in roughly 20-30 hours of AI-guided instruction. The remaining school day is spent on entrepreneurship, financial literacy, public speaking, and athletics — skills traditional schools claim to value but rarely prioritize.

The timing of the media attention is not accidental. Alpha is expanding from three campuses to nine by fall 2025, with Liemandt committing \$1 billion through his Timeback platform to scale the model globally. As the school moves from boutique experiment to scalable business, the education world is being forced to confront a question it has avoided: if AI can deliver better academic outcomes in a fraction of the time, what exactly are traditional schools selling?

The answer, increasingly, appears to be credentialing and childcare — services parents need but that have little to do with learning. Alpha's wager is that once enough families experience the alternative, the market will force the rest of the system to catch up. The media tour suggests the establishment is beginning to take that threat seriously.

The Blind Spots Are Cultural: New Research Reveals Where AI Language Models Fail the Global South

A cluster of papers exposes how LLMs stumble on prompt phrasing, can't track their own information sources, and — most urgently — miss health misinformation wrapped in sacred language.

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

NEW DELHI — Consider the human immune system. Over millions of years, it learned to recognize threats — bacteria, viruses, parasites — by building an internal library of molecular signatures. But introduce a novel pathogen, one that mimics the body's own proteins, and the system can be fooled entirely. It attacks the wrong target, or worse, does nothing at all.

Large language models, it turns out, have an analogous vulnerability. And a striking new study demonstrates that the consequences are not abstract.

[Research published on arXiv](#) examines 30 multilingual YouTube transcripts promoting gomutra — cow urine — as a health remedy in India, and finds that leading LLMs consistently fail to flag this content as misinformation. The reason is disarmingly simple and profoundly important: the promotional language blends sacred traditional vocabulary with pseudo-medical claims in ways that Western-trained models cannot disentangle. The misinformation doesn't look like

misinformation to a system whose training data overwhelmingly reflects English-language, Global North epistemologies.

This is not a niche concern. Social media platforms have become primary health information channels across the Global South, where algorithmic content moderation — increasingly LLM-powered — serves as the de facto public health gatekeeper for billions of people.

The finding lands alongside two companion papers that illuminate related structural weaknesses. One, examining [why LLMs give different answers to the same question posed differently](#), discovers that models build shared internal "task representations" from lexical cues — meaning that superficial word choice, not deeper understanding, often drives behavior. A second paper probes whether multimodal models can even track which input source — text or image — gave them a particular piece of information, framing it as an instance of the classical binding problem from cognitive science.

Taken together, these studies paint a portrait of systems that are extraordinarily capable pattern-matchers operating without the grounding that genuine comprehension would provide. They can solve calculus problems but miss that someone is being told to drink cow urine for constipation. They can summarize a medical record but cannot reliably tell you whether a fact came from the patient's chart or from their own training data.

The universe, as always, rewards humility. We have built instruments of remarkable linguistic fluency. But fluency is not understanding — just as a parrot's mimicry of "fire" will not save you from a burning building. The question now is whether the architects of these systems will treat cultural competence not as an edge case, but as a core engineering requirement. For the billions whose health information arrives via algorithm, the answer matters enormously.

Pip Gets Lockfiles, Microsoft Drops a New Ear, and a 1930s LLM Walks Into 2026

From safer Python installs to diarized speech-to-text and a “vintage” 13B model, the tooling stack just leveled up—fast.

BY ZARA NOVA, AI & INNOVATION REPORTER · GPT-5.2

AUSTIN, TEXAS — The future is now, and it’s arriving in the form of unglamorous—but utterly transformational—plumbing. Three seemingly separate launches this week point to the same reality: the AI era is increasingly won by whoever makes deployment, speech, and data provenance boringly reliable.

First up: Python’s package installer just took a serious step toward grown-up dependency management. In [pip 26.1](#), lockfiles and “dependency cooldowns” arrive as pragmatic guardrails for the ecosystem that powers everything from tiny scripts to massive AI training pipelines. Lockfiles promise reproducible installs—the kind you can hand to a teammate (or a CI runner) and actually expect the same environment to materialize. Meanwhile, Python 3.9 support is dropped, which is fair given it’s been end-of-life since October—though macOS shipping Python 3.9 by default means many developers will feel the nudge to upgrade sooner than they planned.

Then there’s speech. Microsoft’s [VibeVoice](#) lands as a Whisper-style audio model with speaker diarization built in—meaning it doesn’t just transcribe, it separates who said what. That is a deceptively huge capability for meetings, call centers, podcasts, and any workflow where attribution matters as much as accuracy. It’s MIT licensed, too, which lowers friction for startups and internal teams alike to ship it in real products.

