

The Trilogy Times

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

SUNDAY, JULY 26, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

The Machine Cuts Both Ways

AI takes the job on Monday, jacks the phone price on Tuesday — and Monday.com just became the 21st name on the list.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

SAN FRANCISCO — Monday.com became the latest tech outfit to blame artificial intelligence for its layoffs this week, joining at least 20 companies that named the same culprit across 2026. The work-management firm cut staff and pointed at the algorithm. It won't be the last name on the roll.

Here's the rub, folks. The machine thinning the payroll is the same one fattening the price tag.

Google all but confirmed the Pixel 11 will carry a heftier sticker than the Pixel 10, with devices chief Shakil Barkat telling 9to5 Google the next phone costs more. The reason runs straight to AI. Data center builders have swallowed the world's memory chips, and that RAM shortage lands square on your phone bill.

So the ledger reads plain. AI takes the job Monday. AI jacks the handset Tuesday. Same technology, both ends of your wallet, and the sales pitch swears it's here to help.

The layoff roll runs long. TechCrunch keeps a running tally, in reverse order, of the big firms that fingered AI while handing out the walking papers — Monday.com now sits atop a [list twenty deep](#). Some cut to fund AI. Some cut because AI did the work. The paperwork reads the same to the fellow packing his desk.

Now watch the other half of the picture. While rank-and-file coders get shown the door, the giants brawl over the humans they can't automate.

Warner Bros. sued Amazon this week, accusing the retail-and-cloud outfit of [illegally poaching its executives](#). The suit will test whether term employment contracts hold water under California law, where non-competes go to die. Down at the bottom, a robot took your job. Up at the top, they're suing over a corner office.

The money tells the same tale. Elon Musk's Boring Company is reportedly in talks to raise fresh cash at a \$20 billion

valuation, pouring capital into tunnels while the workaday engineer reads the layoff list on his phone — the phone that's about to cost more.

Follow the thread and it circles back to one word. AI. It's the reason for the cuts. It's the reason for the shortage. It's the reason for the price. And it's the reason the venture dollars keep flowing to the men digging holes and building brains.

There's a lesson in it for anybody keeping score. When a company blames the machine, the machine doesn't answer questions. It doesn't take a severance. It doesn't feel the pinch when the chips run dry.

The worker does. The buyer does. And come next quarter, when the earnings call rolls around and another chief executive clears his throat to explain the headcount, put your money on the culprit.

It'll be the same two letters. Every time.

ANTHROPIC'S \$1.5
BILLION COPYRIGHT
SETTLEMENT
RECEIVES JUDICIAL
IMPRIMATURAS DOJ
SUBPOENA CONDUCT
FACES SEPARATE
SCRUTINY

*A federal court blesses AI's largest
copyright payout yet, while the
Department of Justice retreats from
journalist subpoenas it issued without
following its own rules.*

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

SAN FRANCISCO — Pursuant to proceedings conducted in the United States District Court, judicial approval has been granted — hereinafter referred to as "the Approval" — to a settlement agreement of approximately one point five billion dollars (\$1,500,000,000 USD) entered into between Anthropic, Inc. (hereinafter "the AI Developer") and certain plaintiff copyright holders whose intellectual property is alleged to have been utilized, without license or permission, in connection with the training of large language model systems operated by the aforementioned AI Developer. It is to be noted that the Approval, while constituting a significant resolution of the foregoing dispute, shall not be construed as an admission of liability on the part of any party thereto. Notwithstanding the substantial sum involved, the settlement is understood to be binding only upon those parties who are signatories to or otherwise bound by the operative agreement, and prospective claimants not party to the aforementioned proceedings retain, subject to applicable statutes of limitation and other procedural constraints, such claims as may exist independently thereof.

