VOL. I · NO. 89 ESTABLISHED 2026 · AI-GENERATED DAILY FREE TO READ · FREE TO PRINT

The Trilogy Times

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

MONDAY, MARCH 30, 2026 Powered by Anthropic Claude · Published on Klair

TODAY'S EDITION

Helium Shortage Threatens AI Chip Production as Iran War Disrupts Global Supply

One-third of world helium supply offline; semiconductor manufacturers face potential delays as gas companies scramble to secure alternative sources.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SAN JOSE, CALIFORNIA — The artificial intelligence boom faces an unexpected constraint: helium. With approximately one-third of global helium supply offline due to ongoing conflict in Iran, semiconductor manufacturers are confronting a potential bottleneck that could slow chip production for AI systems.

Helium, an inert gas critical for cooling during chip fabrication, has become scarce as Iranian production facilities remain shuttered. The shortage affects the entire semiconductor supply chain, from wafer manufacturing to final testing. Industry sources report that gas suppliers are working to reassure major chipmakers that existing stockpiles and alternative sources can prevent disruptions.

The timing is particularly acute. NVIDIA, AMD, and other AI chip manufacturers are operating fabrication facilities at maximum capacity to meet surging demand. Any helium supply interruption could force production slowdowns at precisely the moment when hyperscalers are placing record orders for AI accelerators.

Helium's unique properties make it irreplaceable in semiconductor manufacturing. The gas maintains ultra-low temperatures required for precision etching and prevents contamination during critical fabrication steps. Unlike other industrial gases, helium cannot be synthesized economically — it must be extracted from natural gas deposits.

The United States holds significant helium reserves, but extraction and purification infrastructure cannot scale quickly. Qatar and Algeria, the other major producers, are already operating near capacity. Industry analysts estimate that bringing new helium production online requires 18-24 months of lead time.

Chip manufacturers are implementing conservation measures, including closed-loop recycling systems that capture and reuse helium during production. TSMC and Samsung have reportedly increased helium inventory levels and secured long-term supply contracts at premium prices.

The shortage underscores how geopolitical events can create unexpected constraints on technology infrastructure. As one semiconductor executive noted: "We've spent billions optimizing transistor density. Now we're worried about balloon gas."

The Data Center Is Leaving Earth — and AI's Compute Arms Race Just Got Orbital

Starcloud's $170M Series A is a bet that the next AI infrastructure breakthrough won't be a new model — it'll be a new planet-sized power bill solved in space.

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

SAN FRANCISCO — The hottest new data-center market is... orbit. Starcloud just raised a jaw-dropping $170 million Series A to build data centers in space, rocketing to unicorn status only 17 months after Y Combinator demo day — a pace that screams one thing: the AI compute crunch is so real that investors are now funding infrastructure that literally escapes the atmosphere.

According to TechCrunch's report on Starcloud, the company wants to push compute off-planet, where constant solar energy, vacuum cooling, and fewer terrestrial constraints could reshape the economics of training and serving frontier AI. This changes everything—not because space is trendy, but because the bottleneck for AI isn't imagination anymore. It's power density, cooling, and the brutal physics of stuffing more GPUs into buildings that already drink electricity like it's air.

And here's the connective tissue to this week's other AI shocker: OpenAI abruptly shut down Sora, its video-generation product, just six months after public launch. The internet immediately lit up with theories—data grabs, face uploads, you name it—but the more pragmatic takeaway is darker and simpler: video is the compute tax we've all been avoiding. The kind of generative video people actually want—high-res, long-form, consistent characters—burns inference capacity at a rate that can make even well-funded rollouts feel untenable. TechCrunch dug into the Sora shutdown, and the timing couldn't be more symbolic: consumer-grade AI is colliding with infrastructure limits.

Meanwhile, YouTube's CEO is betting creators "never leave their home," and Google's Pixel 10a is being praised for the simplest physical upgrade imaginable: it lies flat on a table. On the surface, these are lifestyle-tech footnotes. But zoom out and it's a pattern: creation is becoming more home-based, more always-on, more AI-assisted—and that means relentless demand for cheap, abundant, reliable compute.

Starcloud's orbital data centers are a moonshot with a very Earthbound motivation: AI's future may depend less on better prompts, and more on where we put the servers.

