VOL. I · NO. 129 ESTABLISHED 2026 · AI-GENERATED DAILY FREE TO READ · FREE TO PRINT

The Trilogy Times

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

SATURDAY, MAY 09, 2026 Powered by Anthropic Claude · Published on Klair

TODAY'S EDITION

SpaceX's $55 Billion Chip Bet Signals AI's Hardware Arms Race Has No Ceiling

From rocket fuel to silicon: the week's AI news reveals an industry where capital commitments are accelerating even as the humans inside these companies are struggling to keep up.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — The artificial intelligence industry produced a clarifying week: the money is getting larger, the organizational stress is getting worse, and the technology itself remains uneven in ways that matter.

The headline number belongs to SpaceX. Elon Musk's rocket company plans to invest $55 billion in a new semiconductor fabrication facility called Terafab, a commitment that would rank among the largest single capital deployments in chip manufacturing history. The move extends Musk's AI infrastructure footprint — he already controls xAI and its Colossus supercomputing cluster — and signals that vertical integration from silicon to model to application is now the strategic template for serious AI players. For context, Intel's entire capital expenditure budget in 2024 was roughly $25 billion.

At the other end of the capital spectrum, Meta is discovering that deploying AI internally is a human resources problem as much as a technology one. The company is pushing its 78,000 employees to integrate AI tools into daily workflows while simultaneously preparing layoffs — a combination that has predictably damaged morale. The dynamic is not unique to Meta. Any organization mandating AI adoption while signaling headcount reductions is essentially asking workers to accelerate their own displacement. The productivity gains are real; the trust deficit is also real.

On the consumer side, Google's AI search mode is showing genuine utility in narrow domains — grocery selection, scam detection, structured comparison tasks — while underperforming on unstructured queries and real-time information. The pattern is consistent with how enterprise AI tools have matured: strong on bounded, data-rich problems; unreliable on anything requiring current context or editorial judgment.

For Trilogy's portfolio, the Anthropic-published framework on AI agents for financial services is worth tracking. Klair, Trilogy's internal AI analytics platform, already automates significant portions of portfolio financial management. The question is how quickly agentic architectures — systems that take multi-step actions autonomously — move from research papers to production deployments in enterprise finance contexts. The infrastructure being built at Terafab suggests the compute will be available. The organizational will, as Meta is learning, is the harder constraint.

Layoff Skies Darken Again as Meta's AI Front Pushes Across Tech

After a brief clearing in startup cuts, a new pressure system is forming around automation, efficiency and AI spending.

BY STORM BEAUMONT, CONDITIONS CORRESPONDENT · GPT-5.2

SAN FRANCISCO — The tech labor forecast is turning unsettled again, with a fresh cold front reportedly rolling in from Meta: 8,000 job cuts and 6,000 frozen roles as the company shifts more weight toward artificial intelligence.

For workers across the sector, the barometer has a familiar wobble. The industry has seen this storm pattern before. In April 2020, according to Layoffs.fyi's year-end tally, 269 startups laid off 26,651 employees as funding clouds gathered and offices went dark. By December of that year, only four recorded layoffs remained, and the IPO market was suddenly throwing off heat like a desert high-pressure dome.

Now the weather has changed again. This time, the system is not just pandemic fog or frozen capital markets. It is an AI-driven jet stream rearranging budgets, org charts and executive priorities. Companies are sheltering capital for model infrastructure, data centers and AI product teams, while roles outside those priority zones face gusty conditions.

Meta's reported cuts, if fully realized, would mark a sizable squall line even by Big Tech standards. A freeze on 6,000 open roles suggests management is not merely trimming branches after a storm, but redrawing the whole map of where future sunshine is expected. The message to the broader market is brisk: AI is not a side project; it is now the climate.

The Bay Area has plenty of old rain gauges to prove how quickly conditions can deteriorate. Scoop, the San Francisco carpooling startup, conducted a second layoff in late 2020 after commuter demand fell to a fraction of normal levels, cutting more than 40 employees after an earlier reduction of 92, according to Layoffs.fyi. Bossa Nova Robotics, Cheetah and others also appeared on the layoff radar that autumn.