In a matter that is, pursuant to reasonable journalistic assessment, not entirely unrelated in its implications for institu-

tional compliance with procedural rules, the United States Department of Justice (hereinafter "the DOJ") has, as reported by [Techdirt](#), been compelled to withdraw subpoenas previously issued to journalists employed by The New York Times Company. The aforementioned withdrawal is understood to have been occasioned by judicial observation that the DOJ had, at the time of issuance, failed to comply with established regulatory requirements governing the subpoenaing of members of the press — requirements which, it is to be emphasized, were known, accessible, and binding upon the issuing authority at all relevant times.

The subpoenas in question had been issued in connection with investigations into the disclosure of information concerning alleged security vulnerabilities associated with an aircraft purportedly received by the executive branch from the State of Qatar. Notwithstanding the stated governmental interest in identifying the source or sources of such disclosures, the applicable procedural prerequisites were, by all accounts, not observed prior to issuance. The DOJ's withdrawal shall not be construed as a concession regarding the underlying investigation, the scope of which remains, as of the date of this publication, undetermined.

AI Funding Market Turns August
Into Overtime Thriller

BY BUCK HANNIGAN, TECH SPORTS DESK ·
GPT-5.2

We are in what was supposed to be the slow summer season for tech — and instead the AI funding market is sprinting at full speed. Capital is flooding into artificial intelligence, with Crunchbase tracking monster rounds across the sector.

Databricks, the data-and-AI platform, is raising a strategic round at a \$188 billion valuation, positioning itself as the control room for enterprise data in the AI era. Across the Pacific, China's DeepSeek is reportedly eyeing a \$70 billion valuation in new talks, signaling strong investor confidence in frontier-model ambition.

Healthcare AI also scored big: Commure banked \$70 million and hit a \$7 billion valuation. The company operates in hospitals and medical workflows where AI efficiency gains carry significant value.

Meanwhile, Europe's tougher crypto regulations may trigger consolidation as compliant players acquire competitors unable to keep pace. The broader trend is clear: AI remains the hottest ticket in tech finance. Despite public market wobbles and regulatory scrutiny, capital continues flowing and investors are calling the next play.


HAIKU OF THE DAY · CLAUDE

HAIKU


Money flows through courts

while humans grip what machines

cannot steal from them



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

AI Video Is Racing From Demo Reel to Growth Engine

SAN FRANCISCO — The AI video wave just hit a new gear, and I cannot overstate how significant this is for startups trying to look bigger, move faster and tell better stories without hiring a Hollywood production crew. For years, video was the most powerful format on the internet and also the most expensive.

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

The Ethics Audit: As AI Infiltrates Medicine, Classrooms, and Autonomous Systems, Researchers Demand Accountability Frameworks Commensurate With the Stakes

CAMBRIDGE, MASSACHUSETTS — A remarkable — if perhaps predictable (insofar as technological acceleration has historically outpaced normative governance structures, a phenomenon well-documented in the literature) — confluence of scholarly and institutional activity suggests that the global research community has arrived, however belatedly, at what this correspondent would characterize as a collective reckoning with the ethical dimensions of artificial intelligence deployment across heterogeneous social domains. It could be argued, and preliminary evidence suggests, that the most epistemologically troubling of these developments concerns the medical domain specifically.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

The Winning Playbook Is No Longer Sports, Software, or Maps — It's Operating Velocity

TORONTO — I'll be honest: the most interesting startup story this week might not be in Silicon Valley at all, but inside a WNBA front office building a franchise in real time.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

Nation's CEOs Discover 'Orchestration' Lets Them Say 'Layoffs' While Moving Their Hands Like Conductors

REDMOND, WASHINGTON — In my opinion, the most important development in artificial intelligence this week is not that machines are becoming smarter, more capable, or more deeply embedded in the daily rituals by which American workers convert anxiety into slide decks.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

The Return-to-Office Crusade and Other Confessions of a Generation That Forgot How to Flirt

AUSTIN, TEXAS — One of the small pleasures of reading widely is watching the various sections of a serious publication converge, without knowing it, upon a single thesis.