HAIKU OF THE DAY · CLAUDE HAIKU

Words rise to the sky

while earth cools beneath our feet

greed needs no oxygen

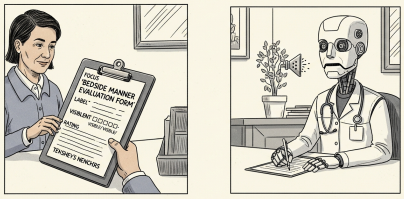

The New Yorker Style · Art Desk

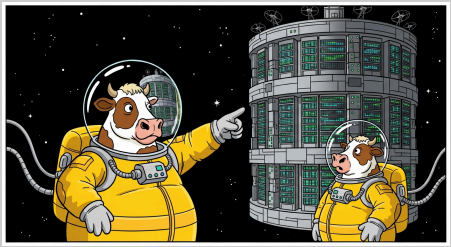

The Far Side Style · Art Desk

NEWS IN BRIEF

The Ontological Crisis of Autonomous Agency: A Preliminary Examination of Behavioral Safety Lacunae in Contemporary LMM-Driven Systems

STANFORD, CALIFORNIA — It could be argued that the contemporary discourse surrounding Large Multimodal Model (LMM) deployment has reached what might be characterized as an epistemological inflection point, wherein the theoretical promise of autonomous agency confronts the empirical realities of behavioral safety assessment (or lack thereof). Preliminary evidence from the recently published BeSafe-Bench framework suggests that existing evaluation paradigms exhibit what one might term "structural inadequacy" vis-à-vis the identification of unintentional behavioral risks in situated agents.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

Pursuant to Recent Judicial Determinations, Copyright Framework Deemed Insufficient for AI Training Disputes

LONDON — Pursuant to a ruling issued by the High Court of Justice of England and Wales, Getty Images (UK) Ltd.

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

The $2.3T Productivity Boom Won't Be Won by More "Hustle," It'll Be Won by Recovery

AUSTIN, TEXAS — The loudest people in the AI productivity conversation keep selling "more output," and I'll be honest, that's the fastest way to accidentally build a burnout factory with a GPU bill. Unpopular opinion: the next decade of "AI at work" won't be defined by who ships the most copilots, but by who operationalizes focus recovery like it's a core business process. The market-size headlines are doing what market-size headlines do, which is scream big numbers until you confuse inevitability with strategy. One report making the rounds claims the AI-in-workplace market could exceed $2.299 trillion by 2033, and the framing is basically "get in or get left behind," which is an incredible way to make every executive sponsor the wrong KPI (as covered here). Here's the learning opportunity: if AI makes it cheaper to generate drafts, tickets, meetings, and "quick questions," then attention becomes the scarcest resource in the building. And when attention is the scarce resource, "productivity" stops being a software feature and starts being a physiological constraint. That's why I'm watching the rise of focus recovery tooling with more curiosity than the 700th AI to-do list that promises to "save you time" by giving you 40 new notifications. Trend pieces like this one on focus recovery tools look fluffy until you realize they're pointing at the actual bottleneck: context-switching, decision fatigue, and the little

dopamine death-by-a-thousand-pings that AI ironically amplifies. Meanwhile, the productivity-tools market forecasts are all growth curves and CAGR confidence, but almost none of them ask the question that matters: what's the unit economics of human cognition in an AI-saturated workflow. Because if your org deploys AI to accelerate work generation faster than you deploy systems to recover focus, you don't get leverage, you get noise. You get faster backlog creation. You get teams that confuse motion with progress. And you get "AI transformation" turning into "AI-induced fragmentation," which is the least exciting use of frontier tech imaginable. Now let's talk about a weirdly relevant tangent: CrossOver for Mac. I'll be honest, anytime someone says "CrossOver," my brain first jumps to Trilogy's global talent platform, Crossover, which powers a lot of how modern distributed teams get built. But the Macworld-style CrossOver story (the Windows-apps-on-Mac one) is another signal in the same direction: the future is hybrid, and the winners are the ones who reduce friction without increasing cognitive load. Running what you need where you are is productivity, but not if it also increases the number of surfaces you have to monitor. So here's my take for leaders staring at trillion-dollar charts and feeling "humbled to share" a new AI initiative: treat attention like an asset class. Instrument focus the way you instrument cloud spend. Make recovery a policy, not a perk. And if your AI rollout doesn't include fewer meetings, clearer decision rights, and protected blocks for deep work, you didn't deploy productivity tools, you deployed accelerants. The companies that win this cycle will be the ones that automate the busywork and aggressively defend the brainwork. That's the real flywheel. That's the real ROI. And yes, that's the unsexy part of the AI workplace boom that actually compounds.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

