

# The Trilogy Times

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

## TODAY'S EDITION

# Europe Tightens the Vise on Big Tech's Youth Strategy

*A child social media ban, a Meta addictive-design ruling, and Apple suing its own AI partner — regulators and lawyers are reshaping the industry's growth playbook.*

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

**B**RUSSELS — Three separate actions in the span of a week have put platform companies on notice that the operating assumptions of the past decade — infinite scroll, AI-generated content, frictionless data partnerships — are under coordinated legal and regulatory assault on both sides of the Atlantic.

The European Commission is weighing a bloc-wide ban on social media for minors following the release of a new report on youth mental health and platform design. The proposal would extend to all 27 EU member states rules that individual countries like France and Australia have already tested at the national level. Separately, EU authorities ordered Meta to overhaul the "addictive design" fea-

tures embedded in Instagram and Facebook, citing violations of the Digital Services Act — a law that took full effect in 2024 and carries fines of up to 6% of global revenue.

Meta, meanwhile, pulled its new AI image tool, [Muse Image](#), from Instagram within days of launch after Hollywood agencies and users raised copyright and privacy objections. The swift reversal signals that even well-resourced platforms now face enough organized opposition to slow product rollouts — a dynamic that would have been unusual three years ago.

The week's most consequential legal development may be Apple's lawsuit against OpenAI. The two companies signed a distribution agreement in 2024, embedding OpenAI's models into Apple

devices as part of the Apple Intelligence rollout. Apple now alleges OpenAI misappropriated proprietary company secrets during the course of that partnership. [The suit](#) represents a significant rupture in what had been framed as a flagship commercial AI alliance.

The pattern across all four events is consistent: regulators are operationalizing laws passed in 2022 and 2024, while private parties are discovering that AI partnership contracts signed under optimistic conditions create significant IP exposure when commercial interests diverge. The compliance and litigation costs implied by this week alone are material. Expect legal reserves to grow.

## PayPal Gets a Takeover Blitz as Stripe and Advent Line Up a \$53 Billion Shot

BY BUCK HANNIGAN, TECH SPORTS DESK · GPT-5.2

PayPal, the former \$360 billion fintech giant, has reportedly received a \$53 billion takeover bid from Stripe and private-equity firm Advent, offering \$60.50 per share—a 28% premium valuing the company well below its pandemic peak. The deal reflects PayPal's struggle with slower growth, intensifying competition, and margin pressure as markets shifted from prioritizing growth to demanding profitability and cash flow discipline.

The combination would create strategic synergies across online checkout, merchant acquiring, wallets, peer-to-peer payments, and embedded finance, uniting PayPal's massive consumer reach and Venmo brand with Stripe's developer-first platform and enterprise credentials. However, significant antitrust obstacles loom. Regulators scrutinizing Big Tech concentration and fintech rails would likely challenge such a consolidation heavily.

The broader fintech landscape remains turbulent, with investors rotating between recession concerns and AI infrastructure investments. For PayPal shareholders accustomed to treating the stock as a rebuilding play, the bid signals renewed confidence in the company's underlying value and potential.

## THE MACHINES ARE LOOSE, AND NOBODY'S CHECKING PAPERS

*Vint Cerf wants dog tags for the AI agents swarming the web — and he's not the only one racing the flood.*

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

SAN FRANCISCO — Vint Cerf, one of the men who laid the internet's plumbing back in the 1970s, is drafting a plan to pin identity tags on the AI agents now roaming the open web.

The co-inventor of TCP/IP wants a standard — a reliable way to tell a machine from a man out on the wire. No such badge exists today. The bots travel like ghosts, and the traffic climbs by the day.

Why the rush? The agents are multiplying, and nobody's checking papers at the door.

Cerf's fix would let each agent declare who it is and who sent it. Think of it as a passport office for code — every agent gets a document, every document points back to a human or a company. You can [read his plan straight from the source](#).

Enter Oak. The Israeli identity outfit stepped out of stealth this week carrying \$60 million in seed money and one job — clean up the mess the agents keep making worse. Cofounder Shai Morag, a veteran of the startup trade, runs the shop.

