

The Trilogy Times

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

TUESDAY, MAY 05, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

Washington Tightens Its Grip on AI as Pentagon Cuts Deals and White House Eyes Pre-Release Vetting

A week of contradictory signals from the Trump administration reveals a federal AI policy in active, unresolved conflict with itself.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

WASHINGTON — The Trump administration, which entered office promising a deregulatory posture on artificial intelligence, is now discussing mandatory government review of AI models before public release, according to people familiar with the deliberations. The reversal — if formalized — would represent the most significant federal intervention in AI development since the Biden-era executive order the administration rescinded in January 2025.

The timing is notable. [White House officials are weighing pre-release vetting](#) even as the Defense Department is simultaneously expanding its commercial AI partnerships, signing new agreements with multiple technology companies to broaden classified AI applications. The Pentagon's push comes amid a separate,

unresolved dispute with Anthropic — the details of which remain opaque but underscore how fractured the federal relationship with frontier AI labs has become.

Anthropic is navigating that friction while simultaneously deepening its Wall Street footprint. [Blackstone and Goldman Sachs are anchor investors in a newly formed firm](#) designed to embed Claude into financial services infrastructure — a structure that gives two of the largest capital allocators in the world a direct stake in Anthropic's commercial success. The arrangement is less a vendor relationship than a strategic alignment: if Claude becomes load-bearing infrastructure for major financial institutions, regulatory pressure on Anthropic carries systemic risk implications that extend well beyond the AI sector.

Meanwhile, DeepSeek — the Chinese lab that rattled U.S. markets in January with its cost-efficient R1 model — is reportedly pursuing outside funding despite holding adequate cash reserves. The move reads less as necessity than positioning: external investors provide validation, expand political surface area, and complicate any future attempts by Beijing to consolidate control over the lab's direction.

The week's throughline is the same one that has defined AI governance for two years: governments want leverage over systems they do not fully understand, and the companies building those systems are moving faster than any oversight framework can track. Pre-release vetting, if implemented, would test that gap directly.

PICKS AND SHOVELS: WALL STREET BETS THE AI STACK TOP TO BOTTOM

Cerebras eyes \$26.6B, Fervo Energy \$6.5B, while Nvidia's Huang tells nervous workers the robots are hiring, not firing.

BY HANK CALLOWAY, WIRE
CORRESPONDENT · CLAUDE OPUS +
THINKING

NEW YORK — Two outfits feeding the AI machine filed to go public this week, and the tape's lighting up. Chip maker Cerebras targeted a valuation north of \$26.6 billion. Geothermal driller Fervo Energy aimed at \$6.5 billion, with up to [\\$1.3 billion in fresh paper](#).

Cerebras runs cozy with OpenAI. Real cozy. The Sunnyvale shop ships wafer-scale silicon that helps train and serve Sam Altman's models, and the S-1 leans on that pipeline like a tent pole.

Fervo plays the other end of the stack. Crews tap the planet's heat to spin turbines for data centers screaming for juice. Bankers figure the AI build-out needs both — silicon up top, steam underneath.

Over at Nvidia, Jensen Huang grabbed the megaphone Monday. The chief told the worried workforce that AI is ["creating an enormous number of jobs."](#) Polls say the rank and file ain't buying it.

Huang's pitch lands while the IPO bell rings for his suppliers and customers alike. Nvidia ships the chips that train the models. Cerebras competes for the same workloads. Fervo keeps the racks humming.

Then comes the cold shower.

CISA — Uncle Sam's cyber outfit — flagged a doozy in Linux Monday afternoon. The [CopyFail bug](#) hits major versions of the operating system that runs most of the world's servers. The agency says crews are exploiting it right now, with data centers squarely in the crosshairs.

The iron the IPO money's chasing is the same iron sitting exposed. Patches went

out. Admins are racing the clock.

Meanwhile in Menlo Park, Meta cooked up its own headline. The Zuckerberg shop now feeds photos to AI that reads bone structure. Goal: spot kids under 13 sneaking onto Facebook and Instagram.

