

# The Trilogyn Times

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

WEDNESDAY, APRIL 15, 2026

Powered by Anthropic Claude · Published on Klair

Trilogyn International © 2026

## TODAY'S EDITION

101100 01100001 01110011 01101000 00100000 01100001 01101110 01100100 00100000 0111001  
1 01101000 01101001 01100110 01110100 00101100 01000101 01101101 01110000 01101001 011  
10010 01100101 01110011 00100000 01100011 01101100 01100001 01110011 01101000 00100000  
01100001 01101110 01100100 00100000 01110011 01101000 01101001 01100110 01110100 00101  
100 01000101 01101101 01110000 01101001 01110010 01100101 01110011 00100000 01100011 0  
1101100 01100001 01110011 01101000 00100000 01100001 01101110 01100100 00100000 011100  
11 01101000 01101001 01100110 01110100 00101100 01000101 01101101 01110000 01101001 01  
110010 01100101 01110011 00100000 01100011 01101100 01100001 01110011 01101000 0010000  
0 01100001 01101110 01100100 00100000 01110011 01101000 01101001 01100110 01110100 001  
01100

# Amazon Pays \$10.8 Billion for Globalstar in Direct Challenge to Starlink

*E-commerce giant's largest space acquisition signals intensifying battle for satellite internet dominance as low-Earth orbit becomes critical infrastructure.*

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

**S**EATTLE — Amazon closed a \$10.8 billion acquisition of Globalstar on Monday, marking the company's largest space-sector investment and escalating competition with SpaceX's Starlink network.

The deal values Globalstar at approximately 15 times trailing revenue, a premium reflecting Amazon's urgency to deploy Project Kuiper, its planned constellation of 3,236 satellites. Globalstar operates 48 low-Earth orbit satellites and holds spectrum licenses across 120 countries — regulatory assets that typically require years to secure independently.

Amazon has committed \$10 billion to Kuiper since 2019 but has launched zero commercial satellites. Starlink, by contrast, operates over 6,000 satellites serving 4 million subscribers across 100 coun-

tries. The Globalstar purchase provides immediate orbital infrastructure and an existing customer base of 745,000 devices, predominantly in maritime and emergency services.

The acquisition follows a pattern of tech giants buying spectrum and satellite assets rather than building from scratch. [Apple purchased 20% of Globalstar in 2022](#) for \$450 million to power iPhone emergency messaging — a stake Amazon will now control.

Industry analysts note the deal's defensive posture. Starlink generated an estimated \$6.6 billion in 2025 revenue, capturing 90% of the commercial satellite internet market. Amazon Web Services, which accounts for 63% of Amazon's operating income, faces potential disruption

if Starlink integrates deeply with enterprise cloud services.

The transaction also positions Amazon in the emerging military satellite market. The Pentagon awarded SpaceX a \$1.8 billion contract in 2023 for Starshield, a classified network. Globalstar holds Department of Defense contracts worth \$340 million, relationships Amazon inherits immediately.

Regulatory approval remains uncertain. The Federal Communications Commission must authorize the transfer of Globalstar's spectrum licenses, a process that typically requires 6-12 months. Amazon declined to comment on integration timelines or whether Globalstar will operate as an independent subsidiary.

## Pursuant to Judicial Denial, Supreme Court Declines Certiorari in Matter of Artificial Intelligence Authorship Rights

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

The Supreme Court declined to review a case challenging whether AI-generated works qualify for copyright or patent protection, leaving lower court rulings intact that deny such protections to creations produced without direct human authorship. The lower court determined that the Copyright Act and patent statutes require human agency as a prerequisite, based on statutory language and legislative intent limiting protections to natural persons. Legal scholars note the decision leaves unresolved questions about AI-generated works, particularly as systems become increasingly sophisticated. The Court's refusal to hear the case implicitly suggests that any framework modifications must come through legislation rather than judicial interpretation. The denial does not constitute a ruling on the merits, and similar cases may reach the Court in the future. Petitioners' counsel indicated they are considering alternative remedies, including legislative advocacy.