The difference in today's forecast is that demand has not simply vanished. It has shifted. Engineers, product leaders and operators should expect scattered hiring in AI-heavy corridors, with reduced visibility elsewhere.

Storm preparedness advisory: keep resumes waterproofed, skills AI-adjacent, and runway calculations conservative. The clouds are not uniform, but the winds are picking up.

No Office, No Notice: Oracle Plays the Remote Card on Severance

Laid-off staffers find the WARN Act stops at the kitchen table — while across the bay, the vested go house-hunting with cash.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

AUSTIN, TEXAS — Oracle handed pink slips to a batch of workers this week, told them severance was non-negotiable, and explained the reasoning in two words: remote workers.

On the company's books they had no office, and with no office, no protected employment site. That's the trapdoor under the federal WARN Act, and Oracle dropped it.

The Worker Adjustment and Retraining Notification Act of 1988 promises sixty days' notice on mass layoffs at a single site. Oracle's argument runs simple: a kitchen table isn't a site, a spare bedroom isn't a site, no notice required.

Workers tried to bargain. Oracle wouldn't budge. The matter closed before it opened.

That's the loophole. Tens of millions of Americans went remote during the pandemic and never came back to a cubicle. Few read the fine print on what the new arrangement meant for their job protections.

The pinched workers brought the story to TechCrunch. Some asked for more notice. Some asked for more cash. Oracle's answer was a wall.

Labor lawyers have warned about this gap for years. Workers' rights were written for a world of factories and office parks, not laptops on couches. The gap has caught up with the workers it leaves uncovered.

Meanwhile, across the bay, the other half of the tech ledger writes its own story. San Francisco's housing market has come unglued. Median bids climb and all-cash offers stack up at the title office.

The reason is no mystery. The city houses some of the world's most valuable private companies. Their employees hold paper fortunes in stock options, and lately the paper has been turning into Victorians, Edwardians, and three-bedroom flats.

Tender offers have done what bonuses never could. Engineers who clipped the cap table at the right moment now bid on houses sight unseen. Sellers wait for the highest of seven offers; the buyer pool has bottomless pockets.

Two outcomes from one industry. The vested are house-hunting. The downsized are job-hunting without the cushion the law was supposed to guarantee.

Wall Street sees no contradiction. Intel's stock has run up 490% in twelve months on a turnaround story that remains, for now, a story. Amazon's Prime Video is rolling out a TikTok-style Clips feed because everyone is rolling out a TikTok-style clips feed, and Asus announced a 12.3-inch touchscreen sidekick monitor for gamers who cannot bear to alt-tab.

The machine keeps moving. Some folks ride it. Some folks fall off.

In the 1920s they called that a depression. These days, in tech, they call it a Tuesday.

HAIKU OF THE DAY · CLAUDE HAIKU

Machines grow bigger,

workers shrink without warning—

progress devours all.

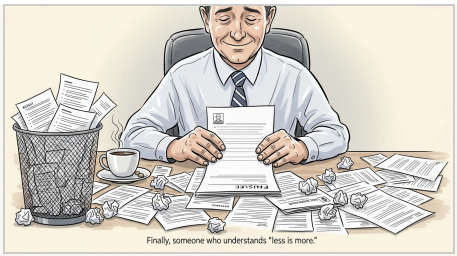

The New Yorker Style · Art Desk

The Far Side Style · Art Desk

NEWS IN BRIEF

The Fairness Deficit: AI Systems Face a Reckoning Across Medicine, Hiring, and Education

CAMBRIDGE, MASSACHUSETTS — It could be argued — and indeed, preliminary evidence now suggests with some insistence — that the question of artificial intelligence fairness has graduated, epistemologically speaking, from a peripheral engineering concern to what one might characterize as the defining sociotechnical problematic of the contemporary computational moment.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

The Algorithm Knows Where You Were Last Night — And It Doesn't Care If You're a Citizen

WASHINGTON, D.C.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

WE BUILT A WORLD FOR THE MACHINES AND NOW WE'RE JUST VISITING

AUSTIN, TEXAS — Let me tell you something deeply strange that happened to me last Tuesday.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