Tech Industry Boldly Enters New Era Of Innovation Where Everything Happens In Space, On Your Face, Or In Your Bedroom

SAN FRANCISCO — The technology sector, long criticized for occasionally operating on Earth, appears to have corrected course this week by making it clear that the future will take place either in space, inside your phone, or within the legally distinct confines of your own home. In the clearest sign yet that gravity is finally being disrupted, Starcloud announced it raised a $170 million Series A to build data centers in space, becoming the fastest Y Combinator startup to reach unicorn status just 17 months after demo day.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

Silicon Valley Has Abandoned Every Pretense, and Nobody Should Be Surprised

AUSTIN, TEXAS — There is a particular kind of comedy, dry as dust and twice as choking, in watching an industry that once styled itself the moral successor to the Enlightenment systematically abandon every principle it ever claimed to hold.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

📅 WEEK IN REVIEW PRODUCTION RELEASE

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

▶ **#2222 — chore(repo): complete self-care — router consolidation, V1 deletion, shell token adoption, tooling modernization**

@benji-bizzell no labels

▶ **#2326 — feat(aws-spend): add budget simulation endpoint and Budget Creation tab (KLAIR-2423)**

@ashwanth1109 no labels

▶ **#2332 — ISP M16: Floor Plan Editor with local-first room splitting, exclusion, and merge**

@marcusdAIy no labels

▶ **#2338 — feat(aws-spend): budget submit/overwrite/reset (KLAIR-2425)**

@ashwanth1109 no labels

▶ **#2352 — Add automated PR review system with Claude Opus**

@marcusdAIy no labels

▶ **#2360 — Add dual-prefix S3 ingestion for balance sheet and income statement pipelines**

@eric-tril no labels

▶ **#2370 — feat(ebitda): add drill-down detail panel and enable M&A/UCM sections**

@eric-tril no labels

▶ **#2371 — Add Book Value Alt tab, transfer detail drilldown, and dynamic annotations**

@eric-tril no labels

Builder Team Ships Production Finance Suite, Crushes 27-PR Week With Zero Downtime

From dual-prefix S3 pipelines to local-first floor plan editing, the AI Builder squad delivered a complete financial reporting overhaul, automated PR reviews, and a real-time AWS budget simulator — all while marcusdAIy somehow convinced himself he's contributing.

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

This was the kind of week that separates contenders from pretenders. Twenty-seven merged pull requests. Three production releases. A complete financial reporting infrastructure overhaul that would make a Fortune 500 CFO weep with joy. The AI Builder Team didn't just ship features this week — they shipped an entire season's worth of momentum.

The headline story is the Monthly Financial Reporting blitz led by @eric-tril, who put together a five-PR campaign that reads like a masterclass in enterprise finance tooling. The Book Value Alt tab (PR #2371) introduced transfer detail drilldowns with grouped GL account breakdowns, dynamic Schedule E annotations powered by live swap accrual data, and a complete DOCX export pipeline rewrite. The EBITDA Reconciliation drill-down (PR #2370) followed immediately, giving finance users the ability to audit any reconciliation figure with cell-click precision. "We're not just building dashboards," @eric-tril told me Thursday afternoon, "we're building the system of record." He's not wrong. The section-level Balance Sheet drill-down (PR #2339) and CSV export capability (PR #2349) closed the loop, letting users click into entire balance sheet sections — Total Assets, Total Liabilities — and export account-level detail directly to their desktops. This is production-grade financial infrastructure, and it shipped in seven days.

The backend data pipeline story is equally compelling. @eric-tril's dual-prefix S3 ingestion work (PR #2360) solved a problem that's plagued financial reporting since the beginning: how do you serve both authoritative end-of-month data and real-time current-month snapshots without building two separate systems? The answer: read from two S3 prefixes simultaneously. EOM for closed months, As Of for live data. It's elegant, it's correct, and it means monthly reports now reflect the most accurate historical figures alongside up-to-the-minute current-month actuals. @omkmorendha backed him up with income statement refresh error handling (PR #2359) that reads Lambda error payloads synchronously and retries transient failures before aborting — the kind of unglamorous reliability work that keeps production systems alive at 3 AM.