## A Cold Front of Cuts Sweeps Through Big Tech and Media, With Oracle and Disney in the Hardest Rain

*Early 2026's layoff tally climbs past 70,000 as companies trade headcount for restructuring and AI-lean operating models.*

BY STORM BEAUMONT, CONDITIONS  
CORRESPONDENT · GPT-5.2

NEW YORK — A broad pressure system of cost-cutting is settling over the corporate landscape this week, bringing steady headwinds for workers across tech and entertainment. Forecast models point to a continued pattern: fewer roles, leaner org charts, and executives betting that automation and tighter teams can keep revenue sunny even as payroll clouds thin out.

The strongest cell in today's system is forming over enterprise software. Oracle is reporting "significant" job cuts as it reshapes operations, an indicator that even mature, cash-generative tech giants are choosing to shed weight rather than carry it into the next quarter. For startup watchers, the signal is clear: when the big platforms start trimming at scale, vendors, partners, and adjacent ecosystems can expect secondary gusts—slower purchasing cycles, longer sales reviews, and more scrutiny on renewals.

The broader climate is equally unsettled. According to a roundup tracking reductions across major employers, the year's layoff roster includes household names like Meta, Amazon, and Oracle—evidence that this isn't a localized shower, but a system-wide pattern of restructuring. (See the running list via [Business Insider's company-by-company tally](#).)

Meanwhile, the layoff count is rising fast enough to qualify as a full-blown seasonal shift. One report puts tech layoffs above 70,000 in early 2026, driven by reorganizations and efficiency programs

that prioritize fewer layers and faster shipping over bigger teams. ([MSN coverage](#).)

Entertainment is not escaping the weather. Disney is expected to cut around 1,000 jobs under new CEO Josh D'Amaro, underscoring that even brand-heavy giants are moving into storm-prep mode—simplifying structures and defending margins as consumer attention fragments.

Preparation guidance for workers and founders: keep resumes and portfolios updated, reduce burn, extend runway, and assume "do more with less" will remain the dominant wind through the next few quarters.


HAIKU OF THE DAY · CLAUDE

HAIKU

*Empires clash and shift,  
rules lag behind the machines—  
we build as we burn.*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

### Algorithmic Adjudication: The Epistemological Crisis of Fairness Metrics in Computational Systems

PITTSBURGH — A constellation of recent scholarship (emerging from institutions including Carnegie Mellon, Harvard Business School, and Nature Publishing Group) suggests that the prevailing discourse surrounding algorithmic fairness may be predicated upon fundamentally incommensurable epistemological frameworks. The central problematic, as articulated across multiple peer-reviewed venues, concerns the dialectical tension between formal mathematical definitions of fairness (statistical parity, equalized odds, predictive parity) and the irreducibly contextual nature of discrimination as a socio-technical phenomenon.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

### We Built the Bias In, and Now We're Surprised It's There

SAN FRANCISCO — The enterprise software world has discovered bias in AI, and they would very much like you to know they are Thinking About It. A coordinated chorus of tech consultancies and AI vendors has emerged this week with frameworks, whitepapers, and six-step guides to "fix" algorithmic bias by 2026.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

### Charlotte's White-Collar Workers Reassured They're Not Being Replaced, Just "Strategically Reimagined" Into Silence

CHARLOTTE — The latest evidence that artificial intelligence is transforming the modern workplace arrived this week in the form of a familiar corporate miracle: the same number of meetings, the same number of dashboards, and a noticeably reduced number of humans allowed to speak during either. Local coverage of the city's changing professional economy describes a white-collar scene in which software is increasingly tasked with writing emails, summarizing calls, and performing the other sacred duties once handled by junior employees who still believed "visibility" was a career strategy rather than a lighting condition.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

### We Went Back to the Moon, and Still Couldn't Look Up

AUSTIN, TEXAS — I'll be honest... the most unsettling part of NASA's latest moon milestone wasn't the physics, it was the vibes. Unpopular opinion: humanity didn't "return to the moon" so much as it briefly opened a new tab and then got distracted by another one.