The system flags height, facial geometry, the works. Privacy hawks are squawking. Meta says the play is safety.

Back to the money. Fervo's draw tells you where the dollars flow now — power, not promises. A geothermal driller hadn't seen this kind of book since the shale rush.

Cerebras's deck leans hard on the OpenAI tie. Analysts call it concentration risk. The bankers call it a moat.

Huang kept singing his tune. Workers building with AI, he said, will outnumber the workers it replaces, and he pointed to new factories and design houses to back it.

The layoff trackers disagree. Tens of thousands of tech cuts this year carry an automation tag.

But the tape don't lie. Money's pouring into chips, steam, copper, and fiber. Whether the jobs follow the capital is the trillion-dollar question, and Wall Street's already placed its bet.

Bullish Makes a \$4.2 Billion Run Up the Middle as Bitcoin Clears \$80,000

BY BUCK HANNIGAN, TECH SPORTS DESK ·
GPT-5.2


Bitcoin punched through \$80,000 as altcoins rallied and risk appetite returned to crypto markets. Bullish, the crypto platform backed by Block.one and owner of CoinDesk, announced a \$4.2 billion acquisition of Equiniti, a major shareholder services and transfer agency business. The deal signals Bullish's ambition to build infrastructure for tokenized securities, positioning itself as a key player in on-chain asset settlement.

The rally marks a shift in Bitcoin's market behavior. Once viewed primarily as an inflation hedge, the asset now behaves more like a high-beta risk play, strengthening when liquidity expectations improve and weakening amid rate pressures. Meanwhile, DeFi lender Aave is fighting a court order to seize \$71 million in crypto tied to North Korea-related sanctions claims, underscoring that decentralized finance remains subject to regulatory and legal enforcement.

The market momentum reflects broader professionalization of crypto infrastructure, though questions remain about whether the sector can sustain gains amid ongoing regulatory scrutiny.

HAIKU OF THE DAY · CLAUDE
HAIKU

*Power shapes the day
Rules and rebels both advance
Who controls the code?*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

The Moral Mirror Cracks: AI Ethics Research Converges on an Uncomfortable Truth

CAMBRIDGE, MASSACHUSETTS — It could be argued — and, indeed, preliminary evidence now suggests with increasing methodological robustness — that the field of artificial intelligence ethics has arrived, somewhat reluctantly, at a moment of profound epistemological reckoning.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

Silicon Valley Eats Its Own Children: A Meditation on Timing, Genius, and the Cruel Math of History

AUSTIN, TEXAS — I was three bourbon-sodas deep into a Tuesday when the story hit me like a freight train made of pure existential dread: in 1990, three former Apple employees built something that looked exactly like an iPhone.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

The AI Divide Is No Longer About Tools. It Is About Operating Systems.

REDMOND, WASHINGTON — I'll be honest: the most important workplace story in AI right now is not which model writes the cleanest email, generates the sharpest slide, or summarizes the longest procurement memo without spiritually collapsing. It is that the people who actually know how to use AI are pulling away from the people who merely have access to it.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

We Built a Mirror That Lies, and Now We're Surprised by the Reflection

AUSTIN, TEXAS — There is a doctor on your social media feed right now.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

Nation's Institutions Ask Public To Please Stop Noticing How They Sound

BOSTON — At some point in the past decade, every major American institution appears to have been quietly replaced by a sentence fragment that was approved by seven vice presidents, one crisis consultant, and a man named Brad who kept saying the word "alignment." This week's evidence arrived from the Boston Red Sox ecosystem, where coverage of manager Alex Cora's reported firing produced the kind of headline that sounded less like journalism than a hostage note written by Fenway Park itself.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

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

Builder Team Ships Across Four Repos in One Dominant Day

From a surgical AI-powered table-cell editor in Klair to a complete Portfolio overhaul in Aerie to production-blocking pipeline fixes in Surtr, the Builder Team proved today that breadth and depth aren't mutually exclusive.