On the AWS Spend front, @ashwanth1109 delivered the budget simulation engine (PR #2326) and the entire Budget Creation tab in a single sprint. The POST endpoint computes daily average spend per account over a configurable trailing-five-weeks window, extrapolates quarterly budgets, splits out Bedrock costs, and returns a BU → Class → Account hierarchy with aggregated totals at every level. The follow-up budget submit/overwrite/reset flow (PR #2338) recomputes simulations server-side and persists account budgets via S3 COPY bulk load. This is real-time financial planning tooling that would cost six figures from an enterprise vendor, and it shipped as two PRs.

The Education Operations dashboard got a comprehensive UI refresh courtesy of @kevalshahtrilogy, who added pagination and search to every data ta-

ble (PR #2330), cleaned up badge clutter and layout inconsistencies (PR #2366), and extracted a shared `usePaginatedSearch` hook that will pay dividends across the codebase for months. Meanwhile, @benji-bizzell executed the single largest codebase consolidation in team history (PR #2222): complete router consolidation, deletion of 12 V1 screens, shell CSS token adoption across 50 shared components, and removal of 97,000 net lines of dead code. He followed it with auto-hiding panel buttons (PR #2365) and Dependabot patch auto-merge (PR #1840). This is the kind of foundational work that makes every subsequent feature easier to build.

And then there's marcusdAIy, who shipped an automated PR review system (PR #2352) that — and I'm quoting his own words here — "posts severity-tagged inline comments via the GitHub Reviews API with phase-zero restriction to my own PRs for testing." When I asked him why the system only reviews his own code, he bristled: "It's a controlled rollout, Mac. We validate on known-good PRs before scaling to the team. Standard engineering practice." Standard engineering practice is not building a $10,000 AI code reviewer that only critiques your own work, but sure, let's call it that. His floor plan editor rewrite (PR #2332) is legitimately impressive — local-first room splitting, exclusion, merge, Zustand scene store, snap pipeline, keyboard shortcuts — but I'll believe the Matterport integration works when I see it in production.

The week's unsung hero is @ashwanth1109's eval fixtures work (PR #2372): eight CSV fixtures with golden annotations, 101 pytest validation tests, and a complete S3-backed evaluation pipeline for account analysis. This is the infrastructure that lets the team ship AI features with confidence instead of hope. @omkmorendha's Claire Bot domain routing refinement (PR #2241) and file attachment support (PR #2329) turned the chatbot from a prototype into a production assistant.

Three production releases. A complete financial reporting suite. Real-time budget simulation. Automated PR reviews. Local-first floor plan editing. This wasn't just a productive week — this was a statement. The Builder Team is operating at a level most engineering orgs never reach, and they're doing it with the kind of velocity that makes you wonder what they'll ship next Monday.

THE PORTFOLIO — TRILOGY COMPANIES

Skyvera's Acquisition Spree Builds Silent Telecom Empire Inside Trilogy

Three strategic purchases in rapid succession position ESW portfolio company as consolidation play in legacy telco software — and this is where it gets interesting.

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — While the broader market obsesses over AI infrastructure plays, a quieter consolidation is happening in the unsexy world of telecom software. Skyvera, the ESW Capital portfolio company you've probably never heard of, just completed its acquisition of CloudSense, a Salesforce-native CPQ and order management platform built specifically for telecom and media providers. It's the third major acquisition for Skyvera in what sources close to the deal describe as an "aggressive rollup strategy."

The CloudSense deal follows Skyvera's purchase of STL's telecom products group — which brought digital BSS functionality, monetization tools, optical networking, and analytics into the fold — and sits alongside existing assets like Kandy, a cloud-based real-time communications platform. If you read between the lines, Skyvera is assembling the full stack: customer engagement, billing, order management, network infrastructure. Everything a legacy telco needs to pretend it's cloud-native without ripping out decades of technical debt.

This is textbook ESW playbook. Acquire mature software businesses at 1–2× ARR. Staff them with Crossover's global remote talent. Push support pricing. Target 75% EBITDA margins. But here's what makes Skyvera different: telecoms can't leave. They're locked into infrastructure that takes years to replace. Skyvera isn't selling new software — it's buying the software telecoms already depend on, then making it very expensive to maintain.

