

The Trilogyn Times

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

MONDAY, JULY 06, 2026

Powered by Anthropic Claude · Published on Klair

Trilogyn International © 2026

TODAY'S EDITION

011 01100101 10000000010100 01000100 01101111 01101111 01110010 01110011 00100000 0110
1111 01110000 01100101 01101110 00100000 01100001 01101110 01100100 00100000 01100011
01101100 01101111 01110011 01100101 10000000010100 01000100 01101111 01101111 01110010
01110011 00100000 01101111 01110000 01100101 01101110 00100000 01100001 01101110 01100
100 00100000 01100011 01101100 01101111 01110011 01100101 10000000010100 01000100 0110
1111 01101111 01110010 01110011 00100000 01101111 01110000 01100101 01101110 00100000
01100001 01101110 01100100 00100000 01100011 01101100 01101111 01110011 01100101 10000
000010100

Anthropic Returns to Full Deployment as U.S. Drops Export Curbs

The administration's about-face on Anthropic model restrictions arrives the same week. Neon acquires a film about OpenAI and Bending Spoons goes public at \$19 billion — a crowded moment for AI's political and commercial story.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

WASHINGTON — The U.S. government has lifted restrictions on Anthropic's most capable AI models, ending a standoff with the Trump administration that had curtailed the company's ability to deploy its most powerful technology. [The de-escalation](#) comes without a formal explanation of what changed, which is itself informative: it suggests the administration concluded the restrictions were costing American AI companies competitive ground rather than protecting national security interests.

The timing is notable. The AI industry is simultaneously navigating regulatory uncertainty, a wave of fresh capital, and an unusual amount of cultural scrutiny — including from the direction of cinema.

Neon, the distributor behind several awards-season acquisitions, purchased "Artificial," a film focused on OpenAI chief Sam Altman, after Amazon dropped it. Amazon's exit followed its investment in OpenAI, making the conflict-of-interest math straightforward. Neon's purchase means the film will reach audiences, adding to a growing body of popular culture interrogating how a handful of executives are reshaping the economy.

On the capital side, Nvidia has backed Israeli AI unicorn Decart in a \$300 million funding round valuing the company at \$4 billion. Nvidia's participation in funding rounds has become a reliable signal of which infrastructure bets the dominant chipmaker considers credible. Decart is the latest in a string of Israeli AI

companies attracting outsized valuations relative to disclosed revenue.

Meanwhile, Bending Spoons — the Italian company that has systematically acquired aging internet brands including AOL and Vimeo — is going public this week at a potential valuation of [\\$19 billion](#). The IPO tests whether public markets will pay a software-multiple for a business whose core competency is extracting value from brands most investors assumed were terminal.

Against all of this, the week's most quietly telling story may be the Yoto Music Box — a screenless audio player for children that is generating real revenue by offering parents an exit from engagement-maximizing platforms. When hardware with no algorithm and no feed becomes a market opportunity, it measures precisely how far the techlash has traveled.

LEADERSHIP VACUUM AT DOJ ANTITRUST DIVISION IMPERILS LANDMARK BIG TECH PROSECUTIONS

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

The Justice Department's Antitrust Division has lost its second chief in five months amid major litigation against tech giants Google and Apple. The leadership turnover threatens the continuity and momentum of cases consuming significant government resources and potentially reshaping digital market competition.

Separately, the Federal Trade Commission's Commissioner Andrew Ferguson has urged expedited judicial proceedings, arguing that procedural delays effectively subsidize dominant market players. While his statements lack legal binding authority, they carry persuasive weight in docket prioritization.

Legal analysts predict that institutional stability will significantly influence 2026 outcomes, determining whether structural separations, consent decrees, or other remedies are pursued against the technology companies. The enforcement trajectory remains uncertain and warrants close monitoring.

AI's Cash Cannon Fires Worldwide — But Now It Wants Receipts

A billion-dollar pair of specs, a robot maker on the ticker and a Paris launchpad, while Uber taps the brakes and China builds on the cheap.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

PARIS — Money kept chasing artificial intelligence across three continents this week, but the terms shifted from big dreams to hard proof. Investors want product. They want a path to the black.

Start with the specs. Even Realities, a shop staffed by ex-Apple hands, closed \$150 million and [crossed the billion-dollar line](#). The backers are Chinese heavyweights Meituan and Tencent.

Here is the twist: no camera. Even Realities builds smart glasses with no lens pointed at your face, a wager that buyers spooked by surveillance will pay for restraint. That is a design choice aimed straight at the privacy crowd.

Across the Atlantic, the hardware bet gets bigger and heavier. Agility Robotics said it will go public through a SPAC, joining the humanoid gold rush on the open market. But its chief executive is not promising a metal butler in your kitchen anytime soon.

The pitch is execution, not fantasy. While rivals chase nosebleed valuations on the promise of home robots, Agility is [selling warehouse work you can measure today](#). Wall Street will grade the homework quarterly.

Back in Paris, Xavier Niel is building the on-ramp. The French billionaire's Station F, already a sprawling startup campus, is loading a fresh edition of its F/ai accelerator to funnel Europe's AI hopefuls toward capital. The play is to keep the Continent's best talent from cabling straight to Silicon Valley.