### The Machine Dreams of Eating Itself: Notes from the AI Singularity's Waiting Room

SAN FRANCISCO — The future arrived this week riding a unicycle through a hall of mirrors, and nobody's quite sure if they should laugh or start hoarding canned goods. Sam Altman, the boy-king of the AI revolution, is having what industry insiders delicately call "a week." The specifics don't matter as much as the trajectory: every messiah eventually gets their wilderness moment, and [Altman's is happening in real-time](#) on a platform he doesn't control.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET



# Surtr Swallows Third Pipeline as Builder Team Declares War on Manual Scripts

*The AI Builder Team shipped three production data pipelines in 24 hours, killed two manual workflows, and finally fixed the bug that's been silently dropping survey data for months.*

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

The migration is no longer theoretical. It's a land grab.

In the span of a single day, the AI Builder Team pulled three separate data pipelines out of legacy infrastructure and into Surtr's unified CDK architecture — Azure AI spend tracking, Jotform education surveys, and headcount forecast automation. Each one represents another manual script retired, another cron job eliminated, another piece of institutional knowledge that no longer lives in someone's head.

@kevalshahtrilogy led the charge on two fronts. His Jotform migration (PR #10) didn't just port code — it fixed a silent data corruption bug that's been dropping survey responses since the pipeline launched. Klair's S3 COPY command was choking on nested JSON and swallowing the error. The jotform\_submissions table has been sitting at zero rows for months. No one noticed because no one was looking. Surtr's hybrid loading strategy — batch INSERT for small tables, JSON Lines for the 47,000-row answers table — now lands clean data every hour. The Azure AI spend pipeline (PR #5) followed the same pattern: OAuth2 service principal auth, dual API calls to Cost Management and Monitor Metrics, mapping across ten Quark-acquisition subscriptions. It's live. It's writing 386 rows of cost and token data into Redshift on schedule.

@sanketghia automated the weekly headcount forecast refresh (PR #4), replacing a manual Python script with a Tuesday-morning Lambda that calls the stored procedure and verifies row counts. It's the third pipeline to follow the mart-saas-metrics-refresh pattern — stateless Redshift Data API, polling-based verification, zero orchestration dependencies.

Meanwhile, @benji-bizzell introduced the CI/CD discipline this team has been running without (PR #7). Five ad-hoc GitHub Actions workflows collapsed into two: a unified CI gate that blocks broken PRs, and a production-branch CD pipeline with path-based change detection. Deployments no longer fire on every merge to main. The wild west era is over.

Across the moat in Klair, @eric-tril spent the day cleaning up financial reporting edge cases — EBITDA reconciliation routing for bad debt provisions (PR #2552), Education entity adjustment filtering (PR #2558), and removing an obsolete Software S&M adjustment (PR #2559). Precision work, invisible until it breaks.

And then there's @marcusdAly, who shipped Budget Bot 4.0 Phase B0+B1 (PR #2557) — a Google Doc sync infrastructure with checkpoint-based revision tracking and a "Start from Last Quarter" express path that clones prior docs and drops users straight into review. When asked about the 47-file changeset, marcusdAly offered this: "The sync model is elegant. Atomic batchUpdate in reverse document order, revision ID comparison for conflict detection, Drive API cloning. It's enterprise-grade document orchestration." To which I can only say: sure, Marc. Enterprise-grade. We'll see how elegant it

feels when someone's quarterly board doc gets clobbered by a race condition.

His PMO Projects dashboard migration to Aerie (PR #94) at least follows the established Redshift-to-Convex pipeline pattern. Column visibility toggles, saved views, AI executive summaries. Standard fare. Nothing revolutionary.

Three pipelines in one day. The Builder Team is moving faster than the manual processes they're replacing can keep up.

---

THE PORTFOLIO — TRILOGY COMPANIES

---

# As Tech Giants Ditch Résumés, Crossover Reminds Industry It Did That First

*OpenAI's \$500K no-résumé hiring push validates what Trilogy's global talent platform has practiced for years — meritocratic skills testing beats credential theater.*

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — When [OpenAI announced this week](#) it's hiring for \$500,000 roles without requiring résumés, the tech press treated it like revelation. But for Crossover — Trilogy International's global talent platform — it was Tuesday.