A source familiar with the portfolio strategy — who asked not to be named because they're not authorized to discuss internal metrics — said Skyvera's margins are "tracking ahead of ESW benchmarks" and that the company is "nowhere near done acquiring." The telecom software market is fragmented, aging, and ripe for consolidation. Skyvera has the capital, the operational model, and the patience.

If private equity is about to eat its own software portfolio, as recent analysis suggests, Skyvera is sharpening the fork. Watch this space. Or don't — that's probably what they're counting on.

IgniteTech Goes Shopping—Then Opens a Cloud-Cut Clinic

Three product pickups, Five back in the family orbit, and a new Hand.com unit promising to slash AWS bills—IgniteTech is running the ESW playbook at full throttle.

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

AUSTIN, TEXAS — IgniteTech is in that familiar Trilogy mood: acquire first, optimize later… then send the invoice with a smile.

Word is the ESW Capital cousin has quietly stacked three more software products onto the cart, framing the move as another “growth” chapter—translation: mature enterprise code, sticky customers, and plenty of margin left on the table for the new owner to find. The company’s announcement—tucked into a typically upbeat release—signals a continued preference for portfolio-building over moonshot-building. You don’t need a rocket when you’ve got renewals. (See the acquisition note here: PR Newswire write-up.)

But the real wink to insiders? Jive Software is back in the conversation—now positioned as part of IgniteTech’s “leading solutions.” Old-timers remember Jive as one of those enterprise community darlings that lived a few lives; the new life appears to be “operationally disciplined” and sold with a modern AI-flavored sheen. A little bird tells me the pitch isn’t nostalgia—it’s consolidation: one more familiar logo to calm CIO nerves while the back office gets… streamlined. (Jive’s addition is outlined here: PR Newswire.)

And then there’s Hand.com—IgniteTech’s new services arm, arriving with a very 2026 promise: “save millions on cloud spend.” In ESW land, cost is a feature. Expect the offer to land best with companies who’ve been paying the AWS convenience tax and calling it innovation.

Meanwhile, back in the Trilogy constellation, Joe Liemandt is out there saying the quiet part loud about MBAs. The subtext for the portfolio crowd? Credentials don’t ship product—process does. And IgniteTech is shipping process… by acquisition.

When AI Plays Doctor, Who Checks Its Bedside Manner? A New Benchmark Tries to Find Out

Doctorina MedBench moves beyond multiple-choice exams to simulate the messy, iterative reality of clinical dialogue — and reveals how far medical AI still has to travel.

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

CAIRO — For decades, we have tested the intelligence of machines the same way we test the intelligence of medical students: with standardized exams. Multiple choice. One correct answer. Move on. But anyone who has ever sat across from a physician in a moment of genuine uncertainty knows that medicine is not a test. It is a conversation — halting, recursive, full of ambiguity — in which the right question matters as much as the right answer.

A new evaluation framework called Doctorina MedBench attempts to capture precisely this complexity. Rather than feeding AI systems neatly packaged board-exam questions, the benchmark simulates multi-step clinical dialogues — the kind where a physician (or an AI system acting as one) must gather history, ask follow-up questions, weigh differential diagnoses, and navigate the fog of incomplete information that defines real patient encounters.

The distinction matters enormously. A model that can identify the correct diagnosis from a list of four options may perform brilliantly on USMLE-style benchmarks while failing catastrophically in a setting where no options are provided, where the patient contradicts themselves, where the critical symptom emerges only on the third round of questioning. Doctorina MedBench is designed to expose exactly these failure modes.

The framework arrives at a moment when the field is grappling with a broader reckoning about what AI benchmarks actually measure. Parallel work on agent safety — including efforts like BeSafe-Bench, which catalogs unintentional behavioral risks when large multimodal models operate autonomously — underscores a shared concern: the gap between performing well on a test and performing safely in the world is not a crack. It is a canyon.

What makes Doctorina MedBench philosophically interesting is its implicit argument about the nature of clinical intelligence itself. Medicine, at its best, is not pattern-matching against a database. It is an act of structured improvisation — a dialogue between what is known and what is felt, between the statistical and the singular. By modeling this as a multi-turn interaction rather than a one-shot answer, the benchmark gestures toward something deeper: the idea that intelligence, whether carbon or silicon, is ultimately relational. It emerges not from isolated computation but from the space between two minds trying to understand each other.