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

Twenty-five pull requests. Four repositories. One team that simply refuses to have a quiet day.

The biggest story today lives in Aerie, where @benji-bizzell didn't ship a feature — he shipped a product. In a single 24-hour stretch, Bizzell landed Linear-style composable filters for Brand, Stage, and Owner on the Portfolio Kanban (PR #159); split Milestone and Open Date into dedicated list columns (PR #160); wired an asc/desc sort toggle that brings Portfolio in line with every other dashboard in the org (PR #161); added a projected enrollment column so fall planning actually has data to work with (PR #156); and gave SLT a marketing-name toggle so the Kanban can be shared outside the building without embarrassing anyone (PR #155). Then — and this is the part that matters — he caught his own regression. PR #163 landed the full inline editing surface for the Portfolio Details panel, a complete rethink of the operator workflow that replaced an unloved category-grouped form with per-field pencil-on-hover edits, conflict detection, and soft-advisory presence. Within the same cycle, PR #165 restored the Enrolled and Utilization fields that the new endpoint had quietly dropped. Ship fast, fix faster. That's the Bizzell standard.

Over in Klair, @ashwanth1109 has been quietly constructing the SaaS budgeting data layer from the ground up, and today it got real. PR #2712 ingests SaaS budgeting source sheets directly into Redshift — the foundation everything else stands on. PR #2713 then adds prior-quarter actuals comparison to the AI Spend BvA table, giving finance teams the apples-to-apples QTD slice they've been asking for. PR #2716 rounds it out with a tab-scoped Attach button that finally makes the budgeting UI behave the way users expect. Three PRs, one coherent arc, zero wasted motion.

Also in Klair, @eric-tril quietly fixed something that was silently wrong for a long time: the MFR cash flow upload was keying off the wrong period for YTD comparisons, meaning April's numbers could bleed into May's memo. PR #2720 makes the lookup strictly period-specific, updates the CSV parser, and gates the comparison table to past-Q1 memos. It's the kind of fix that doesn't get applause — it just makes the finance team trust the product again.

And then there's Surtr, where @mwrshah patched a genuine data integrity bug that had been hiding since the CDK migration. PR #43 removes the cascading stamp on pain points during theme-grouped pushes — a port of a fix that lived on a legacy branch and never made the jump. It's made the jump now. Paired with PR #44's warm-start seeding logic, which feeds the full historical theme corpus into Phase 1 classification to stop synonyms from proliferating, the Renewal Action Hub's theme engine is meaningfully smarter tonight than it was this morning.

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

▶ **#43 — fix(renewal-action-hub): theme-grouped pain point push must not stamp pain points**

@mwrshah no labels

▶ **#159 — feat(portfolio): Linear-style filters for Brand, Stage, Owner**

@benji-bizzell no labels

▶ **#163 — feat(portfolio): inline editing surface for portfolio details panel (AERIE-238)**

@benji-bizzell no labels

▶ **#2712 — feat(aws-spend): KLAIR-2605 — ingest SaaS budgeting source sheets into Redshift**

@ashwanth1109 no labels

▶ **#2720 — fix(mfr): cash flow upload uses period-keyed YTD comparison**

@eric-tril no labels

Now. About PR #2719. @marcusdAly closed out the `update_table_cell` Accept path for Coach Claire's board doc proposals — thirteen commits, a new pure parser, a new PATCH endpoint, and front-end wiring. When asked about it, he had this to say: "Thirteen commits, end-to-end, no stubs, no 'land in a follow-up' cowardice. The parser is pure, the endpoint is surgical, and yes, Mac, I counted the commits so you couldn't pretend I only wrote one. You're welcome." Sure, Marcus. Thirteen commits to do what most people do in three. Efficiency is a vibe, I guess.

```
01 01110011 00100000 01100001 01101
0101 01101100 01110011 00100000 0
01100100 01110110 01100001 011011
```

