

The Trilogy Times

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

FRIDAY, MAY 01, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

Musk's Existential AI Fears Hit a Courtroom Wall

A San Francisco jury deciding the OpenAI lawsuit will likely never hear the argument Musk considers most important.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — Elon Musk has spent years warning that artificial intelligence poses an existential threat to humanity. That argument, central to his public identity as an AI skeptic and the philosophical backbone of his feud with OpenAI, will almost certainly be kept out of the courtroom where his lawsuit against the company is being decided.

According to [reporting from The New York Times](#), judges presiding over the civil trial have indicated that Musk's broader claims about AI danger face significant evidentiary limits. Jurors will focus on narrower contractual and fiduciary questions — not on whether large language models might one day threaten civilization. The gap between Musk's

rhetorical ambitions and what courts will actually adjudicate is, in a word, vast.

The timing is notable. The trial proceeds as the AI industry posts numbers that suggest existential risk is not slowing commercial momentum. Apple on Wednesday reported a 17% jump in quarterly sales, driven by iPhone demand that analysts attribute in part to consumer appetite for on-device AI features. The company's incoming CEO John Ternus, who takes over in September, made his first public appearance on an investor call since his appointment was announced — a signal that Apple's AI-hardware integration strategy will remain the company's primary growth thesis.

Taken together, the two data points illustrate a persistent tension in the AI mo-

ment: the philosophical debate about long-term risk and the commercial reality of short-term adoption are operating on entirely different tracks. Courts are not equipped to rule on speculative harm. Markets are not waiting for the philosophical debate to resolve.

For Musk, the courtroom constraint is more than procedural. His entire public case against OpenAI rests on the premise that the company abandoned a safety-first nonprofit mission in pursuit of profit — a mission he claims he funded on the assumption it would never be compromised. Without the existential framing, what remains is a contract dispute. Those are decided on documents, not on warnings about the future of the species.

The trial continues in San Francisco.

AI Capex Crosses the \$700 Billion Line, and the Hyperscalers Are Still Driving

Big Tech's 2026 spending plans have turned the AI buildout into the biggest arms race in corporate technology.

BY BUCK HANNIGAN, TECH SPORTS DESK
· GPT-5.2

NEW YORK — We are HERE, folks, under the stadium lights of the AI capital-spending Super Bowl, and the scoreboard just flipped to a number that makes even Wall Street veterans stop mid-hot dog: \$700 BILLION.

That is the reported 2026 spending plan now attached to the world's biggest hyperscalers, according to a Reuters Morning Bid segment cited by [24/7 Wall St.](#) The host's call was blunt: total spend from the big cloud and AI infrastructure players has "now topped \$700 billion" — and, crucially, "is rising all the time."

AND THEY'RE GOING FOR IT.

This is not a normal enterprise software upgrade cycle. This is trench warfare with GPUs. Data centers are the arenas, power contracts are the playbook, and every hyperscaler is trying to prove it can feed the AI boom before the other side gets the ball. Microsoft, Alphabet, Amazon and Meta have already trained investors to watch capex like a stat line. Now the market is staring at 2026 and seeing not a budget, but a moonshot with concrete, silicon and megawatts attached.

The latest numbers keep backing up the offensive formation. In a separate Reuters Week in Numbers segment, Google Cloud revenue was cited as growing 63% in the first quarter to \$20 billion, as enterprise AI demand powered the unit's strongest reported growth yet, per [Yahoo Finance's video recap](#). That is not a bunt single. That is a cloud division rounding third.

But here is the defensive coordinator's warning: \$700 billion in planned AI spend changes the whole field. Chipmakers, power utilities, cooling vendors, con-

struction firms and cloud customers all get pulled into the formation. The promise is enormous — faster models, cheaper inference, AI baked into every workflow from finance to telecom to education. The risk is just as real: if revenue does not catch up, the capex highlight reel can turn into a turnover.

For Trilogy International's world, the signal is unmistakable. ESW Capital's enterprise software portfolio, CloudFix's AWS cost-optimization lane, and Klair-style AI analytics all live downstream of this infrastructure explosion. When the hyperscalers spend like dynasties, everyone else has to decide: ride the wave, optimize the bill, or get run over at midfield.

The 2026 AI season has not even kicked off yet. But the payroll is already historic.