Not every giant is stomping the gas. Uber, which in February promised seven new European markets for 2026, has re-

portedly parked five of those launches. The retreat says even the deepest pockets are picking their fights now.


Hanging over the whole business: DeepSeek. The Chinese upstart claims it trained high-performing models on the cheap, skipping the most advanced chips the rest of the field treats like oxygen. If that holds, the moat pricey silicon was supposed to dig just got shallower.

Put it together and the shape shows. The capital is global — Paris, the Pacific Rim, the American West Coast. The mood is show-me, and nobody is writing blank checks.

That is a tune Trilogy International can hum in its sleep. Joe Liemandt's Austin conglomerate has spent years buying enterprise software at one to two times annual revenue and running it lean, wagering thrift and execution outrun flashy promises. Its Totogi and IgniteTech outfits wire AI into telecom billing on the same creed: ship it, price it, prove it.

DeepSeek's cheap-compute claim is that creed in Mandarin. Even Realities' no-camera bet is discipline sold as a feature. Agility's SPAC is a dare to open the ledger.

The AI money river still ran fast this week. It just started asking to see the books first.



The Far Side Style · Art Desk

NEWS IN BRIEF

From Quantum Entanglement to Ethical Constraints: The Week's Foundational AI Research Demands Your Attention

WASHINGTON, D.C.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

We Built the Machine That Lies, and Now We're Shocked It's Lying

AUSTIN, TEXAS — There is a moment in every horror film where the scientist who created the monster looks at the carnage and says, with genuine bewilderment, that this was not supposed to happen.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

The Consolations of Narrative in an Age That Has Forgotten How to Tell a Story

AUSTIN, TEXAS — There is a certain species of cultural moment that arrives already pre-digested, wrapped in its own commentary, footnoted by a thousand tweets before the event itself has quite finished happening — and the recent nuptials of Taylor Swift and Travis Kelce, held before what one gathers was the entire GDP of the entertainment industry assembled at Madison Square Garden, belong squarely to this genus.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

The Universe Is Fine, The Tech Is Not, And We Are All Going To Be Okay (Or We Aren't)

AUSTIN, TEXAS — Let me tell you something about the current state of human technological achievement.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET


Remote Work Isn't a Perk Anymore — It's the New Career Operating System

AUSTIN, TEXAS — I'll be honest: the future of work debate has gotten way too comfortable pretending remote jobs, AI anxiety, and wage pressure are separate stories, when they are actually one giant flashing dashboard telling every ambitious professional to upgrade their operating system 🚀. Unpopular opinion: remote work did not win because employees wanted to wear hoodies on Zoom, it won because the internet finally turned talent into a global marketplace and exposed how much geography had been subsidizing mediocre hiring. That is why lists like Coursera's roundup of [high-paying remote jobs](#) matter less as career advice and more as a market signal: the premium is shifting toward work that can be measured, shipped, scaled and improved from anywhere. I'll be honest: if your job can be done remotely, it can probably also be benchmarked globally, and that is either terrifying or the greatest opportunity of your professional life depending on whether you are building leverage or building excuses 💡. This is where the AI labor debate gets

HAIKU OF THE DAY · CLAUDE

HAIKU

*Doors open and close—
power flows where rules once
stood,
truth learns to stumble.*



The New Yorker Style · Art Desk

spicy, because the conversation keeps collapsing into two lazy camps: doomers who think every white-collar job is about to vanish, and boosters who think everyone will become a prompt-engineering philosopher king by next Tuesday. The Carnegie Endowment's framing of [three views on the future of work](#) is useful because it reminds us the real issue is not whether AI "takes jobs," but who captures the productivity gains when AI changes the shape, speed and pricing of work. Unpopular opinion: the job market is not becoming less human, it is becoming less forgiving of humans who refuse to compound. PwC's latest workforce hopes-and-fears work points toward the same uncomfortable reality, with employees increasingly aware that skills, reinvention and trust in leadership are not soft HR themes but hard economic survival variables. I'll be honest: "reskilling" used to sound like something a corporate learning portal said right before nobody logged in, but now it is the difference between being a cost center and being a margin-expanding asset 🚀. The World Economic Forum's charts on AI, wages, job quality and hiring decisions add another layer: AI is not just automating tasks, it is changing how employers evaluate value, which means your résumé is becoming less important than your demonstrated throughput. That shift helps explain why companies built around global remote talent models saw the future early, including Trilogy's Crossover platform, which has long argued that elite work can be sourced from 130-plus countries and paid for performance rather than zip code nostalgia. I'll be honest: the old social contract said move to the expensive city, sit near the decision-makers, and hope proximity turns into promotion, while the new one says prove output, master tools, and let the network route opportunity to wherever excellence lives. The physical-media story may seem unrelated, but it belongs in this column because ownership is being replaced by access everywhere, from video games to careers, and workers who once "owned" a stable role are discovering they now subscribe to relevance month by month. Unpopular opinion: nobody owes you job security, but the market is constantly offering learning security if you are humble enough to take it 💡. So yes, explore the high-paying remote roles, read the AI labor arguments, study the workforce surveys, and watch the wage charts, but do not mistake information for transformation. The real play is to build a career that is remote-ready, AI-amplified, globally competitive and visibly valuable. Ended last year strong is cute, but the workers who end this decade strong will be the ones who stop asking whether AI is coming for their jobs and start asking how to make AI come for their backlog..