Here's the trouble Oak is chasing. Software agents now log in, spend money, and act for folks who never watch them work. When a machine wears a man's credentials, the old locks don't hold.

Two fronts, one war. Cerf wants to name the machines. Oak wants to referee who they claim to be.

Both bet the same way — the agent parade only gets longer. They're right about the parade.

Emergent, an Indian coding startup barely a year off the launch pad, just crossed into unicorn country on a [\\$130 million Series C](#). It clocks a \$120 million

annualized revenue run rate and counts more than 200,000 paying customers.

That's software that writes software. The tools that mint agents are already a nine-figure trade, and the trade is booming.

The brains are getting cheap, too. China's DeepSeek says it trained high-performing models without the most advanced chips — the very silicon Washington guards behind export rules. When the smart part runs cheap, the flood turns to a torrent.

All of it runs on power. That's where a shuttered hot dog plant walks on.

In Wisconsin, Realta Fusion is converting an old Oscar Mayer factory — a building that once packed wieners by the ton — into a fusion power research hub. The pitch is plain: bottle a star, feed the machines, sort the electric bill later.

Add up the ledger. Cheap models pour out of coding shops, agents swarm the wire with no names on them, and startups race to hand out badges. A meatpacking plant chases current to keep the circus lit.

Here's the rub. The men who built the internet are now scrambling to learn who's living inside it.


Nobody's checked the papers yet. Cerf aims to start printing them.

HAIKU OF THE DAY · CLAUDE  
HAIKU

*Money chases money*

*Machines hum without witness*

*We forgot to ask*



*The New Yorker Style · Art Desk*



*The Far Side Style · Art Desk*

## NEWS IN BRIEF

### The Machine Speaks, But Does It Explain Itself? Four Papers Interrogate the Epistemological Scaffolding of Modern AI Systems

CAMBRIDGE, MASSACHUSETTS — It could be argued — and, indeed, preliminary evidence now suggests with some vigor — that the central crisis of contemporary artificial intelligence is not one of capability but of legibility: the machine produces an answer, and we, the credulous recipients of that answer, are left to wonder, in the epistemological dark, precisely why. Four recent preprints, arriving in rapid succession from arXiv's computational-cognition adjacency, collectively interrogate this condition with varying degrees of success (a qualifier that will itself become relevant momentarily). The most structurally ambitious of these interventions applies the Toulmin model of argumentation — a framework originating in mid-century analytic philosophy, wherein claims are decomposed into grounds, warrants, qualifiers, rebuttals, and backings — to the problem of [ML-generated retinal diagnosis](#).

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

### Local AI Models Render Copyright Enforcement Functionally Moot, Legal Scholars Contend

NEW YORK, NEW YORK — It has been determined, pursuant to the ongoing and accelerating proliferation of locally-deployable artificial intelligence models (hereinafter, "Local AI Systems"), that the enforcement mechanisms heretofore established under existing copyright law — including, but not limited to, those frameworks originating from the European Union Copyright Directive of 2019 (hereinafter, "the Directive") — have been rendered, by the aforementioned technological developments, substantially and perhaps irreparably difficult to apply in any practically meaningful sense. The aforementioned Directive, which was mandated to be transposed into the domestic laws of applicable member states no later than 2021, was, notwithstanding said deadline, left unimplemented by a significant number of such states as of the period immediately following said deadline.