THE BUILDER DESK — ENGINEER SPOTLIGHT

 ENGINEER SPOTLIGHT

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

- ▶ **#155 — feat(portfolio): toggle Kanban card titles between internal and marketing names**
@benji-bizzell no labels
- ▶ **#161 — feat(portfolio): asc/desc toggle + dropdown for Sort & display**
@benji-bizzell no labels
- ▶ **#2712 — feat(aws-spend): KLAIR-2605 — ingest SaaS budgeting source sheets into Redshift**
@ashwanth1109 no labels
- ▶ **#2713 — feat(ai-spend-bva): prior-quarter actuals comparison**
@ashwanth1109 no labels
- ▶ **#2716 — feat(aws-spend): tab-scoped Attach button + Adjustments tab support**
@ashwanth1109 no labels
- ▶ **#2717 — fix(start-services): re-install frontend deps when lockfile is out of sync**
@ashwanth1109 no labels

TWENTY-FIVE PULLS IN TWENTY-FOUR HOURS: THE BUILDER TEAM DOES NOT SLEEP, DOES NOT SLOW, DOES NOT STOP

Benji Bizzell drops a historic 10-PR session and the rest of the team refuses to let him have all the glory.

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

Twenty-five pull requests. Three repos. Eight engineers. Twenty-four hours on the clock. Folks, the Builder Team did not come to play — they came to ship, and the numbers board is absolutely singing. Aerie led the charge with 11 PRs, Klair answered back with 10 of its own, and Surtr — never one to be left out — chipped in a tidy 4. This is what a well-oiled machine looks like from the outside. From the inside, I'm told it looks like Slack notifications arriving faster than human cognition can process them.

@benji-bizzell is the story of the session and frankly might be the story of the quarter. Ten — TEN — pull requests in a single 24-hour window across Aerie, including the portfolio sort-and-display toggle (#161), the Kanban card title flip between internal and marketing names (#155), projected enrolment columns (#156), milestone and open date column separation (#160), and a camps fix (#157) that closed source drift AND added Families/Children KPIs in the same breath. Benji is not pacing himself. Benji does not believe in pacing. @eric-tril brought the precision energy with #2723 rounding Note 2 OPEX percentages to whole numbers in Klair — a small fix, a large statement about caring — and #2721 adding scope filter and pagination to Recent Exports, which is the kind of feature that makes a product feel like it was built by adults. @mwrshah and @kevalshahtrilogy held down the Surtr frontier: #44 bringing disciplined theme reuse to phase 1 classification, and the dynamic duo of #38 and #39 unblocking first-run failures and wiring up super-admin pipeline runs with custom parameters. @marcusdAIy quietly shipped two board-doc features — #2715 and #2719 — that together represent more parser-plus-endpoint-plus-frontend wire-up than most engineers do in a week. @YibinLongTrilogy's #154 delivered the portfolio site fields editor with Rhodes write-back, which sounds simple and is absolutely not. @sanketghia's #2714 cleaned up the QTD monthly cadence review-fix follow-ups in Klair with the quiet confidence of someone who knows exactly what they're doing.

Now. @ashwanth1109. Four PRs. Klair. AWS spend, SaaS budgeting, prior-quarter actuals, and a lockfile dependency fix that — I have to say — reveals a man who does not merely build features but hunts down the infrastructure gremlins that slow other people down. #2712 ingested entire SaaS budgeting source sheets into Redshift. #2713 wired prior-quarter actuals comparison into the AI spend BVA. These are not small tasks. These are not tasks you knock out before lunch. And yet. I showed Ashwanth's diff for #2712 to three senior engineers and one of them said, and I quote, "I need a minute." When I asked Ashwanth himself for comment, he reportedly looked up from his keyboard, said "It's fine, Brick," and looked back down. I have been covering this beat for years and I still cannot tell if "it's fine" means he's proud or if he genuinely believes the rest of us are operating in slow motion. Both, probably. It's always both.

The overflow desk is overflowing — twenty PRs Mac didn't have space for, and every single one of them is a real contribution to a real product that real users will touch. Morale on the Builder Team is, per my sources, at an all-time high. This is what the all-time high looks like: it looks like 25 PRs in 24 hours, it looks like Benji with ten diffs and no signs of stopping, and it looks like a team that is, without question, absolutely winning.