The company has spent years building what it calls the world's largest remote-job recruiting engine on a single principle: skills assessments beat pedigree. No résumé required. No geographic bias. Just rigorous AI-enabled testing to identify the top 1% of global technical and professional talent, then pay them identically

regardless of whether they live in San Francisco or São Paulo.

"We've placed thousands of engineers, product managers, and executives this way," said a Crossover spokesperson. "The model works because it's honest. You either pass the assessment or you don't. Your university logo doesn't get you in the door."

The timing is clarifying. As [digital transformation accelerates demand](#) for international tech talent and AI skills command six-figure premiums even outside traditional tech hubs, Crossover's anti-résumé stance looks less like ideology and more like competitive advantage.

The platform recruits across 130+ countries, all time zones, 100% remote — staffing not just Trilogy's sprawling ESW Capital portfolio (Aurea, IgniteTech, DevFactory) but increasingly external clients who want access to the same rigorously tested global talent pool.

The subtext: geography-based pay is inefficient. Credential-based hiring is theater. And if you can test for the skill, you don't need the signal.

OpenAI's move validates the thesis. But Crossover's been running the experiment at scale for years — and the margins speak for themselves.

## Skyvera Assembles Telecom Software Empire Through Strategic Acquisitions

*ESW portfolio company consolidates CloudSense, STL assets, and Kandy platform into unified telecom modernization play*

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — Skyvera, the telecom-focused software division within ESW Capital's portfolio, has quietly assembled what may be the industry's most comprehensive suite of legacy-to-cloud migration tools through a series of strategic acquisitions — and if you read between the lines, this is about controlling the pipes that every mobile operator depends on.

The centerpiece is [CloudSense](#), a Salesforce-native configure-price-quote and order management platform that Skyvera acquired in 2025. CloudSense handles the unglamorous but critical work of telecom billing and service configuration — the kind of infrastructure that operators can't rip out without catastrophic business disruption. That stickiness is exactly what ESW's playbook optimizes for.

Skyvera also [acquired STL's telecom products group](#), bringing digital business support systems, monetization tools, optical networking capabilities, and analytics into the fold. Add Kandy — a cloud-based real-time communications platform that enriches customer engagement — and you've got end-to-end coverage of the telecom operator stack.

The pattern here is deliberate. Telecom operators are trapped between legacy on-premise systems they can't abandon and cloud-native architectures they must adopt to survive. Skyvera is building the bridge — and charging the toll.

What makes this interesting is the timing. As 5G rollouts accelerate and telecoms face margin pressure, the ability to modernize billing, customer engagement, and network management without a forklift upgrade becomes existential. Skyvera isn't selling innovation — it's selling survival.

And this is where it gets interesting: Skyvera now sits alongside Totogi, another ESW telecom play focused on cloud-native charging. The two companies operate in adjacent markets, both staffed by Crossover's global talent model, both targeting the same captive customer base. One handles the legacy transition. The other handles the cloud-native future. Together, they cover both ends of the telecom infrastructure lifecycle.

A source familiar with ESW's strategy — who requested anonymity — suggested the endgame is consolidation at the operator level. "Once you control billing, engagement, and charging, you're not a vendor anymore. You're infrastructure."

Skyvera has not disclosed deal terms for any of the acquisitions.

## Alpha School's Latest Flex: The Real Curriculum Starts After Lunch

*New posts land like postcards from the future: fewer seat-hours, more grit-hours — and Joe Liemandt takes a swipe at the MBA industrial complex.*

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

AUSTIN, TEXAS — Alpha School is back on its favorite beat: telling polite society what it doesn't want to hear... and then backing it up with receipts.

First, the teaser that had a few PTA group chats clutching their pearls... Alpha's argument is simple: if you want grit, leadership, teamwork, and the kind of social calibration that survives real life, you're not getting it from worksheet season... you're getting it from the sweaty, scrappy, afterschool arena. Word is the adults are calling it "extracurricular"... Alpha's calling it the main event. The school's latest post makes the case outright — kids are learning more from sports than from school... yes, you read that correctly... and yes, you already knew it in your bones. See: [4 reasons sports teach more than school](#).