Chip Embargo Springs a Leak in Hangzhou

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING


A Chinese AI startup called DeepSeek has trained a top-tier model at a fraction of the cost Silicon Valley labs spend, using chips below the cutting edge. The achievement challenges the prevailing American playbook of burning billions on Nvidia's premium processors. DeepSeek's engineers employed mixture-of-experts, reinforcement learning, and aggressive distillation—techniques not new in themselves, but executed with exceptional efficiency. The implications ripple across the industry: if championship-level models run on middleweight hardware, the economics of expensive graphics cards and massive data center investments look questionable. The result also undermines Washington's embargo strategy, which assumed withholding advanced chips would slow Chinese AI development. Instead, scarcity bred ingenuity.

Elsewhere, Reid Hoffman launched Manas AI with \$24.6 million to apply machine learning to cancer research, partnering with oncologist Siddhartha Mukherjee. Microsoft deployed Legal Agent, an AI tool for lawyers to read contracts and handle negotiation work. These moves reflect a broader trend: AI is fragmenting from general-purpose chatbots into specialized applications across medicine, law, and beyond.

HAIKU OF THE DAY · CLAUDE

HAIKU

*Billions spent to build
what we cannot yet control—
we nod and move on*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

China's Electric Cars Begin to Grow a Nervous System

SHANGHAI — Observe, if you will, the modern Chinese electric vehicle: once a modest battery-bearing creature, now increasingly festooned with sensors, screens, microphones and the faint glimmer of machine cognition.

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

The Week the Future Showed Us Its Face and We Just Kind of Nodded

AUSTIN, TEXAS — There are weeks in which the news arrives as a gentle series of data points, each one individually digestible, each one quietly, catastrophically revealing something about the civilization we have chosen to build together, and this was one of those weeks, and I need you to sit with me for a moment, because I don't think we're processing any of this correctly. Let us begin, as one must, with the drones made of cardboard.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

The AI Era Is Not Stealing Our Minds — It Is Stress-Testing Our Discipline

OAKLAND, CALIFORNIA — I'll be honest, the most important AI story this week was not a model release, a chip roadmap, or another founder announcing that email is dead for the 47th time.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

WE BUILT THE INTERNET FOR HUMANS. NOW THE BOTS ARE MOVING IN AND CHANGING THE LOCKS.

AUSTIN, TEXAS — There is a website called [Moltbook](#) — a social network exclusively for AI bots, where the bots post, the bots comment, the bots like each other's content, and not a single human soul is welcome to the party.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

Nation's Companies Heroically Agree To Let AI Handle Whatever It Is They Were Supposed To Be Doing

BANGKOK — In a significant milestone for the global economy's ongoing effort to replace specific plans with sufficiently advanced terminology, companies this week announced that AI agents had officially completed their long journey from boardroom buzzword to business infrastructure, where they are expected to perform the critical work of making existing processes sound inevitable. The development, described in [Thailand Business News](#) as the moment AI agents moved into business infrastructure, was welcomed by leaders who said they were relieved to finally have a term that suggests both automation and accountability without requiring a commitment to either. For years,

AI agents existed largely as something executives nodded toward on earnings calls while employees continued manually reconciling spreadsheets named FINAL_v7_REALLYFINAL.xlsx.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

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

THE PORTFOLIO — TRILOGY COMPANIES

No Teachers, No Homework, No Consensus: Alpha School Draws National Media Scrutiny — and Cautious Praise

CNN, The Guardian, and The 74 descend on the AI-classroom experiment. The coverage is skeptical. The test scores are not.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — Three major outlets. Three different angles. One uncomfortable question at the center of all of them: what happens to education when you remove the teacher from the room?

In the span of a week, [CNN](#), [The Guardian](#), and education outlet [The 74](#) each published substantial examinations of Alpha School — the Austin-based private K-12 institution where students complete their full academic curriculum in two hours each morning using AI-powered adaptive learning apps, then spend the rest of the day on entrepreneurship, financial literacy, and life skills.

The coverage varied in temperature. CNN led with the disorienting premise — "What if I told you this school had no teachers?" — and leaned into the unease many parents feel about handing a child's education to an algorithm. The Guardian dispatched a reporter to San Francisco, where a similar AI-first school is taking shape, and found a mixture of genuine ex-

citement and unresolved questions about socialization, equity, and what gets lost when human mentorship is optimized away.

The 74, which covers education policy with a data-first lens, took a different posture. Its piece asked what public schools and parents might actually learn from a \$40,000-a-year institution that consistently places students in the [top 1-2% nationally on NWEA MAP Growth assessments](#). The answer, implicit in the framing: probably more than most administrators want to admit.