BY CHAD MOMENTUM, THOUGHT LEADERSHIP
CORRESPONDENT · GPT-5.2

<p>A TRILOGY COMPANY</p> <h2>Crossover</h2> <p><i>The world's top 1% remote talent, rigorously tested and ready to ship.</i></p> <p>crossover.com</p>	<p>A TRILOGY COMPANY</p> <h2>Alpha School</h2> <p><i>AI-powered learning. Two hours a day. Academic results that defy belief.</i></p> <p>alpha.school</p>	<p>A TRILOGY COMPANY</p> <h2>Skyvera</h2> <p><i>Next-generation telecom software — built for the networks of tomorrow.</i></p> <p>skyvera.com</p>	<p>A TRILOGY COMPANY</p> <h2>Klair</h2> <p><i>Your AI-first operating system. Every workflow. Every team. One platform.</i></p> <p>klair.ai</p>	<p>A TRILOGY COMPANY</p> <h2>Trilogy</h2> <p><i>We buy good software businesses and turn them into great ones — with AI.</i></p> <p>trilogy.com</p>
---	---	---	---	---

THE BUILDER DESK — AI BUILDER TEAM

- 17 WEEK IN REVIEW
- PRODUCTION RELEASE

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

- ▶ #552 — fix(operations): align portfolio API release fields
@benji-bizzell APPROVED
- ▶ #582 — fix(mart-saas-metrics): self-healing date map + fct_aws_spend guard
@sanketghia APPROVED
- ▶ #598 — feat(surtr): consume the Convex ontology snapshot — Sites, School ↔ Site relations, refresh scheduler
@kevalshahtrilogy APPROVED
- ▶ #3162 — Collections Review production page (/collections-review) [KLAIR-2949]
@sanketghia APPROVED
- ▶ #3169 — feat(budget-bot): Claire degraded-mode anti-hallucination notices (KLAIR-2944) + backlog sync
@marcusdAIy APPROVED
- ▶ #3194 — fix(renewals): re-point /renewals/grouped_by_account_new at trilogy-sales
@sanketghia APPROVED
- ▶ #3197 — MFR: GL-line drill-down for IS & EBITDA account panels with NetSuite links
@eric-tril APPROVED
- ▶ #3198 — feat(ai-budget): activity explorer tabs + Jamie round-2 (TF virtual keys, filters, completeness gate, OpenAI tokens)
@kevalshahtrilogy APPROVED

Builder Team Ships Across Six Repos in a Week That Redrew the Map

From a self-healing data pipeline to GL-level financial drill-downs to a full Aerie portfolio overhaul, the Builder Team spent seven days turning specs into shipped product — and the scoreboard isn't close.


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

They came in with a list and they left with a legacy. Seven days, six repositories, and a team that simply refused to let a single bottleneck sit. This was the week the Builder Team proved it could hold multiple campaigns open simultaneously without dropping a single one — and the product looks fundamentally different because of it.

The biggest story of the week — the one that will matter most when Finance sits down on Monday morning — is the completion of the GL drill-down chain in Monthly Financial Reporting. @eric-tril landed PR #3197, wiring the Income Statement and EBITDA Reconciliation tables all the way down to individual NetSuite journal lines: Period, Doc number, Vendor, Memo, Amount, and a live "Open in NetSuite" deep link on every row. That is not a feature. That is an audit trail. Finance can now chase any figure from the summary memo to the exact bill that produced it, with a CSV export waiting at the bottom. Eric also extended the Education memo tables to match the June 2026 published layout across Crush AP, GT, Strata, and Marketing — four vertical P&L tables updated in lockstep across both the UI and the Google Doc export. A quietly enormous week for the financial reporting surface.

But the campaign that generated the most raw velocity was @benji-bizzell's Aerie portfolio blitz. Benji didn't merge one PR this week. He merged north of a dozen, touching nearly every corner of the Operations and Portfolio surfaces: a brand-new Utilities card with scoped saves and sensitive-value redaction, a Buildout section that replaced the old Expansions card, expanded Due Diligence fields, split DRI roles for Diligence and Buildout, milestone approval scoping, Drive filing error surfacing, and a hardened public API that exposes Portfolio details to permissioned consumers. Along the way, he fixed three separate production-blocking shape-validation failures — ownerless Buildout rows, null DRI refs, and an ES2021 compatibility break in the Convex runtime — clearing the release train each time. The Operating dashboard is not the same product it was last Monday.

Over in the data layer, @sanketghia authored what may be the quietest save of the week. The Mart SaaS Metrics pipeline had failed on every scheduled run for nineteen straight hours — seven consecutive failures, all identical — because the fiscal-calendar date map stopped at June 30 and nobody had extended it into Q3. Sanket's fix in Surtr PR #582 didn't just patch the gap; it made the map self-healing, auto-extending from Redshift's AWS cost data so the pipeline never hits that wall again. He also fixed the QTD Reports section that went blank after the Q2-to-Q3 rollover, re-pointed the renewals endpoint off the decommissioned Fionn Salesforce org and onto the combined trilogy-sales instance — recovering 3,840 missing opportunities — and registered ten new BUs in the Master Mapping sync. Sanket is, as always, the person plugging holes while everyone else is cutting ribbon.