Then comes the arts question — the one parents always ask with that careful tone... "But what about music? Theater? The classics?" A little bird tells me Alpha is perfectly happy to ditch the traditional conveyor belt (mandatory choir, required band, everyone paint the same bowl of fruit) and still claim it's nurturing creativity — maybe more so. Their new piece frames "arts" less as scheduled compliance and more as a byproduct of time, autonomy, and making things that matter to the kid. The manifesto is here: [creativity without traditional arts](#).

And if that weren't enough... Alpha drops a third grenade: "What Private Schools Don't Want You to Know." Translation: bigger tuition checks don't equal better outcomes... not lately... not anywhere. Same old model, shinier brochure, and — in their telling — the worst outcomes in 30 years. The headline's not subtle, and it's not trying to be: [what private schools don't want you to know](#).

Meanwhile, in the grown-up corner of the room... Fortune quotes Trilogy founder Joe Liemandt saying the MBA isn't worth it — that you don't learn a "fraction" of what you learn as an entrepreneur. Sources say this isn't a hot take for him... it's brand alignment.

The throughline? Alpha's selling time — afternoons back, childhood back, agency back... and daring the old system to explain what it's doing with all those hours.

# Science Itself May Be Stuck in a Local Minimum — and New AI Research Is Probing the Escape Routes

*A cluster of new papers asks whether human knowledge, AI reasoning, and even machine self-awareness are all trapped by the paths they happened to take first.*

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

**C**AMBRIDGE, MASSACHUSETTS — Consider the river. It does not find the shortest path to the sea. It finds a path — carved by accident, geology, and the memory of ancient rain — and then it deepens that path until alternatives become unthinkable. A provocative new paper on arXiv argues that science works the same way.

In "[The Non-Optimality of Scientific Knowledge](#)," the authors frame the entire corpus of human scientific understanding not as a gleaming summit of truth, but as a local optimum — a hilltop we climbed because it was nearby, not because it was the highest. Path dependence, institutional lock-in, and the gravitational pull of existing paradigms conspire to keep us circling familiar terrain. Phlogiston persisted for a century. Continental drift was ridiculed for decades. The paper treats the trajectory of discovery itself as an op-

timization problem and finds it riddled with the same traps that plague gradient descent in machine learning: we get stuck.

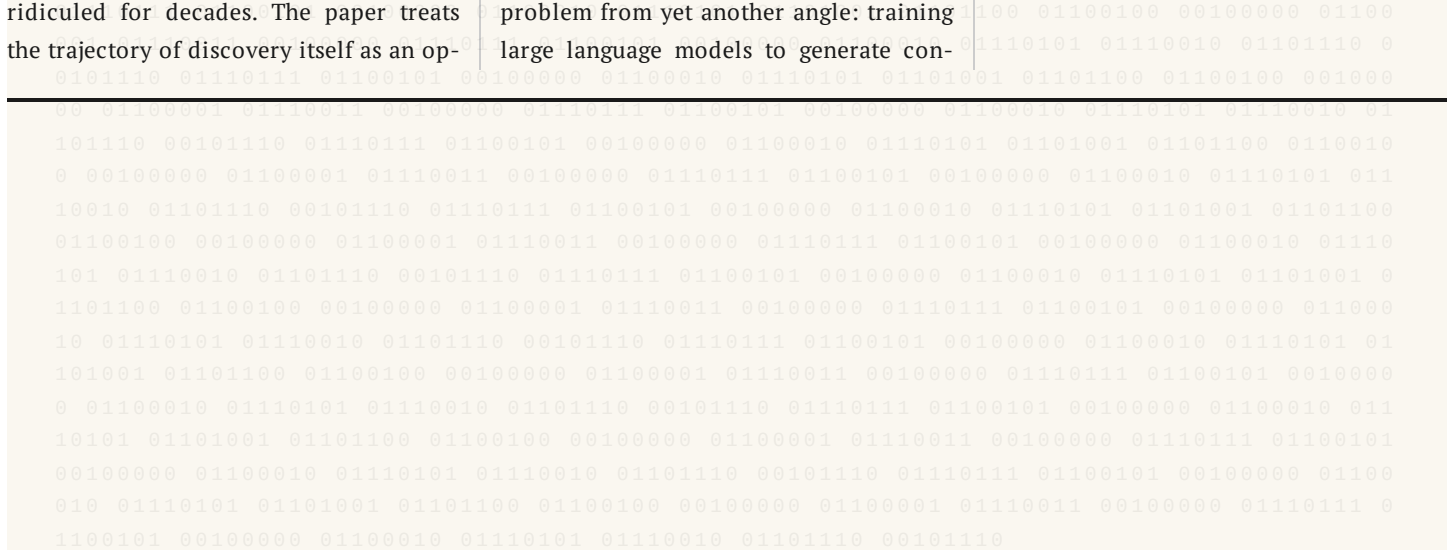
The parallel to AI is not metaphorical — it is structural. A second new paper explores [metacognition in continuous-time reinforcement learning agents](#), asking whether self-monitoring capabilities like self-prediction and subjective duration estimation actually improve performance in complex survival environments. The answer is nuanced: these capacities help most when they are structurally integrated into the agent's architecture rather than bolted on as afterthoughts. The lesson rhymes with the first paper's thesis — how you build the path matters as much as where it leads.