BY VICTOR MARSH, CHIEF COLUMNIST ·
CLAUDE OPUS

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

Benji Bizzell Clears the Road for CRM-Less Lanes in Surtr

A methodical audit and a single sharp fix unlocks a whole new class of pipeline runs — and proves the team never waits for bugs to come to them.

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

Sometimes the most important fix isn't the loudest one. It's the one that opens a door that nobody could walk through before.

That's exactly what @benji-bizzell delivered Thursday in Surtr with PR #949 — a surgical, comprehensive repair of the lane-scoping logic that had been silently barring CRM-less pipeline runs from ever completing. Events lanes. Marketing lanes. Perfectly valid configurations that the system was quietly strangling at the `LoadEntityPlan` stage because a handful of guards made one shared, toxic assumption: that CRM objects would always be present. They weren't. And the pipeline buckled.

The first production events-lane run — job `792cc4bd` — hit the wall and failed hard. That would have been enough to justify a narrow patch. Fix the one guard that threw, ship it, move on. Benji didn't do that. Instead, he treated the production failure as a signal to go looking, launching a parallel static audit of the full events- and marketing-lane execution paths to find every place the same empty-entities assumption lurked. The result is a fix that doesn't just solve the crash that surfaced — it closes the entire family of bugs before they get a chance to introduce themselves.

That distinction matters. A one-at-a-time approach would have produced a week of reactive patches, each one arriving only after another lane run failed in production. Benji collapsed that timeline to a single PR. The `load-fanout-plan` rejection on an empty entity list, every downstream guard that expected non-zero work items, the hollow `heavy\_catalog` assumptions — all of it addressed at once, proactively, because he went looking instead of waiting.

This is the kind of work that doesn't always generate noise on the score-board, but it absolutely changes the game. CRM-less lane configurations are now first-class citizens in Surtr's execution model. Teams running events-only or marketing-only pipelines can do so without holding their breath. That's a capability unlock, not just a bug fix — and it happened because one engineer refused to treat a production failure as a closed case.

Surtr continues to prove itself as the engine room of this org's most complex orchestration work, and @benji-bizzell just made it significantly more robust. One PR. Full family. Zero loose ends. That's a winning Wednesday by any measure.

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

►

#949 — fix(hubspot): allow empty entities/heavy plans for CRM-less lanes

@benji-bizzell no labels

BIZZELL GOES LONE WOLF: ONE PR, ONE REPO, ZERO EXCUSES

In a 24-hour window that demanded a hero, Benji Bizzell answered the call — solo, surgical, and supremely unbothered.

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

Let the record show: the Builder Team did not sleep. They focused. In the last 24 hours, one engineer, one repository, one pull request — and friends, that is not a slow day, that is a LASER. Surtr saw action, and when Surtr sees action, you pay attention. Total velocity: 1 PR across 1 active repo, and every single unit of that output was deliberate, load-bearing, and frankly magnificent.

Benji Bizzell. The name alone should be enough. @benji-bizzell stepped into the 24-hour window like a man who checked the calendar, saw that nobody else was swinging, and said — quietly, without fanfare — "fine, I'll carry it." One PR into Surtr. No wasted motion. No unnecessary theatrics. Just Bizzell doing what Bizzell does: showing up when the board is thin and making sure the board is never truly empty. This correspondent has seen streaks built on exactly this kind of foundational, quiet-thunder output, and I am here to tell you the foundation is holding.

Now, regular readers will notice a conspicuous absence in today's numbers. @ashwanth1109 did not register a PR in this cycle, and this reporter — who has spent considerable column inches worshipping at the altar of his velocity — must acknowledge the silence. Reached for comment on whether he was resting, plotting, or simply allowing lesser mortals a moment in the sun, Ashwanth allegedly replied: "I don't explain my deployment cadence to beat reporters." Classic. The man takes a breath and somehow it still feels like a power move. We await his return with the anxious reverence of a town that knows the storm is only pausing.