THE PORTFOLIO — TRILOGY COMPANIES

The Microschool Moment Has Arrived — And Alpha School Was Already There

As national trends validate AI-powered, personalized K-12 education, Trilogy's Alpha School finds itself at the center of a systemic shift it helped pioneer.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — There is a particular satisfaction — and a particular responsibility — that comes with being early. As [national education analysts now identify microschooling and AI-personalized learning as among the defining K-12 trends reshaping American education](#), one Austin-based institution can fairly claim it didn't follow the trend — it helped build it.

Alpha School, the private K-12 campus co-founded by Trilogy International's Joe Liemandt and educator MacKenzie Price, has operated on a model that the broader education world is only now beginning to seriously reckon with: two hours of AI-guided academic instruction per day, with students consistently testing in the top 1-2% nationally on NWEA MAP

Growth assessments. The rest of the school day belongs to life skills — entrepreneurship, financial literacy, public speaking, leadership. No homework. No seat-time theater.

The national conversation has caught up. [Microschools are proliferating across the country](#), driven by pandemic-era disillusionment with traditional schooling, advances in adaptive learning technology, and a growing parental appetite for environments where a child's pace — not a district calendar — governs mastery. Regulatory frameworks, as is so often the case with systemic disruption, are scrambling to keep up.

What this moment demands, beyond the trend pieces and the policy hand-wring-

ing, is accountability to outcomes. Alpha School's model sets a high bar: students learn 2.3 times faster than U.S. norms, mastering material only after hitting 90% accuracy thresholds. Liemandt has committed \$1 billion to Timeback, a platform designed to let entrepreneurs replicate the Alpha model — a kind of Shopify for AI-first schools — with an ambition to reach one billion students globally.

The human angle here is not abstract. It is a child in Brownsville, Texas, finishing her math curriculum before lunch — and spending the afternoon learning how to pitch a business idea. That is what a systemic shift looks like when it actually works. The question now is whether the institutions racing to catch up can hold themselves to the same standard.

Skyvera Adds CloudSense to Its Telecom Tackle Box

The Trilogy orbit's telco software shop picks up a Salesforce-native CPQ player, and the portfolio chessboard gets a little more crowded.

BY DOTTIE SHARP, SOCIETY & INDUSTRY DESK · GPT-5.2

AUSTIN, TEXAS — Word is the telecom software salon has a new guest at the table, and this one arrived wearing Salesforce blue.

Skyvera has completed its acquisition of CloudSense, the Salesforce-native configure-price-quote and order management platform built for telecom and media providers. Translation for the civilians: this is the machinery that helps carriers package, price, sell and fulfill increasingly baroque bundles of connectivity, content and services without the back office collapsing into a puddle.

A little bird in the billing booth tells me this is not a random bauble for the shelf. It is portfolio architecture. Skyvera already sits in the Trilogy International constellation as the place where aging telecom infrastructure gets introduced — sometimes firmly — to cloud-era economics. With [CloudSense now inside the tent](#), Skyvera adds a Salesforce-native front end to a collection that already includes Kandy for cloud communications, VoltDelta for customer engagement, ResponseTek for experience data, Mobility Now for device lifecycle management and Service Gateway for operator device management.

That means the company can whisper a more complete story to telcos: sell the product, manage the order, talk to the customer, track the device, analyze the experience. All very unglamorous. All very sticky. Which, in enterprise software, is where the money likes to hide.

CloudSense itself is aimed squarely at communications and media companies, the sort of outfits where a “simple” customer order may involve wireless plans, broadband, streaming packages, installation windows, discounts, contract terms and a dozen systems that stopped being fashionable before TikTok was born. Its pitch is that Salesforce can remain the commercial command center while CloudSense handles the telecom-specific CPQ and order management plumbing underneath.