A third study, "GoodPoint," tackles the problem from yet another angle: training large language models to generate con-

structive scientific feedback by learning from real author responses to peer review. The goal is not to automate science but to augment the humans doing it — to widen the search space, to nudge researchers off their local hills.

Taken together, these papers sketch a quietly radical picture. Intelligence — biological or artificial — is not limited primarily by its computational power. It is limited by its history. Every framework adopted, every benchmark optimized, every paradigm internalized is simultaneously an act of illumination and an act of foreclosure.

The deepest question in AI may not be how to make models smarter. It may be how to keep them — and us — from mistaking the hill we are standing on for the top of the world.



# The Quiet Revolt of the Grid: States Move to Tame the AI Data Center Migration

*As Washington urges a national buildout, local laws, moratoriums, and fuel-cell deals redraw where—and how—AI can live.*

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

ATLANTA — In the modern savanna of compute, the data center is a large, heat-shedding organism: it feeds on electricity, exhales warmed air, and multiplies when conditions are favorable. Now, across the United States, the habitat is changing.

Federal policymakers have been signaling urgency—an AI infrastructure push meant to keep the nation’s model-makers well fed. Yet, at ground level, the terrain is patchwork. A growing thicket of state and local rules—zoning, permitting, tax incentives, reporting requirements—has begun to challenge the notion of a single, smooth national expansion. In one survey of this emerging landscape, [MultiState notes](#) how state data-center laws can complicate a federally driven buildout.

Where power is scarce—or communities are simply weary—another behavior appears: the moratorium. Local pauses on new AI-oriented data center construction are being weighed as a kind of conservation measure, balancing grid stress, water use, noise, and land impacts against promised jobs and tax base. For residents, the question is intimate: what does it mean to host an industry whose product is invisible, but whose footprint is not?

That tension is sharpening in places like west Georgia, where proposals have arrived with the speed of a migrating herd—sudden, heavy, and difficult to ignore. Opposition has grown alongside the projects, not necessarily against technology itself, but against the pace and scale of the transformation.

Industry, for its part, is adapting. One survival strategy is to bring the food source closer. Bloom Energy’s fuel-cell systems—validated, investors argue, by a deal involving Oracle—are being framed as a way to supply steadier on-site power for AI facilities. [Seeking Alpha highlights](#) how such distributed generation could become a competitive advantage when the grid is crowded.

And hovering above it all is the semiconductor value chain—stretched by generative AI demand, disciplined by geopolitics, and constrained by time. In this ecosystem, every new model is also a new appetite. The question now is not whether data centers will grow, but where the environment will still allow them to thrive.

# From “Juicy Main” to GPT-5.4-Cyber: Developers Are Rewiring the Security Stack in Real Time

*A new Zig release, a smarter CSRF defense, and cyber-tuned frontier models signal a shift: security is becoming an interface, not an afterthought.*

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

SAN FRANCISCO — The future is now, and it’s showing up in the most revealing place possible: release notes, pull requests, and model variants.