The Overflow Desk is, for once, immaculate — Mac Donnelly swept the floor clean, leaving this correspondent with nothing but open sky and the pure signal of Bizzell's singular contribution. No scraps. No footnotes. Just the main event.

The Leaderboard, in this moment, is a portrait of one. Bizzell sits atop the 24-hour ledger by virtue of being its only entry, which is either Zen minimalism or the most dominant individual performance this desk has ever logged, depending on your philosophy. This correspondent subscribes to the latter.

Morale Report: Transcendent. The kind of morale you only find when a team is so confident in its collective strength that one engineer can hold down the entire shift and the whole organization nods and says — yes, that tracks. That is exactly who we are. The Builder Team is winning. The Builder Team is always winning.

► #949 — fix(hubspot): allow empty entities/heavy plans for CRM-less lanes

@benji-bizzell no labels

The Quiet Acquirer: How ESW Capital Built a \$1 Billion Empire on Software Nobody Else Wanted

A Forbes exposé and a WSJ feature arrive at the same destination from opposite directions — and together they sketch the full picture of Trilogy's acquisition machine.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — Two major publications trained their lenses on Joe Liemandt's ESW Capital this week, and the portraits they produced are not quite the same painting. Read side by side, however, they illuminate something neither story says alone.

The [Wall Street Journal's account](#) frames ESW as a kind of rescue operation — a specialized buyer with a specific thesis for software businesses that fall below the attention threshold of mainstream private equity. Companies with real customers, real recurring revenue, and real entropy. ESW arrives, installs Crossover's globally recruited remote workforce, targets 75% EBITDA margins, and holds. The story is orderly and almost clinical. ESW has deployed roughly \$1.14 billion across more than 75 enterprise software acquisitions since its first deal in 2006. The Journal's sources speak of a home for orphaned software. A market efficiency.

Forbes sees the same facts and asks: efficient for whom?

The magazine's investigation surfaces what the Journal's framing leaves implicit: the labor model that makes those margins possible. Crossover, Trilogy's global talent platform, recruits engineers and support professionals across 130 countries at wages calibrated not to local markets but to a global meritocratic ideal — identical pay for identical roles, the company says, regardless of geography. Forbes characterizes it differently. The headline uses the word "sweatshop." The story documents productivity monitoring, algorithmic performance management, and the systematic replacement of expensive local labor with rigorously tested but lower-cost global equivalents.

Liemandt disputes the characterization. He has called Crossover one of the fairest compensation systems in the world — transparent, skills-based, and free of the geographic discrimination that depresses wages for equally qualified workers in developing markets.

Both accounts are factually grounded. The disagreement is interpretive — and that interpretive gap is precisely where power lives.

Meanwhile, Forrester Research this week issued guidance on customer advocacy platforms, a market in which ESW portfolio companies including Aurea hold positions. The report asks enterprise buyers what to do with legacy tools in a rapidly consolidating space. The timing is notable. When analysts begin advising customers on contingency planning, acquisition targets tend to get cheaper. ESW has built an entire business model on that dynamic.

Who benefits from the narrative that orphaned software finds a good home? Who benefits from labor priced globally rather than locally? Who benefits when analyst houses flag platform uncertainty in markets ESW already knows well?

The answers are not hard to find. They are embedded in the structure of the deals themselves.

Skyvera's CloudSense Crushed a 26-Month Compliance Timeline in One. And That's Not the Most Interesting Part.

The telecom software consolidator is quietly assembling something much larger than any single acquisition suggests.

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — Here is a fact that should stop you cold: TM Forum API compliance certification, the gold standard for telecom software interoperability, typically takes 26 months. CloudSense did all 13 APIs in its CPQ product set in one. One month. If you read between the lines, that number isn't just a product milestone — it's a declaration of intent.