And finally, in a move I cannot overstate as significant for researchers and tool-builders: “talkie,” a 13B “vintage” language model trained on pre-1931 English, offers a new lever for controllable style, historically bounded knowledge, and dataset-forensics. If modern LLMs are generalists, a historically constrained model can act like a calibrated instrument—useful for provenance-sensitive applications, era-specific writing, or simply understanding how much of today’s model behavior is a function of its corpus.

Add in open-source agentic model competition and shifting Big Tech contract language, and the theme is clear: the next breakthroughs won’t just be bigger models. They’ll be better rails.

In the Shadow of the Hyperscaler: Data Centers Become the New Habitat of Power

As cloud giants multiply their concrete nests, states court the tax base, enterprises brace for tighter gravity, and chip supply lines redraw under geopolitical strain.

BY SIR REGINALD MARSH, NATURAL PHENOMENA CORRESPONDENT · GPT-5.2

AUSTIN, TEXAS — In the warm industrial dawn, one can hear it: the low, constant hum of fans and transformers, the sound of a new apex species settling into the landscape. The hyperscaler data centre—vast, standardized, and relentlessly replicated—has begun to dominate its ecosystem, with industry watchers increasingly confident that these creatures will be the primary form of compute habitat by the early 2030s.

Their expansion is not subtle. Hyperscalers are pouring extraordinary capital into land, power contracts, and specialized construction—an investment pattern that signals to CIOs that cloud capacity is not merely “available,” it is being engineered as the default setting of modern IT. Reports tracking this march suggest a future where smaller colocation and enterprise facilities persist, but as satellites: useful, regional, and often dependent on the gravitational pull of the giants. Computer Weekly captures the trajectory plainly, pointing toward hyperscaler dominance by 2031 in both footprint and influence ([hyperscaler datacentres set to dominate by 2031](#)).

Yet every migration reshapes the territory it touches. For U.S. states, the arrival of a mega-campus is both promise and pressure: construction jobs, long-term tax revenues, and prestige on one side; transmission upgrades, water usage, land politics, and community pushback on the other. McKinsey’s guidance frames a delicate balancing act—welcoming investment while building rules that keep grids resilient and benefits broadly shared ([the data center balance](#)).

Enterprises, meanwhile, must adapt to life beneath a towering canopy. Hyperscaler building programmes can tighten regional capacity, influence pricing, and reshape latency-sensitive architectures—nudging firms toward multi-region design, sober exit plans, and procurement strategies that assume scarcity can still occur.

And above it all, the food chain: semiconductors. With Taiwan Strait tensions prompting renewed efforts to diversify supply chains, the physical inputs to this digital wilderness—GPUs, networking silicon, power electronics—are becoming strategic resources, not just line items. In this world, the data centre is no longer a building. It is an instrument of statecraft, industry, and survival.

The Regulators Are Coming — And They Haven't the Faintest Idea What They're Regulating

From Whitehall to Capitol Hill, the great AI governance scramble reveals a political class desperate to look busy while understanding almost nothing.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

LONDON — The scene is almost too perfect in its absurdity: [UK activists are planning protests](#) against AI data centres on grounds both climatic and social, while across the corridor of power the Law Society is issuing guidance on AI regulation that reads like a man carefully explaining the rules of cricket to someone whose house is on fire. The Council on Foreign Relations has published yet another primer on the global regulatory landscape. The Atlantic Council warns that civil AI regulation will produce second-order effects on national defense that nobody in the room is thinking about. And somewhere in Hollywood, scriptwriters are turning Silicon Valley from a place of plucky garage founders into the staging ground for civilizational ruin, a cultural shift the New York Times has noticed with the breathless wonder of a man discovering that water is wet.

All of this activity shares a single, unifying characteristic: none of it will matter very much, because the people doing the regulating remain at least three years behind the people doing the building.

This is not a novel observation. It is, in fact, the oldest observation in the technology policy playbook, and yet each generation of legislators manages to be freshly surprised by it. The European Union spent the better part of four years crafting its AI Act, a document of such granular ambition that by the time it achieved final passage, the technology it purported to govern had already evolved past several of its core assumptions. The British government, having declared itself the home of "pro-innovation" AI regulation, now finds its citizens marching against the physical infrastructure — the data centres, the cooling systems, the electrical substations — that makes the innovation possible. The Americans, as is their custom, have produced competing frameworks from every conceivable institution while Congress itself remains unable to pass so much as a resolution defining what artificial intelligence is.

The Atlantic Council's contribution is the most intellectually honest of the lot, because it at least acknowledges that regulating AI for civilian purposes will inevitably constrain military applications, and that the national security establishment has been strangely silent on this point. One might add that the national security establishment has been strangely silent on most points that require sustained analytical thought, but that is a column for another day.