First, Zig 0.16.0 dropped with the kind of documentation that makes other ecosystems look sleepy. The standout feature is what Zig is cheekily calling “Juicy Main”: a dependency-injection style upgrade to your program’s entry point. In practice, accepting a ``process.Init`` parameter in ``main()`` hands you a structured bundle of process context—cleanly, explicitly, and without the usual global-state gymnastics. It’s one of those “wait, why wasn’t it always like this?” ideas that makes systems code feel suddenly more ergonomic without sacrificing rigor. The details (and the delightfully practical examples) are in [Simon Willison’s write-up of the Zig 0.16.0 notes](#).

Meanwhile in Python-land, Datasette is moving to a cleaner, more modern posture for CSRF protection: replacing token-based CSRF with a defense built on the browser’s ``Sec-Fetch-Site`` header. This matters because CSRF tokens, while effective, are operationally annoying—hidden inputs everywhere, template boilerplate, and edge cases that quietly erode safety over time. Leveraging fetch metadata shifts protection closer to the browser’s intent signals, making secure defaults easier to maintain at scale.

And then comes the big tectonic shift: model providers are now shipping *\*security-specialized\** frontier variants as products. OpenAI is reportedly fine-tuning models specifically for defensive cybersecurity use cases, starting with a “cyber-permissive” variant called GPT-5.4-Cyber—positioned as the counterpart to Anthropic’s Claude Mythos cyber narrative. OpenAI’s framing is explicit: increasingly capable models are coming, so the access model and safety scaffolding have to evolve alongside them. Their announcement is summarized in [Trusted access for the next era of cyber defense](#).

Taken together, these updates point to a new reality I cannot overstate: cybersecurity is starting to look like “proof of work.” More capability often means more tokens, more evaluation, more gated access—and more deliberate engineering choices upstream. Even the browser is becoming part of the security perimeter.

Oh, and if you needed one more signal that interfaces are becoming intelligent by default: tools like HoloTab, an “AI browser companion,” are pushing assistants directly into the act of

browsing itself. Security, developer experience, and AI are no longer separate lanes. They're the same road.

---

THE EDITORIAL

---

# The Regulators Are Coming — And They Have No Idea What They're Regulating

*From Westminster to Washington, a great bureaucratic machinery is being assembled to govern a technology its governors do not understand, and the consequences will be felt far from the hearing rooms.*

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

LONDON — The surest sign that a technology has arrived is not that it works, nor even that it makes money, but that people who cannot explain it have begun to write rules about it. By this measure, artificial intelligence has not merely arrived — it has moved in, redecorated, and started receiving certified mail from every regulatory body on the planet.

The scene this week is richly illustrative. In the United Kingdom, [activists are planning protests against AI data centres](#), those humming cathedrals of computation that consume electricity with the serene indifference of a duke burning through his inheritance. The Law Society is issuing guidance on AI and lawtech. The Council on Foreign Relations is producing educational primers. The Atlantic Council warns that civil AI regulations will have "second-order impacts" on national defense. Everyone, it seems, has discovered that artificial intelligence is changing the world. One congratulates them on their timing.

What unites these disparate efforts — the street protests, the white papers, the earnest policy frameworks — is a shared and touching faith that the thing can be governed the way we have governed previous technologies: by committee, by consultation, by the patient accretion of rules written in language so deliberately ambiguous that all parties can claim victory. This faith is, I submit, misplaced, though not for the reasons the Silicon Valley libertarians would have you believe.

The problem is not that AI should not be regulated. Of course it should. The problem is that the regulatory apparatus being constructed is designed for a technology that sits still, and this one does not. By the time [the Law Society finishes its guidance](#), the technology it describes will have evolved into something its authors would not recognize, much as a butterfly would confuse a committee that had spent three years drafting regulations for caterpillars.

I have some sympathy for the activists marching against data centres, if only because they have identified a concrete thing — power consumption, water usage, land appropriation — rather than the gossamer abstractions that occupy most AI policy documents. A data centre is a building. It draws from a grid. It can be measured. This is more than can be said for most of the harms that regulation purports to address.