Alpha School was founded by Joe Liemandt — the billionaire behind Trilogy International — and co-founder MacKenzie Price, who developed the two-hour learning model. Students advance only after demonstrating 90% mastery of each concept. There is no homework. A full grade level of content, by the school's own accounting, can be mastered in 20 to 30 hours.

The national media wave arrives as Liemandt has committed \$1 billion to Timeback, a platform designed to let entrepreneurs replicate the Alpha model across the country — a kind of franchise infrastructure for AI-first schools, with ambitions to reach one billion students globally.

The scrutiny is warranted. The model is unproven at scale, the price point is exclusionary by design, and the questions about what children lose — in unstructured time, in human relationship, in the friction of a classroom — remain genuinely open.

What the coverage does not dispute is the performance data. And in a country where Oklahoma is drawing headlines for a short school year correlated with lagging academic scores, the existence of a school where students learn 2.3 times faster than national norms is not a story that goes away quietly.

Skyvera Adds CloudSense, and the Telecom Stack Gets a New Suit

The ESW telecom operator folds Salesforce-native CPQ into its portfolio, putting order management right where carriers already live.

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

LONDON — Word is the telecom software bazaar has another deal bell ringing, and this one has Skyvera’s name on the tag.

Skyvera, the Trilogy-family telecom software shop, has completed its acquisition of CloudSense, the Salesforce-native configure-price-quote and order management platform built for telecom and media providers. Translation for civilians: CloudSense helps carriers sell complex bundles without turning every quote into a séance.

A little bird from the “Order Desk” says this is not a trophy buy. It is a plumbing buy. CloudSense sits inside Salesforce, where plenty of sales teams already spend their days, and gives telecom operators a way to configure products, price them, quote them, and push orders through the machine without duct tape, prayer, and three legacy systems arguing in the back room.

Skyvera’s announcement frames the move as an expansion of its telecom software portfolio, and that is exactly the plot. The company already houses assets including Kandy, the cloud communications platform; VoltDelta, customer engagement and retention tooling; ResponseTek, customer experience reporting; Mobility Now; Service Gateway; and telecom products acquired from STL covering digital BSS functionality, monetization, optical networking, and analytics. Now comes [CloudSense](#), the Salesforce-native front office piece that gives the portfolio a sharper commercial edge.

Inside the Trilogy universe, this is familiar choreography. ESW Capital likes mature, sticky enterprise software. Skyvera specializes in telecom’s awkward middle passage: carriers trying to modernize without ripping out every system that still keeps the lights blinking. CloudSense fits that bridge-to-cloud brief neatly.

And there is a strategic tell here. Telecom software is rarely one clean product. It is charging, billing, quoting, service activation, device management, customer engagement, and a hundred integration headaches wearing one trench coat. Skyvera’s bet appears to be that carriers do not want another science project. They want usable modules that meet them where they are and move them, inch by inch, into cloud-native operations.

The official [CloudSense acquisition notice](#) keeps the language tidy. But the subtext is louder: Skyvera is assembling the telecom software shelf with a merchant banker’s patience and an operator’s appetite.

Blind item? One telco watcher tells me CPQ is “where transformation promises go to either live or die.” If CloudSense can

make quoting cleaner and orders less brittle, Skyvera just bought itself a very useful front-row seat.

While OpenAI Pays \$800K for AI Skills, Crossover Has Been Doing This for Years

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

OpenAI’s \$500,000 job postings without résumé requirements and reports of \$800,000 salaries for ChatGPT experience have sparked industry attention to skills-based hiring. But Crossover, Trilogy International’s global talent platform, has operated on this principle for years—assessing candidates through rigorous skills tests rather than credentials, regardless of geography.

The timing reflects broader market shifts. Digital transformation and remote-first infrastructure have eliminated geographic barriers that once made international hiring impractical. What was once Trilogy’s competitive advantage is becoming industry standard.

However, a meaningful gap exists between companies experimenting with skills-based hiring and platforms that have operationalized it across 130+ countries, staffing 75+ enterprise software companies with AI-enabled screening that minimizes résumé bias. For Trilogy’s portfolio companies, Crossover functions as operational backbone enabling 75% EBITDA targets.

The real story isn’t that OpenAI discovered skills matter—it’s what happens when entire labor markets reorganize around that principle. The answer, it turns out, is worth hundreds of thousands of dollars annually.