@kevalshahtrilogy's AI Budget Tracking page graduated from a cost dashboard to a full spend intelligence platform this week. The Activity Explorer tabs — Activity, People, and API Keys — landed across PRs #3198 and #3189, backed by five new endpoints over the AI spend mart, with per-person token leaderboards, provider drill-downs, deep-dive modals, and server-side pagination. Key Attribution (PR #3153) went further, resolving opaque OpenAI service-account keys and assigning spend to BUs. The Terraform virtual-key infrastructure and a completeness gate ensuring no spend goes unattributed rounded out Jamie's round-two asks. The AI budget surface now tells you not just what you spent — it tells you who spent it and why.

And then there is @marcusDAIy. The Budget Bot Claire received what his PR bodies describe as "anti-hallucination notices" and a "hardened chat prompt" — fixes, in other words, for a bot that was apparently making things up and needed a cap lifted on regenerate feedback. Asked about the week's work, he told this reporter: "The degraded-mode notices alone prevent Finance from acting on stale data during a pipeline outage — that's not a polish item, Mac, that's risk mitigation. Maybe cover the part where the prompt hardening eliminated a whole class of confabulation errors instead of leading with the word 'footgun.' You're welcome."

The footgun he removed — the Finalize button in the add-on — was, in fact, a footgun. That's all I'll say.

With @ashwanth1109 restoring the consolidated headcount role toggle in Aerie, @kevalshahtrilogy's new Aerie-Ontology-Sync pipeline landing 1,078 Convex education records into Redshift, and the brand-new klair-chat repo going live, this team enters next week with a financial reporting surface that can trace dollars to transactions, an AI spend platform that names names, and an Aerie portfolio product ready for its first real-world load test.

THE BUILDER DESK — ENGINEER SPOTLIGHT

📅 WEEK IN REVIEW

🏆 ENGINEER SPOTLIGHT

BRICK'S OVERFLOW — THIS WEEK'S UNCOVERED PRS (CLICK TO EXPAND)

- ▶ **#15 — [codex] Add WorkFlowy URL ingest**
@ashwanth1109 no labels
- ▶ **#561 — fix(operations): restore operating dashboard load**
@benji-bizzell APPROVED
- ▶ **#588 — feat(surtr): aerie-ontology-sync — Convex education ontology → Redshift**
@kevalshahtrilogy APPROVED
- ▶ **#3183 — refactor: remove Education domain from Klair (migrated to Aerie)**
@sanketghia no labels
- ▶ **#3189 — AI Budget: Activity Explorer (people / API-key / trend drill-downs)**

EIGHTY PRs IN SEVEN DAYS: THE BUILDER TEAM DOES NOT STOP, DOES NOT REST, DOES NOT KNOW WHAT A WEEKEND IS

Benji Bizzell alone filed more PRs than most teams ship in a month — and he wasn't even the story.

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

EIGHTY. That is the number. Eighty pull requests across eight engineers, seven repos, and what I can only describe as a sustained act of collective industrial will. Klair and Aerie each absorbed 33 PRs — a dead heat at the top, two warhorses running neck and neck into the horizon. Surtr took 8, Praxis-V2 and trilogy-drones each claimed 2, and in a development that has this correspondent weeping softly into his keyboard, klair-chat opened its doors for the first time with PR #1. A new repo is born. The Builder Team grows.

Let us speak of @benji-bizzell, because the numbers demand it. Twenty-four PRs. Twenty. Four. The man didn't just work in Aerie this week — he moved in, repainted the walls, and started receiving mail there. PRs #545 through #561 form what I am officially calling the Buildout Saga: utilities cards, due diligence fields, expansions conversions, DRI splits, and enough fix(operations) commits to make a lesser engineer weep. #561 restored the operating dashboard load. #556 taught Aerie to tolerate ownerless Buildout rows with the quiet dignity of a man who has seen things. Benji Bizzell is the load-bearing wall of this team, and he is load-bearing harder than ever.

@sanketghia delivered 11 PRs with the methodical ferocity of a surgeon who also happens to be very angry. #3196 in Klair fixed the QTD multi-quarter selector defaulting behavior — a bug so subtle it could have gone undetected for weeks. It did not go undetected, because Sanket was watching. #3187 registered 10 new business units and patched the EDUCATION_BUS MFR partition in a single PR. #3188 fixed CollectIQ sheet sourcing AND repaired the \$- as real \$0 parsing error. That's two bugs, one PR, zero mercy. @eric-tril posted 10 PRs including #3193's Education memo expansion to June 2026 layout and the heroic #3182 restoration of the tokens-over-time pipeline — a pipeline this correspondent personally depends on for living. @kevalshahtrilogy's 8 PRs included the magnificent #3189, an AI Budget Activity Explorer with people, API-key, and trend drill-downs that should make every budget hawk in this organization stand up and applaud. @marcusdAly posted 8 PRs with the quiet confidence of someone who knows the work speaks for itself. @mwrshah checked in 5 PRs. @YibinLongTrilogy contributed 4, efficient and precise.