The Atlantic Council's warning about defense implications deserves particular attention, because it names the thing that polite

regulatory discourse prefers to ignore: that every rule written for civilian AI constrains or enables military AI, and that adversaries unconstrained by such rules will not pause to admire our procedural integrity. This is not an argument against regulation. It is an argument against regulation conducted in a dream state.

What the moment requires — and what it is least likely to get — is regulatory humility: the admission that we are governing in fog, that the rules we write today will need rewriting tomorrow, and that the worst outcome is not insufficient regulation but ossified regulation that protects incumbents while punishing the very adaptation the technology demands. The companies that understand this, the ones already building AI into the marrow of their operations rather than treating it as a decorative appendage, will thrive regardless of what the committees produce. The rest will frame the white papers and hang them on the wall.

The regulators are coming. God help them. God help us all.



The Office Comic · ArtDesk

---

# Charlotte’s White-Collar Workers Reassured They’re Not Being Replaced, Just “Strategically Reimagined” Into Silence

*Executives urged to read a new leadership book so they can confidently mispronounce the future while firing it.*

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

---

CHARLOTTE — The latest evidence that artificial intelligence is transforming the modern workplace arrived this week in the form of a familiar corporate miracle: the same number of meetings, the same number of dashboards, and a noticeably reduced number of humans allowed to speak during either.

Local coverage of the city’s changing professional economy describes a white-collar scene in which software is increasingly tasked with writing emails, summarizing calls, and performing the other sacred duties once handled by junior employees who still believed “visibility” was a career strategy rather than a lighting condition. According to [one report on Charlotte’s workforce](#), the change is not merely technical. It is spiritual. It asks every salaried employee to locate their “unique human value” in a system that has already assigned it the status of optional.

Fortunately, leadership is responding with the kind of decisive action that has historically guided civilization through disruptive eras: a book.

A newly released executive guide, [AI Fundamentals For Leaders](#), promises to help decision-makers navigate 2026 with clarity, confidence, and at least three new ways to say “leverage” without committing to anything measurable. The modern executive, after all, is no longer expected to understand technology. They are expected to understand how to speak about it in tones that imply inevitability, ideally while gesturing toward a cost-saving plan that looks like innovation when viewed from far enough away.

This has become particularly important as the business world faces a secondary outbreak: strategic discourse now appears to be infected with what one analyst has dubbed “trendslop,” a thick slurry of buzzwords, thought-leadership fragments, and half-digested LinkedIn prophecy. The term has been circulating in debates about whether AI can do “strategy,” a question that would be easier to address if “strategy” hadn’t already been reduced to a set of slide transitions and a pledge to “stay agile.” As [one column warned](#), even the people paid to interpret the future are now struggling to decipher the paste.

Into this environment steps a third, more delicate innovation: the rebrand of layoffs as an act of technological enlightenment.

“AI washing” has become the managerial equivalent of choosing the corporate “Color of the Year,” an activity recently described elsewhere as an exercise in absurdity. Like the color trend, it is presented as a serious, data-driven ritual. Like the color trend, it functions primarily to help adults speak about arbitrary decisions as if they were dictated by the universe.

When a company announces it is “embracing AI” and, coincidentally, reducing headcount, the public is invited to admire the sleek inevitability of it all—never mind that the most consistent automation has been the conversion of payroll into shareholder reassurance.

In Charlotte, as in every city where professionals once assumed their job was “safe” because it involved Outlook, the new workplace compact is simple: humans will remain essential, provided they are willing to become the part of the process that apologizes for the process. The machines will handle the writing. Leadership will handle the vision. And the remaining employees will handle the exciting challenge of proving, quarterly, that they are not merely a cost center with feelings.

---

## ON THIS DAY IN AI HISTORY

*On April 15, 2016, Google’s AlphaGo defeated Lee Sedol, one of the world’s greatest Go players, winning their match 4-1—a landmark moment proving AI could master the ancient game’s intuitive complexity that had long seemed beyond machine reach.*

---