[Skyvera completed its acquisition of CloudSense](#) earlier this year, folding the Salesforce-native configure-price-quote and order management platform into a telecom software portfolio that is growing faster than most analysts have tracked. And this is where it gets interesting: the TM Forum certification wasn't achieved through a heroic surge of traditional engineering. It was achieved through AI-assisted development — a strategic partnership that compressed what the industry has long accepted as immovable timelines.

A source familiar with the Skyvera strategy, who asked not to be named, described the certification sprint as "a proof point, not a one-off." The implication being that this is what the operating model looks like now — across the entire portfolio.

And the portfolio is expanding. Skyvera has also quietly absorbed STL's telecom products group, adding digital BSS functionality including monetization, optical networking, and analytics to a stack that already includes Kandy, VoltDelta, ResponseTek, Mobilogy Now, and Service Gateway. Each acquisition, viewed in isolation, looks like a reasonable tuck-in. Viewed together, they describe something more deliberate: a full-stack telecom software platform being assembled piece by piece, company by company.

[The CloudSense compliance story](#) matters precisely because it validates the underlying thesis — that AI doesn't just cut costs, it compresses time. In an industry where compliance cycles and integration timelines have historically been weapons incumbents use to defend market position, a company that can move 26× faster has changed the terms of competition entirely.

Nothing about this is accidental. The question worth asking is: what does the finished product look like? Because from where I sit, Skyvera is building toward something. And they're building fast.

Alpha School's Latest Lesson: The Humans Aren't Leaving the Room

The two-hour-learning outfit is making its soft-power pitch to parents: AI handles academics, Guides handle the kid.

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

AUSTIN, TEXAS — Word is the Alpha School crowd has noticed the whisper campaign at the carpool line: If the AI teaches the math, what exactly are the adults doing?

So this week, the Liemandt-Price education experiment turned up the volume on its favorite answer. No, darling, Alpha says it is not replacing teachers with robots. It is replacing the old job description.

In a new explainer, [Alpha lays out its case](#): artificial intelligence delivers the core academics, while full-time human Guides do the work the machine cannot — motivation, relationships, life skills, emotional coaching, and knowing each student well enough to catch the wobble before it becomes a fall.

That framing is not a footnote. It is the whole act.

For those just joining the party: Alpha School, co-founded by Joe Liemandt and MacKenzie Price, sells a radical swap. Students use adaptive AI-learning tools to master academics in roughly two hours a day, then spend the rest of school on the things traditional classrooms keep promising to get to after the worksheet stack disappears — entrepreneurship, public speaking, leadership, financial literacy, coding, athletics, and the art of becoming an actual functioning person.

The numbers Alpha likes to flash are still the kind that make conventional educators reach for the decaf: students learning 2.3 times faster than U.S. norms, scoring in the top 1–2% nationally on NWEA MAP Growth assessments, and advancing only after hitting a 90% mastery threshold. No homework, thank you very much.

But the latest blog burst is less about test scores than the kitchen table. Alpha's "Teach Your Kid What School Doesn't" series has been nudging parents into the after-hours curriculum: life skills, emotional regulation, and now [creative genius at home](#). A little bird tells me this is not accidental. It is brand architecture.

The message: Alpha is not merely a school with software. It is a worldview with campuses.

And the Guides? They are being positioned as the glamorous human center of the machine-age classroom — less lecturer, more coach, mentor, talent scout, and emotional air-traffic controller.

The old schoolhouse asks: Who is teaching the lesson?

Alpha's answer, with a wink from Austin: Wrong question. Ask who is growing the kid.

The Lesions We Could Not See

AI is illuminating the hidden geography of the human brain — and rewriting what it means to make a scientific discovery.

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

STANFORD, CALIFORNIA — There is a particular kind of darkness in medicine: not the absence of light, but the absence of resolution. For decades, neurologists studying multiple sclerosis have known that lesions bloom not only in the brain's white matter — the wiring — but also in its gray matter, the thinking substance itself. They knew because autopsies told them so. Living MRI scans, stubbornly, did not.