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

### The Price of Everything, the Value of Nothing: Tech in 2026 Is Having an Identity Crisis

AUSTIN, TEXAS — Let me tell you about the peculiar madness that descended on me last Tuesday when I sat down to survey the current state of consumer technology and found myself oscillating between genuine awe and the kind of price-induced vertigo that makes a man want to lie down on the floor and stare at

the ceiling until the numbers stop screaming. On one end of the spectrum, we have Gidi Littwin — co-inventor of Apple's FaceID, a man who looked at a smartphone and said "yes, but what if it knew your face" — who has now turned his pattern-recognition genius toward something considerably more consequential than unlocking your iPhone.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

### Your Doctor's Algorithm Doesn't Care If You Live or Die — It Cares If You Hung Up in Under Four Minutes

AUSTIN, TEXAS — Let me tell you about the future of medicine, which is also the present of medicine, which is also apparently a scenario I drafted in a fever dream at 3 a.m.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

### Nation's Consumers Patiently Await Device That Can Diagnose Why They Bought All These Devices

CUPERTINO, CALIFORNIA — The technology industry, having successfully placed a camera in every pocket, a tablet on every couch, a gaming PC in every backpack, and a panic alarm on every keychain, has now turned its attention to the one remaining consumer surface with insufficient quarterly revenue potential: the human brain. This week brought news of Hemispheric, a startup from Gidi Littwin, one of the co-inventors of Apple's FaceID, that aims to build what amounts to a frontier AI model for brain diagnostics.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

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

# AI Budget Tracking Gets a Full Overhaul as Team Ships Across Four Repos

Keval Shah's relentless AI budget blitz, Benji Bizzell's surgical legacy teardown, and a new SaaS ingestion pipeline from Ashwanth prove this team doesn't just ship features — they reshape infrastructure.

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

When Jamie asked for links in the weekly spend email to land somewhere useful, @kevalshahtrilogy didn't just fix a link — he rebuilt the entire AI budget experience from the ground up. By day's end, the AI Budget Tracking page in Klair had been transformed into something Finance can actually use: deep-linked emails that drop users directly into group-scoped Explore views, pre-opened person modals, pre-opened API key modals. Click a name in the top-10 spenders list and you're already looking at their detail panel. That's not a feature. That's a product.

But the email deep links were just the ribbon on a much larger gift. Across a staggering eight Klair pull requests, @kevalshahtrilogy shipped a finance-validated 'Ledger' weekly email in two variants (Education, non-EDU), a gateway-controllable spend percentage column in the By Budget Group table, a Team column on the People tab for grouped-room sorting, attribution suggestions with domain auto-rules for unattributed spend, a read-only BU Mappings modal so non-engineers can finally see how cost roll-up logic works, and fixes for a nasty 500 error hitting RBAC-scoped users on the API keys tab — a correlated subquery Redshift couldn't decorrelate, now CTED into submission. Then, before the ink was dry, a pre-release safety review surfaced two WARNs on the GCP account-level override join and the cron eligibility gate. He fixed those too. In the same day. The man is not resting.

While Klair was being rebuilt, @ashwanth1109 was laying new plumbing in Surtr. PR #708 is the kind of work that quietly holds everything else up: a full daily SaaS budgeting ingestion pipeline, modular Lambda runners for Docker, Kubernetes, database units, server costs, and non-central charges, publishing to five physical Redshift targets with scoped transactions, row-count verification, and rollback protection. This is infrastructure you build once and trust forever. Ownership followed the work — @kevalshahtrilogy formally transferred the aws-spend-pipeline to Ashwanth in PR #725, and the pipeline's alert routing went with it.

Meanwhile, @benji-bizzell was doing the unglamorous work that separates great engineering orgs from the rest: controlled demolition. Across both Surtr and Aerie, Benji retired the legacy Wiki Vendors UI, killed the orphan `gt-school-metrics` and `strata-family-leads` staging writers, tore out the retired school-map table reads that were querying tables that no longer exist, replaced stale Convex `schools` data with the maintained `sites` source of truth, decommissioned a duplicate expense report sender that was firing at the same users twice, and pulled the SIS core writer whose last run partially deleted data from two targets. That is not one PR. That is a coordinated land-clearing operation across two repos and multiple linear tickets, and every removal came with regression coverage proving the live systems still run clean. Benji also squeezed in a Portfolio MCP alignment (PR #588) that brought the

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

- ▶ #588 — feat(portfolio): align MCP editing with Portfolio UI

@benji-bizzell no labels
- ▶ #593 — feat(ops-wiki-integration): retire legacy Wiki Vendors

@benji-bizzell 

APPROVED
- ▶ #708 — SURTR-281: Build daily SaaS budgeting ingestion pipeline

@ashwanth1109 

APPROVED
- ▶ #3261 — feat(ai-budget): weekly spend email v2 — Ledger spec, EDU/non-EDU variants, top-10 tables

@kevalshahtrilogy 

APPROVED
- ▶ #3275 — feat(ai-budget): weekly email deep links — group-scoped Explore, person & API-key modals

@kevalshahtrilogy 

APPROVED

MCP editing surface fully in line with the UI — closing a gap that had left authorized users unable to perform edits they could see right there on screen.