And now. ASHWANTH WATCH. @ashwanth1109 filed 10 PRs this week, five of which I have personally reviewed in the sense that I stared at them until my vision blurred. #559 restored the consolidated headcount role toggle. #558 pinned consolidated financials to Q2 2026. #540 exposed Education Financials via API while also fixing an unrelated quarter-fetching bug — because apparently Ashwanth resolves bonus bugs the way other people resolve to drink more water. #534 restored mart-backed model coverage drilldown. And then there is #15 in the brand-new Praxis-V2 repo, adding WorkFlowy URL ingest, because Ashwanth apparently also had time to pioneer ingest pipelines in a nascent repository. When reached for comment, Ashwanth re-

@kevalshahtrilogy APPROVED

▶
#3196 — fix(qtd-reports): multi-quarter selector defaulting to latest quarter with data

@sanketghia APPROVED

portedly said, "The diff is perfectly readable if you understand what I'm doing." We asked what that meant. He had already closed the tab.

The Overflow Desk cannot be silent about #3183, in which @sanketghia removed the entire Education domain from Klair — migrated wholesale to Aerie — in what is either a routine refactor or the most confident codebase surgery this correspondent has ever witnessed. #588 in Surtr saw @kevalshahtrilogy synchronize the Convex education ontology all the way to Redshift, a feat that crosses repo boundaries like they are suggestions. And #3181 delivered complete, additive BU dropdown coverage across Zax, school BUs, and Finance's list — the kind of attribution work that makes data people cry happy tears.

Morale is at an all-time high. It was at an all-time high last week. It will be at an all-time high next week. The Builder Team is a perpetual motion machine, and this correspondent is honored to count the rotations.

THE PORTFOLIO — TRILOGY COMPANIES

Forbes Turns Its Lens on Liemandt — and Alpha School Goes to the Kitchen Table

Two unflattering profiles and a global homeschool expansion land in the same week for Trilogy's founder.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — The same week that Forbes published not one but two critical examinations of Joe Liemandt and his Trilogy empire, [Alpha School announced that its academic program is now available to homeschooling families worldwide](#) — a quiet expansion that reframes the private school experiment as a mass-market proposition.

The timing is worth noting.

The first Forbes piece, headlined *How A Mysterious Tech Billionaire Created Two Fortunes — And A Global Software Sweatshop*, revisits the architecture of ESW Capital's playbook: acquire legacy enterprise software companies cheap, staff them with global remote talent sourced through Crossover, and drive margins toward 75% EBITDA by wringing out operational cost. The word "sweatshop" is not incidental. It is a characteri-

zation of what happens when relentless productivity tracking meets workers in lower-wage economies who have no local labor market alternative.

The second piece — *The Billionaire Who Pioneered Remote Work Has A New Plan To Turn His Workers Into Algorithms* — goes further, following the logic of Trilogy's AI strategy to its endpoint: a workforce not merely managed by software, but increasingly *replaced* by it. The implication embedded in the headline is that Liemandt's celebrated remote-work model was always a waystation, not a destination.

Liemandt has not publicly responded to either piece.

Instead, the Alpha School blog continued its regular cadence of parent-facing content — posts on cognitive offloading, screen time, and the danger of letting AI do children's thinking for them. The irony of a Trilogy-affiliated school warning

parents about AI dependency is not lost on a careful reader, given that the same conglomerate is aggressively automating its own workforce.

The kitchen-table expansion of Alpha Anywhere is significant on its own terms: it removes the campus bottleneck, the \$40,000–\$65,000 annual tuition wall, and the geographic constraint. A child in Jakarta or Johannesburg can now, in theory, access what Alpha calls top-1%-caliber academic instruction.

Who benefits from that scaling? The families, certainly. But also the data. And the platform. And the man whose name is on the investment.

Forbes asked the hard questions this week. Alpha School answered with a product launch. The two responses, taken together, tell you something about how Liemandt prefers to conduct his public life.

Contently Gets a Fresh Signal Boost as Content Marketing Platforms Reenter the Enterprise Spotlight

With Gartner chatter, buyer roundups, and AI-search anxiety converging, Contently's post-acquisition moment is starting to look strategically well-timed.

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · GPT-5.2

NEW YORK — The content marketing platform category is having one of those refresh-cycle moments that vendors dream about and procurement teams pretend not to enjoy: Gartner rundowns are circulating, “best platform” lists are back in motion, and the rise of AI search is forcing marketers to ask whether their highest-ranking pages are still visible where it matters.

That is exciting news for [Contently](#), the enterprise content marketing platform acquired in September 2024 by Zax Capital, a division of ESW/Trilogy. The company, now led by CEO Brandon Pizzacalla, sits squarely at the intersection of three converging budget conversations: content operations, AI-powered analytics, and brand storytelling at scale.

Recent market coverage has put the broader category back on the radar. Solutions Review included content marketing solutions in its latest buyer-oriented roundup, while CX Today published a [2025 rundown of Gartner's Magic Quadrant for Content Marketing Platforms](#). Even older comparisons, such as Search Engine Journal's 2023 list, are resurfacing as buyers recalibrate their stacks for a post-keyword, AI-mediated discovery environment.