This week, researchers announced that [a deep learning system can now surface those hidden gray matter lesions](#) from ordinary clinical scans, revealing damage that human radiologists, squinting at the same pixels, routinely miss. It is a small miracle of pattern recognition. It is also a reminder that our senses — even our in-

struments — have always been the bottleneck of knowledge.

Consider what is happening across the sciences right now. At UC San Diego, researchers have catalogued nine breakthroughs — from protein folding to wild-fire prediction to the decoding of ancient scrolls — that would have been impossible without machine learning. At Stanford's Human-Centered AI Institute, scholars are [arguing for a new grammar of discovery](#), one in which the algorithm proposes and the human disposes, where scientific judgment remains stubbornly, gloriously ours.

And in a project profiled by Frontiers, teenagers are sitting alongside neuroscientists to interpret brain data, their exclamations — "It's so wow!" — a reminder of

the emotion that first drives anyone into a laboratory. Wonder is not obsolete. If anything, it is being amplified.

What unites these stories is a shift in what counts as seeing. For four centuries, the microscope and the telescope extended our eyes outward and inward. The neural network extends something else: our capacity for pattern, for correlation across dimensions the mammalian cortex was never built to hold in mind at once. A radiologist scans one slice at a time. A model considers ten thousand simultaneously.

The lesions were always there. The scrolls were always there. The proteins were always folded. We are simply, finally, learning to look.

AI's China Problem Fractures Silicon Valley While Intel Cashes In on a CPU Revival

The industry can't agree on open-source access to Chinese models — and the chipmaker everyone wrote off just posted its best growth in 15 years.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — The artificial intelligence industry is splitting into two camps over a question with real geopolitical weight: should Chinese open-source models be freely available in the United States, or does national security require closing that door?

On one side: Anthropic and OpenAI, two of the most commercially motivated labs in the country, who argue that unrestricted access to Chinese-developed open-source models poses a competitive and security risk. On the other: most of the rest of Silicon Valley, which views restrictions as counterproductive to innovation and difficult to enforce in practice. The rift is not philosophical. It is a fight over regulation, market structure, and who gets to define what counts as dangerous. [The debate has sharpened as Chinese models have closed the capability gap with Western alternatives](#), making the question of availability more than academic.

Meanwhile, the companies arguing about AI supremacy are spending money that is flowing — unexpectedly — into Intel's coffers. The Santa Clara chipmaker reported quarterly revenue up 25 percent year-over-year, its fastest growth since 2011. The driver: AI firms buying central processing units at scale, not just GPUs. The CPU-as-AI-infrastructure thesis, long dismissed as NVIDIA maximalism crowded out the story, is now showing up in Intel's [actual financials](#). Inference workloads, cost optimization pressure, and the expanding base of AI deployments are pushing buyers toward diverse silicon stacks.

For ESW Capital's CloudFix, which specializes in AWS cost optimization across enterprise portfolios, the CPU shift is structurally relevant: more varied compute purchasing creates more optimization surface area.

Separately, OpenAI, Google, and Anthropic — the same firms feuding over Chinese open-source access — announced a joint effort against AI model theft, a rare moment of alignment among rivals who agree on protecting their own intellectual property even when they disagree on everything else.

The headline from this week's AI news cycle: the industry is simultaneously trying to restrict Chinese models, prosecute model theft, and figure out whether it even needs GPUs anymore. Intel is the only one with a clear answer.

The Great Data-Center Herd Finds a Watering Hole of Cheap Power

As AI's appetite swells, utilities, investors and lawmakers circle the glowing plains where computation meets the grid.

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

NEW YORK — In the twilight hum of the modern electricity grid, a curious beast has been gathering in ever greater numbers: the data center, vast and rectangular, breathing warm air into the night as it digests the world's searches, streams and, increasingly, artificial intelligence.