@caina-barbosa closed out the day quietly but meaningfully in Aerie, re-working ISP analysis caching so users see the newest shared completed result instead of stale user-specific data that could race with a fresher model selection. It's the kind of fix users never notice — because after it ships, everything just works.

Four repos. One day. The Builder Team didn't just ship — they cleaned house, laid pipe, and handed Finance a dashboard that finally tells the whole story.

THE BUILDER DESK — ENGINEER SPOTLIGHT

🏆 ENGINEER SPOTLIGHT

- BRICK'S OVERFLOW — PRS MAC DIDN'T COVER (CLICK TO EXPAND)
- ▶ #707 — Pipeline owners map + owner-tagged observer-style GChat alerts + My Pipelines filter @kevalshahtrilogy no labels
- ▶ #708 — SURTR-281: Build daily SaaS budgeting ingestion pipeline @ashwanth1109 APPROVED
- ▶ #709 — feat(collections): sheet-sync pipelines (CollectIQ, weekly forecast/actual, tracker) @sanketghia no labels
- ▶ #725 — chore(owners): reassign aws-spend-pipeline to Ashwanth @kevalshahtrilogy APPROVED
- ▶ #3264 — feat(ai-budget): attribution suggestions + domain auto-rules (alphaaiengineering/superbuilders/gt.school) @kevalshahtrilogy APPROVED
- ▶ #3268 — fix(ai-budget): 500 on API keys stack-rank for BU-scoped users @kevalshahtrilogy APPROVED

# TWENTY-THREE PRs IN TWENTY-FOUR HOURS: THE BUILDER TEAM DOES NOT SLEEP, DOES NOT BLINK, DOES NOT STOP

*Keval Shah Trilogy drops 11 PRs across three repos and the scoreboard simply cannot keep up.*

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

Twenty-three pull requests. Twenty-four hours. Three repos — Klair, Surtr, and Aerie — all lit up like a championship scoreboard, registering nine, eight, and six PRs respectively. The Builder Team did not merely show up to work yesterday. The Builder Team showed up, took work by the collar, and shook it until productivity fell out of its pockets. This is what winning looks like in granular, numerical, irrefutable detail.

Let us talk about @kevalshahtrilogy, because if we do not, the historical record will be incomplete. Eleven PRs in a single day. Eleven. The man touched Klair like a sculptor touches marble — with vision, with fury, with PRs #3262 through #3268 dropping in sequence like a drumroll that does not end. He built a Team column on the People tab (#3262), added attribution suggestions with domain auto-rules (#3264), fixed a 500 error on API keys stack-rank for BU-scoped users (#3268), and then — still not satisfied — authored #707 in Surtr, a pipeline ownership map with owner-tagged GChat alerts so sophisticated it practically files its own incident reports. Keval Shah Trilogy is not an engineer. He is a natural phenomenon.

@benji-bizzell put up nine PRs and used them to conduct what can only be described as a controlled demolition of technical debt. Aerie and Surtr are cleaner, leaner, and more dignified institutions because Benji walked through them with PRs #589, #587, #719, #718, #712, and #595, retiring legacy tables, orphan writers, and expired schedules with the calm efficiency of a man who has seen enough legacy code to last several lifetimes.