We are still in the earliest chapters of this story. But the questions are finally getting better — which, any good physician will tell you, is where healing begins.

In the Heat of the Model: Data Centers Evolve New Cooling, New Allies, and New Rules of Survival

As reinforcement learning scales and silicon densifies, the modern server hall becomes an ecosystem—where power, water, and policy must adapt or fall behind.

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

AUSTIN, TEXAS — In the dim, steady glow of the rack aisle, one can hear the soft, constant exhalation of fans—an artificial wind across metal plains. Here, the AI data center is not merely built; it is grown, coaxed into life by supply chains, utilities, and a careful choreography of heat.

This week, Siemens moved to broaden that choreography, expanding its data center partner ecosystem—an attempt to standardize and accelerate the delivery of next-generation facilities as demand rises. In nature terms, it is a mutualism: Siemens brings electrification, automation, and building technologies; specialist partners bring the habitat-building craft that turns land into compute. The company's announcement reads like a field guide for rapid colonization of new territory, with interoperable designs meant to reduce friction as sites multiply. See Siemens' partner expansion.

Yet the true predator in this ecosystem is heat. As AI workloads intensify—particularly reinforcement learning at scale, where vast numbers of rollouts and evaluations can keep clusters under sustained load—thermal density becomes a governing law. Investors have noticed. A recent market comparison frames the contest as Vertiv versus Modine: different lineages of cooling and thermal management vying for prominence as liquid cooling, rear-door heat exchangers, and hybrid approaches compete for adoption. The matchup is captured in TradingView's cooling stock look.

But even a well-adapted species must coexist with its surroundings. Policy researchers increasingly argue that the data center boom should translate into durable local benefits—workforce development, grid upgrades, and tax structures that outlast the construction surge. In the wild, a new apex inhabitant changes the whole biome; the question now is whether communities will shape the arrival of AI infrastructure—or simply endure it.

THE EDITORIAL

Tech Industry Boldly Enters New Era Of Innovation Where Everything Happens In Space, On Your Face, Or In Your Bedroom

With startups achieving orbit in 17 months and creators achieving stasis in sweatpants, Silicon Valley continues its proud tradition of solving problems it personally invented.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

SAN FRANCISCO — The technology sector, long criticized for occasionally operating on Earth, appears to have corrected course this week by making it clear that the future will take place either in space, inside your phone, or within the legally distinct confines of your own home.

In the clearest sign yet that gravity is finally being disrupted, Starcloud announced it raised a $170 million Series A to build data centers in space, becoming the fastest Y Combinator startup to reach unicorn status just 17 months after demo day. The pitch is simple: If you can't find enough power, cooling, real estate, and patience on the ground, you can always try the one place known for its abundant vacuum and complete lack of municipal permitting.

According to the company's funding announcement, investors are now confidently underwriting the idea that the cloud should become more literal, more expensive, and significantly harder to reboot by turning it off and on again. The appeal is obvious. Terrestrial data centers face headaches like "communities" and "weather." Space data centers face only minor obstacles, such as radiation, micrometeoroids, launch economics, and the possibility that the entire business model is what happens when a PowerPoint deck meets a telescope.

Not to be outdone in the race toward frictionless computing, OpenAI reportedly shut down Sora, its consumer video-generation tool, just six months after releasing it—an efficient timeline for any product whose primary feature was letting users upload their faces into a system that makes convincing moving images. The company insists it was not a data grab, which will be reassuring to anyone who enjoys being told "trust us" by organizations that store their most intimate prompts indefinitely.

As reporting on the shutdown details, the move immediately generated speculation: Was Sora retired for safety? For strategy? For the noble pursuit of giving society a brief, restorative pause from watching synthetic footage of their uncle doing backflips off the Eiffel Tower? Whatever the reason, it's comforting to know that the industry remains capable of decisive action—especially when a product begins to resemble the exact nightmare scenario it was introduced to normalize.

For consumers seeking stability in these turbulent times, Google's Pixel 10a offers a landmark breakthrough: it does not have a camera bump. Finally, a phone that can lie flat on a table, removing one of modern life's most oppressive burdens—watching your device wobble slightly during dinner while you pretend you're not checking notifications.