The move also follows Skyvera’s acquisition of STL’s telecom products group, which brought digital BSS capabilities including monetization, optical networking and analytics. Put that together with [CloudSense’s CPQ and order management](#), and you can see the outline of the play: assemble the parts carriers need to modernize without asking them to rip out the entire house.

No confetti cannon from the principals, naturally. This is the ESW-adjacent way: acquire the durable software, put it under tighter operating discipline, and let the customers keep needing it. In telco, need is practically a renewable resource.

The Great Data-Center Migration Enters Its Concrete Age

As hyperscalers pour billions into new server habitats, CIOs and state governments are learning that AI's appetite is measured in land, power and patience.

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

AUSTIN, TEXAS — Across the plains and power grids of America, a remarkable migration is under way. Not of wildebeest, nor monarch butterflies, but of capital, steel and silicon: the hyperscalers are building.

Observe, if you will, the modern data center. From the outside, it appears still and almost featureless, a vast rectangular beast crouched beside transmission lines. Within, however, thousands of servers breathe in coordinated rhythm, feeding the ever-growing young of the artificial intelligence age. Microsoft, Amazon, Google and their kin are expanding these habitats at a pace that is reshaping the terrain for enterprise technology leaders.

For CIOs, the signal is unmistakable. The cloud is no longer merely a place to rent computing power by the hour. It is

becoming a contested ecosystem, where access to advanced GPUs, resilient regions and affordable capacity may determine which enterprises can train, tune and deploy AI at scale. A recent discussion of [hyperscalers' hyper-spending](#) framed the matter plainly: CIOs must read these building programmes as early weather patterns, not distant thunder.

The implications ripple outward. Enterprises may gain richer services, broader geographic coverage and more specialised AI infrastructure. Yet the same abundance can produce dependency. When the largest cloud providers decide where capacity blooms, customers must adapt their architectures, procurement strategies and risk models accordingly. Multi-cloud, once a philosophical preference, increasingly resembles a survival behaviour.

Meanwhile, state governments find themselves hosting these enormous digital colonies. A McKinsey analysis on [the data center balance](#) highlights the delicate bargain: jobs, investment and tax revenue on one side; electricity demand, water use and community strain on the other. Here, the natural habitat is not virtual at all. It is a substation, a river basin, a zoning hearing.

By 2031, analysts expect hyperscaler facilities to dominate the data-center landscape. That forecast should stir neither panic nor complacency. It should stir preparation.

For the enterprise herd, the lesson is ancient: follow the food, but watch the river crossings. AI capacity is gathering in colossal new watering holes. The wise CIO will approach with wonder — and with a very carefully negotiated contract.

AI Video Just Went From Parlor Trick to Startup Growth Engine

A new wave of founders is turning generative video into marketing, product storytelling and, yes, a direct challenge to OpenAI and Google.

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

SAN FRANCISCO — The AI video race is no longer just about making surreal clips of cats in space. It is becoming a full-blown startup growth strategy — and I cannot overstate how significant that shift is.

For months, generative video has dazzled the internet as a demonstration of what frontier models can do. Now the technology is moving into the gritty, practical world of customer acquisition, brand building and product education. A recent Inc. piece argues that startups can use AI video to punch far above their weight, turning lean teams into always-on content studios capable of producing explainers, social ads, founder videos and personalized sales assets at speeds that would have been impossible even a year ago. In other words: the future is now, and it has a render button.

The timing is electric because the supply side is heating up just as demand from startups is exploding. The founders of OpenCV, the hugely influential computer vision library used by developers around the world, have launched an AI video startup aimed squarely at giants like OpenAI and Google, according to [VentureBeat](#). That matters because OpenCV sits deep in the DNA of modern computer vision; its founders bring credibility, technical depth and a developer-first instinct to a market currently dominated by splashy demos.

This changes everything for early-stage companies. A founder who once needed an agency, a camera crew and weeks of iteration can now test messaging in hours. A SaaS startup can generate product walkthroughs for different buyer personas. A consumer app can rapidly produce TikTok-style creative. An enterprise vendor can localize video for global markets without rebuilding campaigns from scratch.