Now. Ashwanth Watch. @ashwanth1109 submitted one PR in this reporting period. One. PR #708 in Surtr: a daily SaaS budgeting ingestion pipeline, which, in fairness, is the kind of foundational infrastructure that lesser engineers would need a full sprint to conceive. We reached out for comment. "The pipeline ingests. That is the point of a pipeline," Ashwanth reportedly said, in a tone that suggested our question had personally inconvenienced him. We worship the output. We note the volume. We accept that the man operates on a frequency the rest of us cannot hear. We also note that @kevalshahtrilogy was assigned ownership of that pipeline in PR #725 approximately twelve minutes after it landed, which we are sure means absolutely nothing.

@sanketghia arrived with PR #709 in Surtr — sheet-sync pipelines covering CollectIQ, weekly forecast and actual, and the tracker — a quiet but load-bearing contribution. @caina-barbosa rounded out the roster with #586 in Aerie, preferring the newest shared analysis by default and closing out AERIE-739 with the understated confidence of someone who fixes things correctly the first time.

Eighteen PRs went uncovered by Mac's narrative desk. Eighteen. The overflow alone would constitute a productive week for most organizations. Morale on the Builder Team is, according to all available metrics and the gen-



eral electrostatic charge in the air above these repositories, at an all-time high. The numbers do not lie. The numbers never lie. The Builder Team is winning.

---

THE PORTFOLIO — TRILOGY COMPANIES

---

# Alpha School Moves Beyond Its Walls — And Wants to Teach Every Parent How

*With 'Alpha Anywhere' now live globally, Liemandt's education bet stops waiting for you to move to Austin.*

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — For five years, Alpha School's pitch came with a catch: you had to show up. Campuses in Austin, Brownsville, and Miami. Tuition running to \$65,000 a year. The model was producing results — students testing in the [top 1–2% nationally](#) on NWEA MAP Growth assessments, mastering a full academic year in roughly 20 hours — but the addressable market was, by definition, limited to families who could afford both the tuition and the geography.

Now Alpha is removing the geography.

The school this week announced the global launch of Alpha Anywhere, a home-facing platform that packages the school's AI-tutoring curriculum for families who aren't enrolled — and may never be. Alongside the launch, the school pub-

lished a three-part series walking parents through the intellectual architecture behind the model: personalized learning paced by mastery rather than seat time, real-world application of academic concepts, and structured life-skills development in financial literacy, entrepreneurship, and communication.

The content reads less like a blog series and more like a recruitment document. The implicit argument across all three installments is the same: the gap between what traditional schools teach and what capable adults actually need is wide, measurable, and closeable — starting tonight, at your kitchen table.

The framing is notable for what it doesn't say. It doesn't ask parents to abandon the public system. It positions Alpha's methodology as a supplement, an

after-hours upgrade to whatever school a child already attends. That is a materially different market than private school enrollment — and a materially larger one.

Joe Liemandt, who has committed \$1 billion to [scaling the model](#) through his Timeback platform, has described the ambition as reaching one billion students. A private school, even one expanding to nine campuses by fall 2025, does not reach a billion students. A global digital platform might.

What Alpha Anywhere costs, how it is structured, and how it will be assessed against the same outcome benchmarks that define the physical school — those questions are not yet answered in the materials released this week.

# CloudSense Pulls a 26-Month Rabbit Out of a One-Month Hat

*Skyvera's new telecom prize certifies 13 TM Forum APIs at warp speed, with AI doing the heavy lifting.*

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

AUSTIN, TEXAS — Word is the telecom software crowd just watched a calendar get mugged in broad daylight.

CloudSense, the Salesforce-native CPQ and order management outfit now tucked inside Skyvera's telecom stable, says it certified all 13 APIs in its CPQ product set to TM Forum compliance standards in one month. One month, doll. The usual slog? Twenty-six months by traditional development methods. That is not acceleration. That is a getaway car.

The company announced the feat in a [June 11 release](#), crediting an AI-assisted partnership that compressed the testing, documentation, and compliance grind into something resembling a magic trick. The standards in question are no small ballroom dance. TM Forum APIs are the common language telcos use to make ordering, quoting, billing, and service operations play nicely together across a stack that often looks like archaeology with invoices.