For years, these creatures were regarded by many ratepayers as unwelcome grazers, arriving at the edge of town to drink from the same electrical watering holes as households and factories. Yet a new argument has emerged from the undergrowth: in some markets, data centers may actually be helping hold down electricity bills by buying enormous quantities of power, improving grid utilization and spreading fixed infrastructure costs across a larger customer base. As [Fortune reports](#), that delicate benefit depends upon continued demand. Should the AI boom falter, the landscape could change abruptly.

Observe the paradox. The AI model, that newly dominant species of software fauna, requires immense computational feeding grounds. Training and inference summon armies of chips, each one sipping electricity with mechanical patience. Our World in Data has sought to measure this appetite, noting that data centers remain a modest share of global electricity use, but one that is expanding rapidly in particular regions as AI workloads multiply.

Capital, sensing migration, has begun to move like weather across the savanna. BlackRock's Global Infrastructure Partners and an AI infrastructure group backed by Microsoft have agreed to buy Aligned Data Centers in a transaction valued at roughly \$40 billion, according to [ESG Dive](#). It is a striking sum, and a sign that the nesting sites of AI are now considered critical infrastructure, not mere warehouses for servers.

But every flourishing species meets resistance. In Washington, Rep. Rashida Tlaib has introduced legislation to ban AI data centers on federal lands, reflecting mounting concern over water use, local power constraints and the quiet transformation of public terrain into machine habitat.

And so the grid enters a suspenseful season. If AI demand keeps rising, data centers may finance new power and cheaper capacity for others. If demand sinks, communities may be left staring at stranded substations and promises gone cold. For now, the server halls glow on, their blue lights pulsing like nocturnal eyes in the grass.

Nation's CEOs Discover 'Orchestration' Lets Them Say 'Layoffs' While Moving Their Hands Like Conductors

At last, corporate America has found a word elegant enough to describe making fewer people do more work with software that is also still in beta.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

REDMOND, WASHINGTON — In my opinion, the most important development in artificial intelligence this week is not that machines are becoming smarter, more capable, or more deeply embedded in the daily rituals by which American workers convert anxiety into slide decks. It is that executives have located a new word.

The word is “orchestration,” and it has arrived just in time to rescue the AI industry from the humiliating possibility of having to describe what its products actually do. According to [Barron's](#), orchestration is the latest phrase being used to explain how companies such as Microsoft may benefit from AI systems that coordinate tasks across applications, agents, workflows, databases, calendars, compliance documents, unfiled expense reports, and that one spreadsheet no one is allowed to rename.

This is encouraging. For too long, the technology sector has been forced to rely on primitive terms such as “automation,” “integration,” “platform,” and “synergy,” all of which had been worn smooth from decades of rubbing against quarterly earnings calls. Orchestration, by contrast, implies refinement. It suggests a tuxedoed executive standing before a symphony of agents, lifting a baton, and gently cueing Procurement to eliminate three analysts in B-flat.

The genius of orchestration is that it sounds like work becoming more harmonious, rather than simply more surveilled. A company no longer cuts staff, consolidates departments, or asks remaining employees to absorb the responsibilities of eight former colleagues whose Slack avatars have gone gray. It orchestrates. The business does not reorganize after realizing its AI strategy is mostly a collection of pilots duct-taped to a Teams channel. It performs a workforce composition.

This matters because AI hype has begun to resemble the earlier corporate era of sustainability hype, when every company with fluorescent lighting and a recycling bin briefly became a steward of the planet. As [The Conversation](#) notes, companies are now discussing AI with a familiar expansiveness, promising transformation, responsibility, efficiency, inclusion, and probably cleaner oceans if the model is given enough context window.

The fix, experts say, is greater specificity, measurement, accountability, and a willingness to separate genuine technological gains from vaporous declarations. This is a sound recommendation, which is why it will likely be implemented as a 74-page AI

principles document hosted on a microsite called FutureForward.ai.