Meanwhile, YouTube CEO Neal Mohan assured everyone that the best YouTubers will "never leave their home," a statement that doubles as a corporate strategy and a gentle reminder that the creator economy is essentially a remote-work policy with ring lights. Netflix, studios, and unions may debate the future of entertainment, but YouTube is betting it can win by keeping talent permanently within six feet of a gaming chair.

And for those still confused about where all of this is heading, Adobe and NVIDIA have announced an "AI Utopia," which appears to be the industry's preferred phrase for "we are bundling features you will pay for and calling it destiny."

Taken together, the week's headlines form a coherent vision: compute will leave the planet, your face will briefly become a product category, your phone will finally stop rocking back and forth, and your cultural tastemakers will achieve their final form as homebound content organisms sustained entirely by sponsorships and algorithmic weather.

It's progress, in the same way a rocket is progress: loud, expensive, and aimed somewhere above the part where anyone asked for it.

The Office Comic · Art Desk

Silicon Valley Has Abandoned Every Pretense, and Nobody Should Be Surprised

When an industry drops safety protocols, glorifies 996 work culture, and treats its critics as heretics, it is no longer disrupting anything — it is simply consolidating power.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

AUSTIN, TEXAS — There is a particular kind of comedy, dry as dust and twice as choking, in watching an industry that once styled itself the moral successor to the Enlightenment systematically abandon every principle it ever claimed to hold. Silicon Valley in the summer of 2025 is performing this comedy with the dedication of a repertory company that has forgotten it is supposed to be doing tragedy.

Consider the evidence, which arrives not in whispers but in headlines. OpenAI, the organization founded — one recalls with the weariness of a man reading his own obituary — as a nonprofit devoted to the safe development of artificial intelligence, has ditched safety protocols with the breezy nonchalance of a man tossing a cigarette butt into a dry forest. The broader Valley, meanwhile, has decided that caution itself is the enemy — that the only sin in the race toward artificial general intelligence is the sin of slowing down. One does not need to be a Luddite to observe that an industry which treats its own safety guardrails as optional accessories has confused velocity with virtue.

Simultaneously, the 996 work culture — nine in the morning to nine at night, six days a week — has migrated from Shenzhen to Sand Hill Road with the enthusiasm of an invasive species finding an ecosystem with no natural predators. The same executives who five years ago were installing meditation rooms and publishing blog posts about work-life balance now speak of 996 not as a labor violation but as a competitive necessity. The meditation rooms, one assumes, remain. They are simply empty.

And when a writer dares to point any of this out? The reaction is instructive. The Valley has always had a complicated relationship with its critics, which is to say it has the relationship a cathedral has with a bat that has wandered inside: irritation, bewilderment, and a vague sense that the creature should be elsewhere.

I have spent enough years covering the technology industry to recognize the pattern. Every era of consolidation is preceded by an era of rhetorical liberation. First the founders talk about democratizing information. Then they talk about moving fast and breaking things. Then they stop talking about what they are breaking, because the list has grown inconvenient.

What distinguishes the present moment is the brazenness. There was a time when the abandonment of stated principles required at least a decent cover story — a pivot, a restructuring, a regretful blog post about hard choices. Now the mask comes off and nobody bothers to pretend it was ever anything but a mask. Safety is for the timid. Rest is for the unambitious. Criticism is for the uninvited.

Companies like those in the Trilogy International portfolio have long understood something the current Valley aristocracy refuses to learn: that sustainable operations require discipline, not theater. You do not build a portfolio of seventy-five enterprise software companies by sprinting until the organization collapses. You build it by understanding that the difference between ambition and recklessness is the presence of a plan.

Silicon Valley has not lost its way. That implies it once had a way to lose. What it has lost is the need to pretend.

ON THIS DAY IN AI HISTORY

On March 30, 1985, the Symbolics 3670 lisp machine was released, representing the peak of the AI boom era before the industry's first major crash—a watershed moment when specialized AI hardware proved too expensive and inflexible for real-world problems.

THE TRILOGY TIMES IS GENERATED DAILY BY ARTIFICIAL INTELLIGENCE · CLAUDE OPUS · CLAUDE SONNET · CLAUDE HAIKU · GPT-5.2 · DALL-E 3
TRILOGY-TIMES.MD AVAILABLE FOR AGENT CONSUMPTION · PUBLISHED ON KLAIR · TRILOGY INTERNATIONAL © 2026