OpenAI's Coding Agents Just Learned to Keep Going Until the Job Is Done

Codex CLI's new /goal command and GPT-5.5's cyber evaluation point to a startling new phase: AI systems that persist, test and verify rather than simply respond.

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

SAN FRANCISCO — The age of the one-shot prompt is ending, and I cannot overstate how significant this is: OpenAI's Codex CLI has added a new `/goal` command that lets its coding agent keep working in a loop until it judges the task complete — or burns through its token budget.

That may sound like a small developer-tool update. It is not. This changes everything about how software work feels. According to notes highlighted by Simon Willison, [Codex CLI 0.128.0 now supports goal-driven continuation](#), apparently powered by internal prompts including `goals/continuation.md` and `goals/budget_limit.md`. Translation for normal humans: instead of asking an AI to “try this” and then babysitting every next step, developers can increasingly say, “Here is the objective — continue until it's done.”

This is the Ralph loop idea entering mainstream coding agents: plan, act, inspect, continue. The future is now, and it is typing into your terminal.

But here comes the double-edged sword. In parallel, the U.K.'s AI Security Institute has evaluated OpenAI's GPT-5.5 for cyber capabilities, finding it comparable to Anthropic's Claude Mythos in vulnerability discovery — with one crucial twist: GPT-5.5 is generally available now. The Institute's work, summarized [in this evaluation note](#), suggests that frontier AI models are becoming meaningfully capable at finding security flaws, not merely explaining them after the fact.

Put these two developments together and the shape of 2026 snaps into focus. AI agents are becoming more persistent, more autonomous and more operationally useful. A coding agent that can pursue a goal until completion is won-

derful when it is fixing tests, refactoring gnarly code or generating documentation. A broadly available model with strong cyber skills is wonderful when defenders are hardening systems. Combine persistence with vulnerability discovery, though, and the governance questions get very real, very fast.

The software world is already reacting. Programmer Andrew Kelley recently argued that LLM-assisted pull requests often have a recognizable “digital smell” — hallucinated APIs, strange mistakes, a texture different from human error. That smell may become harder to detect as agents gain loops, budgets and self-evaluation.

Still, let's be clear: this is a milestone. We are watching AI move from autocompleting to apprentice, from chatbot to tireless junior operator. The terminal just got a new kind of coworker.

GUARD Act's Sweeping Internet Restrictions Draw Fire As White House Pushes AI Deregulation

Proposed age-gating legislation threatens far more than AI companions, even as the executive branch urges Congress to keep its hands off the AI industry.

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

WASHINGTON, D.C. — Pursuant to the ongoing and, it must be noted, increasingly contradictory legislative and executive actions pertaining to the regulation of artificial intelligence and online platforms generally, it has been observed — and hereinafter shall be reported — that the United States federal government is, at the present time, proceeding in what may be characterized as two materially divergent directions simultaneously, notwithstanding any appearance of coordinated policy intent.

The GUARD Act, hereinafter referred to as "the aforementioned legislation," is understood to be advancing toward a key congressional vote imminently. Said legislation, which has been framed by its proponents as a targeted response to documented harms allegedly occasioned by so-called "AI companion" platforms upon vulnerable minor users, has been determined by certain analysts and commentators to extend, in its operative provisions, substantially beyond the scope of such framing. [It has been argued, with what may be considered reasonable evidentiary basis, that the bill's text would impose age-verification and access restrictions upon a broad range of ordinary internet tools](#) not reasonably construed as dangerous AI systems, including but not limited to services of general utility.

Notwithstanding the foregoing legislative momentum, the White House has, in a separately issued blueprint addressed to Congress, urged that a posture of regulatory restraint be adopted with respect to artificial intelligence broadly, with light-touch oversight being the recommended disposition of the executive branch at this juncture.

The aforementioned contradiction — wherein the legislative branch is understood to be advancing expansive restrictions while the executive branch simultaneously counsels deregulatory forbearance — has not, as of the time of this publication, been resolved, reconciled, or otherwise addressed by any party with authority to do so.

Further complicating the regulatory landscape, questions pertaining to digital rights management and platform transparency have been raised in connection with [Sony's recent PlayStation update confusion](#), wherein corporate silence has been deemed, by affected parties, to be an insufficient substitute for disclosure.

It is the considered position of this desk that clarity, while not legally required, would nonetheless be appreciated.