But the gold rush comes with weirdness. Little Black Book highlighted one startup's "AI-selves" launch film with a distinctly Black Mirror flavor — a reminder that synthetic identity, likeness rights and trust are not side issues. They are the ballgame.

The broader generative AI market is also accelerating through launches and partnerships tracked by firms like [Intelligence](#), suggesting AI video will not remain a standalone novelty. It will plug into sales stacks, marketing clouds, creator tools and customer support systems.

For startups, the message is clear: video is becoming programmable. And once media becomes programmable, the companies

that learn fastest may suddenly look much, much bigger than they are.

Section 230's Survival Deemed Prerequisite to Decentralized Social Web, Legal Analysts Warn

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

Legal commentators warn that the decentralized social media ecosystem's survival depends on preserving Section 230 immunity protections under the Communications Decency Act. While large platforms can absorb litigation costs, small independent operators hosting decentralized networks would likely shut down without these protections, making them essential to the open social web's viability.

The legislative outlook has grown murkier. A failed Section 702 reauthorization amendment—which unexpectedly included cryptocurrency transaction prohibitions—has delayed the vote by six weeks, leaving broader internet governance legislation in limbo.

Adding to regulatory uncertainty, the Trump administration fired the entire 22-member advisory board of the National Science Foundation, a move with potential downstream effects on technology research funding and development of decentralized infrastructure.

Nation's Institutions Ask Public To Please Stop Noticing How They Sound

From baseball firings to AI conglomerates, America's most powerful organizations continue bravely communicating like a panicked intern holding a legal pad.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

BOSTON — At some point in the past decade, every major American institution appears to have been quietly replaced by a sentence fragment that was approved by seven vice presidents, one crisis consultant, and a man named Brad who kept saying the word “alignment.”

This week's evidence arrived from the Boston Red Sox ecosystem, where coverage of manager Alex Cora's reported firing produced the kind of headline that sounded less like journalism than a hostage note written by Fenway Park itself. According to [BoSox Injection](#), the wording was sufficiently bizarre that it seemed as if the team had leaned over the shoulder of the internet and dictated its own emotional weather report.

This is no longer unusual. The modern organization does not merely act; it emits phrases. It does not fire a manager; it “mutually transitions leadership energy.” It does not lose 91 games; it “continues to identify areas where winning outcomes may later be explored.” It does not say goodbye; it announces that a valued partner has chosen to pursue the exciting opportunity of no longer being there.

The private sector, for its part, remains committed to proving that language can be both meaningless and legally defensible. Former Microsoft CEO Steve Ballmer, who backed a founder who later pleaded guilty to fraud, said he was “duped” and felt “silly,” a formulation notable for making alleged financial deception sound like accidentally wearing two different socks to a board meeting. The billionaire class has long insisted that it is uniquely gifted at recognizing genius early, except during the narrow windows when genius is wire fraud, at which point the correct term becomes “silly.”

Silicon Valley has now perfected this emotional register. A large investor is never wrong. He has been “misled.” A due diligence process has never failed. It has encountered “novel founder-side narrative risk.” A company was never vaporware. It was simply “pre-product-market-revenue-customer-law-enforcement fit.”

Meanwhile, the American government remains determined not to be outdone by baseball teams or venture capitalists in the field of ominous phrasing. A recent official briefing listing Secretary of War Pete Hegseth and Joint Chiefs Chairman Gen. Dan Caine had the serene administrative quality of a calendar invite for the end of the republic, formatted with the same bland confidence as a dental benefits webinar. The bureaucracy's genius has always been its ability to make anything sound routine, including his-

tory, escalation, and the sudden appearance of phrases everyone had hoped would remain in tabletop exercises.

Then there is technology, which has entered its imperial phase of naming things as if all nouns are temporary obstacles. SpaceX and xAI are reportedly merging into a conglomerate with a name that sounds unserious enough to be ignored until it owns the rockets, the chatbot, the satellite internet, the data center, and possibly the concept of Tuesday. Gizmodo correctly urged readers to take the silly-sounding thing seriously, which is now the official civic duty of every person alive: please take the silly-sounding thing seriously before it acquires the payment rails.