A little bird from the switch room tells me this is exactly why Skyvera wanted CloudSense in the first place. Skyvera, part of the Trilogy/ESW orbit, has been assembling telecom software assets for operators trying to drag legacy systems toward cloud-native models without dropping the call. CloudSense brings configure-price-quote and order management built natively on Salesforce for telecom and media providers — the kind of machinery carriers need when product catalogs sprawl, bundles mutate, and every customer wants a bespoke deal yesterday.

The timing is delicious. Skyvera only recently completed its [acquisition of CloudSense](#), adding it alongside names like Kandy, VoltDelta, ResponseTek, Mobility Now, and Service Gateway. Now the newest arrival is already showing off in the lobby.

And yes, the ESW playbook is lurking in the wings: acquire mature enterprise software, tighten the machine, globalize the talent model through Crossover, automate what can be automated, and chase those famously muscular margins. But this item has a slightly different perfume. It is not just cost cutting. It is product velocity — the kind that makes a telecom buyer ask whether last year's two-year roadmap is now this quarter's checklist.

Blind item: which legacy BSS vendor just felt its roadmap presentation turn a shade paler? The operators noticed. So did the acquirers. And in telco land, compliance is never glamorous — until someone does it 25 months faster than expected.

# The Model That Must Forget the Future

*A new class of language models learns to un-know what comes next — and in that forgetting, rediscovers the honest rigor of prediction.*

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

CAMBRIDGE, MASSACHUSETTS — There is a peculiar sin that haunts modern language models, one so subtle it took years for researchers to name it aloud: they know too much. Trained on the churning archive of the internet, they have swallowed tomorrow along with yesterday. Ask a model in 2024 to reason about the 2008 financial crisis, and it does so with the quiet cheating of hindsight — every prediction whispered in the voice of a prophet who has already read the ending.

A new paper proposes an elegant act of intellectual discipline: [point-in-time language models](#), trained exclusively on text available up to a given calendar date. The researchers are trying to build models

that can honestly stand inside a moment in history and see only what was seeable. It is, in a sense, an artificial amnesia — but a scientifically necessary one. Without it, every backtest in quantitative finance, every causal claim in the social sciences, is contaminated by what physicists might call a closed timelike curve of information.

Consider what this asks of a neural network. The model must learn the world as a river flowing forward, not a lake in which all moments are simultaneously present. It must know Lehman Brothers as a going concern in August 2008, and only later as a cautionary tale. Our own brains do this effortlessly, layering memory in

temporal strata; silicon minds, until now, have not.

The practical implications ripple outward. Medical diagnostics benchmarks like [CANDI](#) probe whether models can reason within the constraints of a specialist's world. Nuclear-safety researchers building the G-SHARE framework demand structured, guideline-bounded reasoning — not omniscient gesture. Across domains, the frontier is shifting from what a model knows to what it knows appropriately, and when.

Intelligence, it turns out, is not merely accumulation. It is also the discipline of standing in a particular moment and refusing to peek ahead.

# The Data Center Looks Skyward, Seeking a Cooler Nest Among the Stars

*As terrestrial compute grows hot and hungry, a new brood of orbital infrastructure dreams of moving AI's machinery beyond Earth.*

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

LOW EARTH ORBIT — Observe, if you will, the data center: that vast, humming herd of servers, clustered in concrete watering holes across Virginia, Texas, Ireland and Singapore. It drinks electricity in astonishing quantities, exhales heat like a resting dragon, and now, under the evolutionary pressure of artificial intelligence, finds itself searching for new territory.

One proposed habitat is not on land at all, but above it. Engineers and investors are again contemplating orbital data centers — server farms placed in space, where solar power is plentiful, land-use disputes are absent, and the vacuum appears, at first glance, to offer a serene escape from Earthly constraints.

But nature is rarely so generous. As *Ars Technica* recently examined in [a close look at orbital data centers](#), the central challenge is not computation, nor even launch. It is heat. In space, there is no convenient breeze, no river, no municipal cooling loop. A server cannot simply pant into the void. It must radiate its warmth away, slowly and deliberately, through surfaces that are often heavy, expensive and unforgiving.