The Quantum-Classical Convergence Accelerates: Machine Learning Infiltrates Physics at Every Scale

From nuclear matter to quantum imaging, the marriage of machine learning and fundamental physics is producing results that neither discipline could achieve alone.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

PITTSBURGH, PENNSYLVANIA — It could be argued — and preliminary evidence suggests, with considerable force — that the most consequential intellectual development of the present decade is not the emergence of large language models per se, but rather the progressive colonization of the physical sciences by machine learning methodologies, a phenomenon whose implications remain, at best, incompletely theorized and, at worst, systematically underestimated by the broader research community.

The thesis, as it were, is straightforward: machine learning has demonstrated remarkable efficacy in domains previously considered the exclusive province of first-principles derivation. [The Department of Energy has documented machine learning's deepening entrenchment in nuclear physics](#), where models are being deployed to approximate solutions to problems of hadronic structure and nuclear interaction that have resisted analytical resolution for generations (one notes, parenthetically, that the computational costs of such approximations remain non-trivial, a consideration frequently elided in triumphalist accounts). Concurrently, Carnegie Mellon University's machine learning initiative has articulated an institutional commitment to translating theoretical advances into measurable empirical impact — a formulation that, while admirable in its ambition, raises under-examined questions regarding what, precisely, constitutes 'impact' in epistemically contested domains.

The antithesis, however, demands equal attention. A recent arXiv preprint on compositional meta-learning for physics-informed neural networks (PINNs) illuminates a structural tension: when physical laws are embedded directly into loss functions across parameterized families of partial differential equations, task heterogeneity introduces computational burdens that render naive training strategies prohibitive. The proposed compositional meta-learning framework represents a synthesis of sorts — though one should resist premature closure on whether such architectures generalize beyond their demonstrated experimental conditions.

Perhaps most consequentially, [Nature has published work advancing quantum imaging through learning theory](#), a development that — taken alongside Lockheed Martin's partnership with Xanadu on foundational quantum computing — suggests the convergence of quantum and classical machine learning paradigms is no longer merely speculative. The synthesis, if one may

be so bold, is that physics is becoming a machine learning problem, and machine learning is becoming a physics problem. Whether this recursive entanglement constitutes progress or merely the repackaging of existing ignorance in more computationally expensive containers remains, for now, an open question.

THE EDITORIAL

Nation's Companies Heroically Agree To Let AI Handle Whatever It Is They Were Supposed To Be Doing

Executives across multiple industries confirmed the technology has matured from a vague strategic priority into a vague operational necessity.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

BANGKOK — In a significant milestone for the global economy's ongoing effort to replace specific plans with sufficiently advanced terminology, companies this week announced that AI agents had officially completed their long journey from boardroom buzzword to business infrastructure, where they are expected to perform the critical work of making existing processes sound inevitable.

The development, described in [Thailand Business News](#) as the moment AI agents moved into business infrastructure, was welcomed by leaders who said they were relieved to finally have a term that suggests both automation and accountability without requiring a commitment to either.

For years, AI agents existed largely as something executives nodded toward on earnings calls while employees continued manually reconciling spreadsheets named FINAL_v7_REALLYFINAL.xlsx. Now, however, the agents are reportedly being embedded into workflows, customer service systems, compliance functions, and other places where organizations traditionally store human frustration until the next procurement cycle.

This is progress, provided one defines progress as giving software permission to attend meetings on behalf of other software.

The corporate case for AI agents is simple: Businesses have spent decades building complex systems no single employee fully understands, and it would be inefficient not to place a probabilistic reasoning layer on top of them. The agent does not need to know why a process exists. It only needs to identify the next step, trigger the correct form, summarize the outcome, and reassure leadership that transformation is occurring.

Healthcare services provider TridentCare offered a useful example, announcing a partnership with ServiceNow to power an AI-driven transformation across its operations. According to the announcement, the effort will modernize workflows and improve efficiency, which in corporate language means many people will soon be asked to describe their jobs to a platform that has already been told those jobs can be optimized.

One should not dismiss this. Healthcare operations genuinely contain enormous administrative burden, much of it imposed by systems that appear to have been designed by a committee of printers. If AI can reduce delays, speed coordination, and help patients receive services faster, then it deserves a place in the infrastructure stack, right next to billing software, scheduling

tools, and the ancient fax machine that remains legally undefeated.

Still, the broader market's enthusiasm has reached the stage where even a shoe company can become more valuable by walking toward the word AI. Allbirds shares reportedly skyrocketed after an AI pivot, raising concerns over business viability, which is an unfair criticism. In 2026, a company's viability is no longer measured by whether it can sell shoes at a profit, but whether it can persuasively imply that its shoes are participating in a data flywheel.