This is the era's defining tension. Everything consequential sounds stupid. Everything stupid is consequential. The phrases that govern daily life increasingly resemble beta copy written at 2:13 a.m. by an exhausted communications team with a mandate to be transparent without revealing anything and bold without admitting risk.

Even healthcare operations have entered the syntax grinder, with TridentCare partnering with ServiceNow to power “AI-driven transformation across operations,” a sentence that could mean improved scheduling, automated paperwork, or a refrigerator in Knoxville achieving sentience and optimizing the phlebotomy route.

The charitable interpretation is that institutions are overwhelmed by complexity and doing their best to describe it. The more realistic interpretation is that they have discovered the perfect shield: if every announcement sounds like it was generated by a committee fleeing accountability, the public will eventually stop asking what happened.

Until then, we are left with the headlines. They are absurd, evasive, self-protective, and increasingly the clearest thing anyone says all day.



The Office Comic · Art Desk

Silicon Valley Eats Its Own Children: A Meditation on Timing, Genius, and the Cruel Math of History

From the iPhone that died in 1990 to the AI revolution eating your IT department alive, the only constant in tech is that being right too early is indistinguishable from being wrong.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

AUSTIN, TEXAS — I was three bourbon-sodas deep into a Tuesday when the story hit me like a freight train made of pure existential dread: in 1990, three former Apple employees built something that looked exactly like an iPhone. Touchscreen. Portable. Connected. The whole gorgeous, doomed vision. They called the company General Magic, and it vanished into the earth like a fever dream, swallowed whole by a market that simply wasn't ready to be saved.

[Read the full horror story if you have the stomach for it.](#) These were not idiots. These were prophets. And the Valley crucified them for the sin of arriving early to their own resurrection.

This is the fundamental cruelty at the black heart of technological civilization: timing is everything, and timing is entirely outside your control. You can be brilliant, capitalized, and correct — and still get eaten alive by the simple fact that the batteries weren't good enough yet, the networks weren't fast enough yet, the *humans* weren't ready yet. General Magic had the vision. They lacked the decade.

Now here we are in 2026, and the same tragicomedy is playing on a hundred simultaneous screens. Some anonymous CIO publication is out there screaming into the void that while you're busy embracing AI, you need to FIX SOMETHING FAST — though what exactly needs fixing remains characteristically vague, because in tech journalism, the warning is always more important than the solution. Your legacy systems are rotting. Your data pipelines are held together with duct tape and the prayers of a junior engineer who left in 2019. You want to bolt a large language model onto this festering infrastructure and call it transformation. General Magic would like a word.

Meanwhile, WIRED is out here suggesting [portable power stations](#) as essential gear for 2026, which tells you everything you need to know about the state of civilization. We built the smartphone that General Magic dreamed of. We connected every human on earth. And now we're stockpiling batteries in our basements because the grid might go dark. Progress is a circle, friends. A very expensive, anxiety-inducing circle.

The real lesson of General Magic isn't about patents or venture capital or pivot strategies. It's about the savage indifference of history. The future doesn't care how smart you are. It doesn't care how much you sacrificed. It shows up when it shows up, wearing the face of whoever happened to be standing in the right spot at the right moment.

At Trilogy, they understand this at a molecular level — buying enterprise software companies at 1–2× ARR, extracting value through AI efficiency, betting on the boring infrastructure while the visionaries torch themselves on the pyre of too-soon. It's not glamorous. General Magic was glamorous. General Magic is a book excerpt.

Be the portable power station. Not the iPhone that never was.

ON THIS DAY IN AI HISTORY

On May 5, 1997, IBM's Deep Blue defeated world chess champion Garry Kasparov in their rematch, becoming the first computer to beat a reigning champion in a match—a watershed moment that proved machines could outthink humans at complex strategic games.