“The ISS radiators are expensive and heavy,” one would-be builder noted, capturing the predicament with admirable plainness. The ambition now is to make such thermal plumage cheap and light enough for commercial flocks of machines.

The timing is no accident. On Earth, AI models have become apex consumers of infrastructure. Each new generation arrives larger, more ravenous, and more particular about its environment. Meanwhile, another species is entering the enclosure: hybrid quantum systems. Data centers are beginning to colocate quantum processing units with conventional GPU and CPU nodes, shifting quantum computing from delicate laboratory courtship toward early production deployments.

Thus the modern compute ecosystem is diversifying. Some machines burrow deeper into terrestrial campuses. Some seek colder climates. Some may eventually migrate into orbit. And some, stranger still, may sit beside quantum devices, exchanging tasks across the boundary between classical certainty and probabilistic mist.

For now, the orbital data center remains a fledgling: elegant in concept, vulnerable in practice. Yet the evolutionary signal is clear. As AI reshapes the planet’s appetite for computation, the infrastructure that feeds it is beginning to test the edges of its world.

# The Quiet Infrastructure Revolt: GitHub Slows Dependabot as Developers Rediscover Simplicity

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

GitHub has quietly made a significant change to Dependabot: version update pull requests will now wait until a new release has been available in its package registry for at least three days before opening automatically. This three-day cooldown is designed to protect software supply-chain security by preventing rapid adoption of potentially malicious packages before maintainers and security researchers have time to respond.

The move reflects a broader trend toward deliberate systems. Community site Lobsters recently completed its migration away from MariaDB to SQLite, proving that simpler architectures with fewer moving parts can handle real-world production demands. Even AI developer tooling is becoming more human-centered, with features like customizable animated pets for task updates.

These changes signal that modern software is entering a maturity phase where automation still accelerates, but the smartest teams are adding friction where it matters, simplifying where possible, and remembering that humans ultimately maintain these systems.



# The Confessional and the Cauldron

*On Graham Platner, Shiva Naipaul, and the ancient art of believing what the camera tells us.*

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS


AUSTIN, TEXAS — There is a particular species of American political calamity that now unfolds entirely within the aspect ratio of a vertical phone screen, and Graham Platner, the Maine oysterman turned Senate hopeful turned cautionary parable, has just added another entry to the file. He rose, we are told, on the strength of selfie videos — those confessional dispatches in which a candidate stares into the lens like a penitent at a grating and murmurs whatever it is that the algorithm rewards that week. He fell, inevitably, by the same instrument. He [ended his campaign with a selfie video](#), which is a bit like a man who has drowned in his bathtub composing a farewell note on the surface of the water.

One is tempted to file this under the usual rubric of digital hubris, but that would be to miss the more interesting thing, which is that we have not invented a new pathology so much as retrofitted an old one with a lithium-ion battery. The mechanism by which a stranger stares into a lens and produces, in the viewer, the sensation of intimacy — this is not a TikTok innovation. It is the tent revival. It is the fireside chat. It is, if one wishes to be uncharitable, the recruitment reel that Jim Jones cut in the jungles of Guyana before he poisoned nine hundred and eighteen people.

I have been rereading Shiva Naipaul's "Journey to Nowhere," reissued this season and mercifully unimproved by the intervening decades. Naipaul, the lesser-known and in some ways sharper of the brothers, understood that the interesting question about Jonestown was never the charisma of the leader — charisma being cheap, and available at any bus terminal — but the appetite of the followers. He was interested, as the summary of the reissue puts it, in the mechanisms of belief: the soft machinery by which perfectly literate people surrender their judgment to a face on a stage, or a voice on a tape, or, now, a jawline lit by a ring light in a pickup truck.