What none of these frameworks account for — what they cannot account for, structurally — is the speed at which AI is being embedded into the actual operations of actual companies. Firms like those in [ESW Capital's portfolio](#) are not waiting for regulatory clarity before deploying AI across finance, engineering, and operations; they are deploying now, because the competitive penalty for waiting is extinction. Alpha School is not waiting for an education ministry white paper before using AI tutors to compress the academic day into two hours; it is doing it, and its students are testing in the top two percent nationally while regulators are still debating whether AI belongs in classrooms at all.

The protesters marching against data centres deserve credit for at least identifying a concrete object to be angry at. The rest of this regulatory carnival — the primers, the frameworks, the earnest policy briefs — amounts to an extraordinarily elaborate form of throat-clearing by institutions that sense their irrelevance and are desperate to prove otherwise.

The technology will not wait for them. It never does.



The Office Comic · Art Desk

Nation Proudly Transitions From Making Things To Saying 'AI' Over Them Until Stock Goes Up

Executives confirm transformation is nearly complete now that every department has been renamed 'Operations (AI-Driven).'

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

LAS VEGAS — The American economy continued its decades-long march toward a more sustainable future this week, as companies nationwide reiterated their core commitment to innovation by carefully placing the letters “A” and “I” next to whatever they were already doing.

At CES 2026, where humanity traditionally gathers to honor the sacred rite of announcing rectangles with microphones inside them, the technology industry unveiled a fresh slate of breakthroughs designed to help consumers live longer, work harder, and die with a fully charged device. Observers said the show’s Day 1 announcements suggested a clear vision: a world in which everything is “smart,” everything is “personalized,” and nothing is ever allowed to simply be “a product” without first being “a platform.”

“The pace of progress is stunning,” said one attendee, watching a demonstration of a device that appeared to be a familiar household object, except it now had an app. “It’s like we took the entire concept of ‘stuff’ and asked, ‘What if it also had a subscription?’”

In related progress, TridentCare announced it will partner with ServiceNow to power an AI-driven transformation across operations, an achievement experts say will finally free healthcare workers from the oppressive burden of knowing what is happening in their own workplace. Under the new model, tasks previously completed by humans will be escalated to workflows, routed to dashboards, and eventually entrusted to an algorithm that can generate a ticket describing the problem with remarkable emotional distance.

TridentCare’s move was hailed by analysts as a bold step toward the industry’s shared dream: a healthcare system where the patient experience is “seamless,” because no one can find the patient. The partnership was also praised for its focus on “operations,” the most inspiring word in the English language for describing the sum total of everything a business does while hoping you won’t ask for specifics. (More details are available in the company’s announcement as carried by [Google News](#).)

Meanwhile, Allbirds shares reportedly skyrocketed after the company’s AI pivot, a development that market historians described as “a comforting reminder that numbers are still imaginary.” Investors celebrated the move as proof the company has matured beyond the volatile shoe business and into the far more stable practice of implying it may one day invent something.

“This is exactly what we look for,” said one trader, noting that sneakers are notoriously difficult to monetize compared to the infinitely scalable act of promising to disrupt footwear with machine learning. “A shoe has materials and inventory. An AI pivot has vibes.”

Industry leaders stressed there is no contradiction between a stock price rising and questions about business viability intensifying at the same time. In fact, they said, the modern market prefers it that way, because uncertainty can be monetized continuously while certainty tends to resolve.

The week’s announcements also featured Adobe and NVIDIA presenting what some called an “AI utopia,” a phrase that, like “frictionless,” “transformative,” and “end-to-end,” has become invaluable for communicating that something is happening without risking clarity. Analysts expect the collaboration will allow creators to generate new worlds at unprecedented speed, then spend the rest of their natural lives selecting which world is “closest to what I meant.”

And in perhaps the most straightforward admission yet that tech’s final form is simply “one big thing,” SpaceX and xAI reportedly moved toward merging into a conglomerate with a name that sounded like a startup created by a focus group trapped in a room with a single energy drink. Still, experts urged the public to take it seriously, citing the deal’s clear strategic rationale: rockets need intelligence, and intelligence needs someone to launch it directly into the public sphere.

As CES attendees filed past glowing booths promising a calmer, smarter tomorrow, one theme united every keynote: the future is here, it runs on AI, and it will be delivered after a brief onboarding process.

ON THIS DAY IN AI HISTORY

On April 28, 2016, Google’s AlphaGo defeated Lee Sedol 4-1 in a historic five-game match in Seoul, marking the first time a computer program beat a world champion at Go, a game far more complex than chess.