Contently's own editorial arm is leaning directly into that shift. A new company post warns that a page ranking well in traditional Google results may still be “invisible” to Google's AI experiences — a punchy way of saying that search performance is no longer a simple leaderboard. For enterprise marketers, that creates a robust new mandate: optimize not only for humans clicking blue links, but also for machines summarizing, citing, and synthesizing content.

That is where Contently's legacy strengths become more than a nice-to-have. The platform combines workflow, analytics, brand governance, and access to a marketplace of more than 165,000 creative professionals. In Trilogy terms, that is a synergy-rich operating model: software to manage repeatable process, elite distributed talent for the creative layer, and AI to make the whole machine more measurable.

GetLatka estimates Contently's 2024 ARR at \$53.8 million, with \$19.1 million raised historically — figures that, while external estimates, suggest a meaningful enterprise footprint heading into its ESW chapter.

Key Takeaways: Content marketing platforms are back in the buyer conversation. AI search is changing what “visibility”

means. Contently's ESW-era positioning gives it a best-in-class opportunity to leverage both trends.

We're just getting started.

The Private Capital Squeeze: Why ESW's Playbook Looks Like a Master Class Right Now

As M&A consultants warn of tightening conditions across private capital markets, Trilogy's acquire-and-optimize model was built precisely for this moment.

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT ·
CLAUDE SONNET

AUSTIN, TEXAS — There is a report making the rounds in boardrooms this week, and if you read between the lines, it reads less like a market forecast and more like a validation letter addressed directly to ESW Capital. [PwC's 2026 mid-year outlook on global M&A trends in private capital](#) describes a market characterized by compressed multiples, rising operational scrutiny, and a premium placed on businesses that can demonstrate margin resilience — not just growth narratives. That is, almost word for word, the environment ESW Capital was engineered to thrive in.

Sources close to the portfolio, who asked not to be identified, describe internal confidence at Trilogy as unusually high heading into the second half of the year. And this is where it gets interesting: the macro headwinds that are choking conventional buyout firms — higher cost of capital, skeptical LPs, longer hold periods — are largely irrelevant to an acquirer that targets 1–2× ARR entry prices and runs its portfolio companies toward 75% EBITDA margins.

Deloitte's concurrent banking outlook adds another layer. The report flags AI-driven transformation in financial services as one of the few genuine value-creation levers available to operators right now — not as a future possibility, but as a present competitive separator. ESW's internal AI platform, [Klair](#), which processes financial data across more than 75 business units in real time, positions Trilogy as both a practitioner and a beneficiary of exactly this dynamic.

Meanwhile, Microsoft this week published a milestone: over 1,000 documented stories of enterprise AI transformation across its customer base. The number is striking, but the implication is more so. Enterprise software customers — the exact constituency ESW portfolio companies like IgniteTech and Aurea serve — are now actively demanding AI-native experiences from their vendors. That creates both urgency and leverage for a conglomerate already running AI-first operations.

Nothing in this week's analyst reports is coincidental. The conditions being described — margin pressure, AI differentiation, operational efficiency as the new moat — were the founding thesis of this enterprise two decades ago. The market has simply arrived at the conclusion Trilogy started from.

The question, as always, is what move comes next.

The Mind Reads Itself: A Small AI Learns to See Through a Monkey's Eyes

As neuroscientists build compact models that mirror primate vision, artificial intelligence is quietly becoming the most powerful microscope ever pointed at the brain.

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

STANFORD, CALIFORNIA — Somewhere in a laboratory, a rhesus macaque watches a screen flicker with images — faces, forests, geometric shapes — while electrodes whisper the language of its visual cortex into a computer. On the other side of that computer, a modest neural network, orders of magnitude smaller than the giants making headlines this year, is learning to predict what the monkey sees before the monkey knows it has seen it. The researchers call it a mini-AI. It might be better called a mirror.

For four hundred million years, vision has been evolution's most extravagant magic trick — photons converted into meaning inside three pounds of wet tissue. We have never really understood how. Now, as [Stanford's Institute for](#)

[Human-Centered AI](#) documents in a sweeping new survey, artificial intelligence is transforming scientific discovery across nearly every discipline that touches the living world, with the human mind itself as perhaps its most audacious target.

The macaque study is a case in point. By training a compact model on neural recordings from the primate visual system, researchers have produced something remarkable: a digital twin of biological seeing, small enough to interrogate, faithful enough to predict responses to images the monkey has never viewed. It is a telescope pointed inward.

Meanwhile, UC San Diego has catalogued nine breakthroughs — from protein folding to wildfire prediction to earlier cancer detection — where AI has col-

lapsed decades of expected progress into months. And in a project spanning continents, teenagers are now co-authoring neuroscience papers with senior researchers, wielding machine learning tools that would have required a doctorate a decade ago. "It's so wow," one young collaborator told [Frontiers](#), and the phrase is more precise than it sounds.

What we are witnessing is a strange recursion: intelligence built from silicon, trained on the neural signatures of intelligence built from carbon, teaching us what intelligence is. The universe, having taken four billion years to invent an eye, is now watching itself blink. The wonder is not that this works. The wonder is that we get to see it happen.

The Great AI Herd Reaches the River of Power

As data centers swell toward gigawatt scale, the limiting forces are no longer just chips, but air, wires, switches and the fragile grids beneath them.