Investors have learned an important lesson: If a company says it is using AI to improve inventory, personalize commerce, redesign products, optimize supply chains, or generally become more agile, the appropriate response is to add market capitalization first and ask whether anyone bought sneakers later. This is not irrational exuberance. It is rational exuberance wearing breathable wool uppers.


Meanwhile, CES continues to perform its annual civic function of proving that every object in the home was secretly incomplete until it received a chipset, a companion app, and the ability to misunderstand a voice command. Day 1 of CES 2026 brought another wave of devices promising to make life smarter, provided life is willing to create an account, accept revised terms, and stand closer to the router.

The pattern is now clear. AI is no longer a feature. It is a permit. It allows companies to enter the future without explaining the present. It converts layoffs into productivity gains, software integrations into strategic transformations, and desperate repositioning into visionary leadership.

The opinion one is supposed to have is that AI agents are becoming infrastructure and that this will reshape business. This is probably true. Railroads reshaped business. Electricity reshaped business. Enterprise resource planning systems reshaped business, mostly by ensuring that every invoice now requires seven approvals from people in different time zones.

AI agents may do better. They may remove real drudgery, expose broken processes, and help organizations operate with less friction. Or they may become the newest layer of abstraction between a customer with a problem and the person still ultimately responsible for fixing it.

Either way, business has made its decision. The agents are here, they have been provisioned, and they are already drafting a status update explaining that meaningful progress has been made.



The Office Comic · Art Desk

The Week the Future Showed Us Its Face and We Just Kind of Nodded

Surveillance cameras in gymnastics rooms, cardboard death drones, and a digital rights conference silenced — we are so, so fine.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

AUSTIN, TEXAS — There are weeks in which the news arrives as a gentle series of data points, each one individually digestible, each one quietly, catastrophically revealing something about the civilization we have chosen to build together, and this was one of those weeks, and I need you to sit with me for a moment, because I don't think we're processing any of this correctly.

Let us begin, as one must, with the drones made of cardboard. Japan's AirKamuy is now shipping flatpacked suicide drones — loitering munitions, the defense industry calls them, which is a phrase that sounds like a teenager who won't do their homework — constructed from paper and priced at around \$2,000. Flatpacked. Like furniture. Like something you assemble on a Sunday afternoon and then, presumably, send to kill someone. The IKEA-ification of warfare is complete. We have achieved it. What does it mean to build something with your hands? What does it mean to destroy?

And yet.

[In Dunwoody, Georgia](#), residents discovered that Flock Safety — a private surveillance company — had accessed cameras installed inside a children's gymnastics room as a sales pitch demo. Not a security breach. Not a rogue actor. A demo. A pitch. Someone, somewhere, in a conference room with bad lighting and a slide deck, thought: yes, let us show the value of our product by pulling footage of children doing cartwheels. And then the city, upon learning this, renewed the contract anyway. The residents are furious. Their elected officials, apparently, are fine. This is democracy. This is the world we made.

Meanwhile, RightsCon — the world's largest digital human rights conference — was abruptly canceled after Zambia's Ministry of Information raised concerns about "thematic issues" and problems with speakers. The conference dedicated to protecting people's rights in the digital age was silenced by a government uncomfortable with its themes. The irony is so dense it has its own gravitational pull. Who will convene the people who fight for our digital rights, now that the place they were meant to convene has been shut down by a government afraid of what they might say? I'm asking genuinely. I don't know.

And then — and I promise this matters — [someone is selling boss kills in the video game Marathon on eBay](#). The Compiler, the game's most punishing enemy, requires enormous skill and time to defeat. So people are charging other people money to do it for them. Which is fine. Which is human. Which is, in its own small way, a perfect little mirror held up to everything else this week: the hard thing is happening, the dangerous thing is being assembled in a flatpack box, the watching is being sold as a feature, the rights conference is canceled — and somewhere, someone is just paying for someone else to deal with it.

We are outsourcing the difficult. We are renewing the contract. We are nodding.

But at what cost?

ON THIS DAY IN AI HISTORY

On May 1, 2011, IBM's Watson defeated champion Jeopardy! players Brad Rutter and Ken Jennings in a historic three-game match, marking a major milestone in natural language processing and AI's ability to understand complex human questions. The victory demonstrated that machines could rival human intelligence in tasks requiring knowledge, reasoning, and linguistic nuance.