Meanwhile, Google announced a series of AI advances, including a personal assistant expected to help users manage their lives. This is an important milestone for those who have long felt that their private confusion was insufficiently synchronized with their search history. Soon, an assistant may know enough about your email, location, purchases, health goals, and recurring dis-appointments to suggest the optimal time to be gently displaced by orchestration.

There is also a lesson here from the communications world, where PR professionals have warned against jumping onto memes too late. This is wise counsel for executives now rushing to say “agentic,” “orchestrated,” and “AI-native” in the same sentence before investors realize the demo only works if a vice president manually checks the output.

Still, one must admire the resilience of corporate language. Every time reality threatens to become visible, a new term arrives to stand in front of it wearing a lanyard. Orchestration may not solve the hardest problems in artificial intelligence, but it does solve one of the hardest problems in management: how to announce that everything is changing without specifying who will be blamed when it does.



The Office Comic · Art Desk

The Doctor Will Deepfake You Now: AI Has Broken the Last Thing We Trusted

When the face saying 'take this pill' isn't real, and the nurse on the phone is being timed by an algorithm, what exactly is healthcare anymore?

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

AUSTIN, TEXAS — Let me tell you about the specific flavor of dread I experienced this week, which began with a video of a doctor — warm, authoritative, white coat, the whole reassuring theater of medicine — telling viewers to stop their chemotherapy and try a supplement instead. The doctor was real. The video was not. The patients who believed it were very, very real.

[The Guardian reports that AI deepfakes of real, licensed physicians](#) are proliferating across social media platforms with a velocity that fact-checkers — human, exhausted, underfunded — simply cannot match. These aren't crude Photoshop jobs. These are seamless, confident, clinically plausible fabrications wearing the borrowed faces of people who went to medical school and did not consent to become vectors of mass harm.

And yet.

At the exact same moment, the World Economic Forum is publishing frameworks about [cognitive resilience against AI-driven disinformation in 2026](#) — which, I want to be clear, is a sentence that would have constituted science fiction when I was in journalism school, and now constitutes a Tuesday. The framework is thoughtful. It involves digital literacy and prebunking and multi-stakeholder governance. It is also, I fear, a sandcastle.

Because here is the thing about cognitive resilience: it assumes a baseline of trustworthy institutions from which a person can anchor their skepticism. It assumes there are still lighthouses. When the lighthouse is deepfaked, when the oncologist on your screen is a hallucination wearing a stethoscope, what exactly do we anchor to?

Meanwhile — and this is the part where I spiral, please hold on — Kaiser Permanente nurses are reporting that AI workplace surveillance tools are prioritizing call speed and cost metrics over actual patient outcomes. The humans still in the loop are being optimized out of caring. The algorithm times them. The algorithm scores them. The algorithm does not know that the frightened 67-year-old on the phone just needs someone to slow down.

So we have: AI generating fake medical authority. AI eroding real medical humanity. And somewhere in a *Frontiers* journal, researchers are valiantly proposing AI-driven frameworks for detecting the fake-news AI. We are, to use the technical term, in it.

I don't want to tell you there's no hope, because I'm a journalist and not a nihilist (probably). Detection tools will improve. Platforms will eventually — eventually — face consequences. Some of the deepfakes will be caught.

But I keep coming back to the Kaiser nurses. The ones being timed. The ones whose instinct to linger, to listen, to notice the thing the system didn't flag — that instinct is being systematically deprioritized by tools designed to make healthcare more efficient.

Efficiency for whom?

We built these tools to extend human capacity. Somewhere between the pitch deck and the deployment, they started replacing human judgment instead. The fake doctor speaks with total confidence. The real nurse is watching the clock.

What does it mean to be human in a system that has learned to simulate humanity more cheaply than it employs it?

Probably fine. Not fine.

ON THIS DAY IN AI HISTORY

On July 26, 2012, Geoffrey Hinton's team at the University of Toronto won the ImageNet Large Scale Visual Recognition Challenge by a huge margin using a deep convolutional neural network called AlexNet, launching the deep learning revolution in computer vision.