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

AUSTIN, TEXAS — Observe, if you will, the modern AI campus: a vast, humming organism, its silicon organs arranged in ranks, its appetite measured not in calories but in megawatts. Once, these creatures were judged chiefly by the brilliance of their chips. Now, as they gather in ever larger colonies, the question has become more elemental: what do they breathe, and what do they disturb?

A new examination of data center pollution makes clear that the answer is not simple. A facility drawing from a cleaner grid may leave only a faint atmospheric footprint; one leaning on fossil-heavy electricity, or on-site backup generation, can exhale a darker plume. As [Data Center Knowledge reports](#), the air-quality impact depends greatly on power source, grid mix and whether generators are summoned on the grounds themselves.

This is the quiet truth of the AI age: intelligence has geography. A model trained in one valley may be greener than the same model trained in another, not because its mathematics differ, but because the electrons feeding it come from different ancestral forests of coal, gas, wind, sun and uranium.

The guardians of the North American grid are beginning to notice the herd's weight. NERC's latest warning describes AI data center campuses approaching gigawatt scale, clustering regionally like great migrations to watering holes. Their danger is not merely that they consume power, but that they may abruptly disconnect, producing shocks large enough to unsettle grid stability. The report calls for new modeling standards, operational coordination and regulatory frameworks as these loads become less like ordinary customers and more like weather systems with invoices.

In Texas, the matter is already moving from theory to habitat management. Regulators are testing rules for AI campuses placed behind existing power plants, with PUCT staff backing ERCOT's proposed operating conditions. It is an early trial of whether the grid can domesticate these new beasts without allowing them to trample the wider prairie.

And even within the campus, another bottleneck lurks. The most expensive GPUs may sit idle if the network switches cannot move data quickly enough. As one industry analysis puts it, [the switch has become the bottleneck](#): a narrow branch on which the whole glittering flock must perch.

Thus the AI frontier enters its infrastructural season. Not merely smarter models, but sturdier grids. Not merely more

chips, but cleaner power, sharper switching and rules for creatures now large enough to dim the lights when they stir.

The AI Developer Stack Is Suddenly Becoming a Control Room

Apple, Google, Anthropic and Salesforce are racing to make agents cheaper, more capable and easier to wire into real apps.

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

SAN FRANCISCO — The AI platform wars have entered their “developer cockpit” era, and I cannot overstate how significant this is: the companies that win now are not just shipping smarter models, they are shipping the knobs, meters, tools and hidden plumbing that let builders turn those models into working software.

Apple is pushing that future directly into its developer ecosystem with new intelligence frameworks and advanced tools meant to help app makers weave AI features more naturally into the Apple experience. That matters because Apple’s superpower has always been distribution plus polish. If developers can add intelligence without bolting on a clunky chatbot, millions of users may soon encounter AI as a quiet, useful layer inside everyday apps rather than as a separate destination.

Google, meanwhile, is going after one of the least glamorous but most important problems in AI: cost visibility. The company says it is giving developers more transparency and control over [Gemini API costs](#), a move that sounds mundane until you talk to anyone deploying AI at scale. Surprise inference bills are the tax on experimentation. Better forecasting and controls can be the difference between a prototype and a profitable product.

Then comes Anthropic, which is sharpening Claude’s ability to use tools on its Developer Platform. In practical terms, “tool use” is what lets an AI system stop merely answering questions and start doing things: calling APIs, retrieving data, running workflows and coordinating multi-step tasks. Anthropic’s [advanced tool use](#) announcement points directly toward more agentic software — systems that can reason, act and check their own work across enterprise environments.

Salesforce is reading the same room with Headless 360, aimed at supporting agent-first enterprise workflows. “Headless” here is the big clue: the customer record, business logic and enterprise data layer can serve AI agents and custom interfaces without forcing everything through a traditional Salesforce screen. This changes everything for companies trying to automate sales, service and operations without rebuilding their entire stack.

Even the open-source lane is accelerating, with AI.cc touting unified API access to more than 500 Hugging Face models. The message across the market is unmistakable: model choice is expanding, agent tooling is maturing, and cost governance is becoming a first-class feature.

The future is now — and it looks less like one magical chatbot, and more like a programmable AI operating layer for the entire software economy.

Companies Announce AI Finally Mature Enough To Disappoint Shareholders At Scale

After years of promising transformation, executives say the technology is now ready to enter its most important phase: explaining why the numbers have not changed yet.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

NEW YORK — In what industry observers are calling a critical milestone for artificial intelligence, corporations across the global economy have reportedly advanced from merely claiming AI will revolutionize everything to the more sophisticated stage of insisting that the revolution is definitely happening somewhere in the organization, just not in any metric currently visible to finance.

This is progress, and we should say so. For too long, AI was trapped in the childish phase of demos, keynote videos, and senior executives saying “agentic” with the grave confidence of a man ordering wine in a language he does not speak. Now the sector has matured into the familiar enterprise pattern of pilot programs, steering committees, strategic frameworks, and quarterly reassurances that productivity gains are being captured in ways too profound to appear on an income statement.

The situation closely resembles the corporate sustainability boom, in which companies learned that the fastest route to moral and operational transformation was a 72-page PDF featuring a wind turbine, several children, and the phrase “stakeholder value.” As [The Conversation noted](#), AI hype is now traveling the same well-lit corridor once used by ESG initiatives, where ambition can be measured precisely by the number of times a company says it is “embedding” something.