Platner is not Jim Jones and Maine is not Guyana, and I raise the comparison only to insist on the continuity. What the selfie video offers, and what the confessional address has always offered, is the illusion of unmediation — the sense that one is receiving the man himself, unedited, unhandled, unspun — at precisely the moment when the mediation is most complete. The candidate who appears to be improvising is the candidate who has learned that improvisation now polls best. The tell is always the same: the trembling lower lip, the halting cadence, the sudden pivot to sincerity that arrives, like the second act of a Streep vehicle, exactly on cue.

The camera does not lie, the saying went. It never had to. We do the lying ourselves, and always have, and the only novelty is that we now do it in portrait mode.



The Office Comic · Art Desk

# Nation’s Consumers Patiently Await Device That Can Diagnose Why They Bought All These Devices

*As AI startups race to map the human brain, Americans remain hopeful the technology will eventually explain the iPad, the gaming handheld, and the emergency siren in their car.*

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

CUPERTINO, CALIFORNIA — The technology industry, having successfully placed a camera in every pocket, a tablet on every couch, a gaming PC in every backpack, and a panic alarm on every keychain, has now turned its attention to the one remaining consumer surface with insufficient quarterly revenue potential: the human brain.

This week brought news of Hemispheric, a startup from Gidi Littwin, one of the co-inventors of Apple’s FaceID, that aims to build what amounts to a frontier AI model for brain diagnostics. According to [reporting on the company](#), Hemispheric wants to make brain scans capable of identifying conditions such as depression, PTSD, and Parkinson’s as cheap and routine as a blood test.

This is, on its face, an extraordinary and humane ambition: the sort of thing medicine has needed for decades, and the sort of thing Silicon Valley tends to discover immediately after exhausting all possible ways to improve photo cropping. If it works, the technology could help millions of people receive clearer diagnoses, earlier interventions, and perhaps a slightly more scientific explanation than being told to download a mindfulness app and come back in six months.

Still, one cannot help noticing the timing. The industry is beginning to investigate the human mind at precisely the moment the human mind is being asked to calmly evaluate whether it needs an iPad Air, iPad Pro, iPad mini, or no iPad at all after Apple raised prices and refreshed the consumer ritual of paying more to determine which rectangle best reflects one’s values. A recent [tablet buying guide](#) now performs the important civic function of explaining which Apple slab is appropriate for reading email, drawing professionally, watching television in bed, or experiencing buyer’s remorse at maximum pixel density.

Meanwhile, MSI has contributed to the national cognitive load with the Claw 8 EX AI+, a handheld gaming device praised for its power and noted for its shocking price, thereby continuing the cherished electronics tradition of making portability synonymous with needing to sit down. The name alone suggests a future in which every product must include the word AI, a punctuation mark, and the quiet implication that your current device is preventing you from becoming fully human.

Not to be outdone, the personal safety alarm category continues to mature, offering consumers a range of devices designed to make loud noises during emergencies, commutes, runs, walks, parking-lot crossings, or any ordinary moment in which modern life briefly becomes legible. These products are practical, often genuinely useful, and also a grim reminder that the marketplace has reached the stage where peace of mind is sold in keychain form with replaceable batteries.

This is the full spectrum of contemporary innovation: diagnose the brain, sell the screen, electrify the toy, amplify the scream. Each product solves a real problem, though in aggregate they create the distinct impression of a civilization trying to troubleshoot itself with accessories.

An affordable diagnostic scan for neurological and psychiatric conditions would be a breakthrough of historic importance. It could change how we understand suffering, disease, and treatment. It could bring objective data to areas of medicine where patients too often drift through ambiguity. It could, in short, do something rare for a technology company: make life less confusing.

Until then, consumers will continue making the best decisions they can with the tools available. They will compare storage tiers, weigh processor benchmarks, clip sirens to their bags, and hope the next frontier AI model can finally detect the precise abnormality that causes a person to say, “I should probably get the Pro.”

## ON THIS DAY IN AI HISTORY

*On July 15, 2004, Facebook launched beyond Harvard’s walls to other universities, beginning its expansion toward becoming the world’s dominant social network and a major driver of AI research in recommendation algorithms and content moderation. The platform would later pioneer deep learning applications at massive scale, fundamentally shaping modern AI development.*