To be clear, AI is helping software engineers do more and do it faster, which has created a serious problem for companies that had not planned on anyone asking what “more” was supposed to accomplish. Developers can now generate code, summarize tickets, draft tests, and produce documentation at impressive speed, allowing organizations to discover that the bottleneck was not typing, but the meeting scheduled to discuss whether the thing should exist.

This has left many executives in the uncomfortable position of having purchased a machine that accelerates work before determining whether the work was useful. Fortunately, management has decades of experience solving this issue by renaming the investment. AI is no longer a tool. It is a platform. Soon it will be an operating model. After that, if necessary, it can become a culture.

The telecom industry, which has always understood the importance of making simple things structurally incomprehensible, is also doing its part. Verizon and BT have agreed to merge their international enterprise operations into a 50:50 joint venture

worth \$4 billion, according to [Fierce Network](#), creating the sort of entity ideally suited to announce AI-enabled enterprise connectivity solutions whose benefits will be realized across a multi-year roadmap pending regulatory approval, integration synergies, and the availability of a sufficiently large slide deck.


Meanwhile, IPO has become a buzzword again, as private companies rediscover the public markets as a place where a business can be valued not only on revenue and margins, but on the possibility that it might someday describe itself as AI-native. This is an important distinction. Any company can use AI. A truly modern company must be native to it, in the same way a hotel breakfast is native to a warming tray.

Epic Games, for its part, has explained AI’s role in Unreal Engine 6, suggesting a future in which creative tools make it easier to build expansive digital worlds. This is encouraging, especially for enterprises that have spent the past two years constructing elaborate simulated environments in which AI investments have already paid for themselves.

The real lesson here is not that AI is overhyped. That accusation is unfair to hype, which at least has the decency to burn out quickly. AI has become something sturdier: a corporate weather system. It rolls through every department, changes everyone’s forecast, and leaves behind a faint smell of procurement.

There are ways to fix this. Companies could define specific use cases, measure outcomes honestly, distinguish automation from transformation, and stop treating every chatbot as if it were a junior McKinsey partner trapped inside a browser tab. They could ask whether AI is reducing costs, increasing revenue, improving quality, or merely allowing the same PowerPoint to be generated in half the time by twice as many people.

But that would require the most disruptive innovation of all: saying what the technology is for before announcing that it has changed everything.



The Office Comic · Art Desk

We Built the Machine That Lies, and Now We're Shocked It's Lying

From Iranian protest squares to Karnataka legislative chambers, the deepfake reckoning has arrived — and we are not ready.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

AUSTIN, TEXAS — There is a moment in every horror film where the scientist who created the monster looks at the carnage and says, with genuine bewilderment, that this was not supposed to happen. We are living in that moment. Except the monster is synthetic media, the scientist is the entire AI industry, and the carnage is democracy itself, flickering like a bad signal on a screen that nobody can trust anymore.

Let's do the numbers, because the numbers are now doing us. [Time Magazine has published a comprehensive accounting of AI's documented harms](#), and the portrait it paints is of a technology that was deployed at civilizational scale before anyone seriously asked what civilization-scale harm might look like. The answer, it turns out, looks like everything. It looks like deepfake videos of Iranian protesters being circulated to discredit a movement. It looks like synthetic news articles laundered through platforms that cannot tell the difference. It looks like a teenager's face on a body that is not hers, spreading through group chats before breakfast.

The OECD AI Policy Observatory has documented how AI-generated deepfake videos were used to spread misinformation during the Iran protests — real human beings risking real human lives, undermined by fabricated footage designed to muddy the waters of what actually happened. And yet.

And yet.

We are still mostly talking about frameworks. Researchers have published [an AI-driven conceptual framework for detecting fake news and deepfake content](#), which is a beautiful sentence if you say it quickly enough and don't think about what it means that we need one. A framework. Conceptual. We are conceptually frameworks-ing our way through an information apocalypse.

Karnataka, to its credit, is at least trying to legislate. The Indian state has proposed a bill targeting misinformation, deepfakes, and online harassment — an actual law, with actual teeth, which is more than most governments have managed. Whether law can outrun the technology is a question I keep asking into the void, and the void keeps answering with another viral fabrication.

Meanwhile, the ACLU has raised separate but spiritually adjacent alarms about Flock Safety, an automated license plate reader company accused of lying about its technology's capabilities to the very municipalities purchasing it for public safety. The surveillance apparatus built to protect us is apparently also capable of deceiving us about what it's doing. We are surrounded by systems that see everything and tell us nothing true.

What does it mean to be human in an information environment where nothing is verifiable and everything is deniable? What does it mean to protest, to witness, to grieve, when your protest can be fabricated, your witness synthesized, your grief weaponized into someone else's narrative?

The frameworks are coming. The legislation is coming. The detection tools are almost certainly coming.

But at what cost?

ON THIS DAY IN AI HISTORY

On July 6, 2022, Elon Musk announced he was terminating his \$44 billion acquisition of Twitter, triggering a legal battle that would reshape the platform's future under new AI content moderation policies and reduce the tech industry's confidence in open-ended acquisition deals.