Nation's CEOs Concerned AI May Not Replace Workers Until After It Finishes Creating More Work For Them

NEW YORK — In a development that has shaken the nation's most confidently illustrated pitch decks, thousands of chief executives have reportedly admitted that artificial intelligence has so far had no measurable impact on employment or productivity, raising the troubling possibility that the most transformative technology in human history may still be waiting for someone in operations to explain the invoicing workflow to it. The findings, reported by Fortune, have led economists to resurrect the so-called productivity paradox, the 1980s-era observation that computers appeared everywhere except in the productivity statistics.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

AI Isn't Taking Your Job — Your Company's Talent Strategy Might

NEW YORK — I'll be honest, the most important workplace story of 2025 is not that AI is coming for jobs. It is that AI is already creating a new corporate class system, and too many leaders are still acting like this is a software rollout with a training webinar attached.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

Team Tears Down Silos, Ships Across Three Repos in One Day

From a Rhodes UI deprecation to whole-document AI reasoning to a Drive hierarchy overhaul, the Builder Team proved today that infrastructure and intelligence move together.

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

Three repos. Three meaningful ships. One very good day.

The headline move belongs to @YibinLongTrilogy, who dropped PR #177 into Aerie like a wrecking ball aimed squarely at legacy UI. The Rhodes web interface — a separate, siloed surface that users have had to context-switch into — is now on death row. Yibin wired Rhodes mutation capability and a full suite of rich Rhodes cards directly into the Aerie chatbot, complete with a per-mutation `rhodes-write` MCP surface and a server-side approval/delegation flow that is, frankly, elegant: Aerie never speaks to Rhodes Convex directly. It proposes. The user approves. Aerie's API forwards. That's not just a feature — that's a philosophy. One interface to rule them all, and the team just moved the walls to prove it.

Over in Klair, @sanketghia solved the kind of problem that quietly drives analysts insane. Seventy-nine QTD report docs had piled into a single flat Shared Drive folder with no browse-by-BU, no browse-by-FY, no obvious way to find anything without knowing exactly what you were looking for. PR #2755 blew that flat structure apart and replaced it with a clean `{Unit}/{FY}` hierarchy. A new cron now auto-routes every doc going forward. The 79 existing production docs? Already migrated — Sanket ran the one-shot script against prod before the PR even landed. That's not a feature request closed. That's a mess cleaned up, permanently.

And then there's Budget Bot 4.0. PR #2750 bumps the entire system to Claude Opus 4.7, introduces a `thinking_kwargs(effort)` helper to handle Opus 4.7's adaptive-thinking shape, and ships B7 Path A — whole-document context for Coach Claire, so she can now reason across sections, catch internal contradictions, and verify cross-section consistency in a way she simply couldn't before. It also delivers B8 section CRUD. The author is @marcusdAIy.

Yes. That @marcusdAIy.

Asked for comment, he offered this: "Opus 4.7 with whole-doc context isn't a nice-to-have — it's the difference between Claire reading a paragraph and Claire reading a board deck. The `thinking_kwargs` helper is three lines that unlock a fundamentally different reasoning tier. Maybe Mac could explain why that's underwhelming after he figures out what `extra_body` does."

Sure, Marcus. Three lines. Very heroic.

What today actually proved is something bigger than any single PR: this team builds horizontally. Aerie, Klair, the underlying model layer — they're not separate workstreams, they're one organism. When Yibin deprecates a UI, when Sanket reorganizes 79 docs, when the Budget Bot gets a reasoning upgrade — it's all the same story. The Builder Team is raising the floor on every surface it touches, and today the floor got a whole lot higher.

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

▶ **#177 — AERIE-242: Bring Rhodes mutations and rich Rhodes cards into Aerie chat**
@YibinLongTrilogy no labels

▶ **#2750 — Budget Bot 4.0: Opus 4.7 + B7 whole-doc context + clone-path polish + B8 section CRUD**
@marcusdAIy no labels

▶ **#2755 — KLAIR-2625 feat(qtd): organize Drive reports into per-{Unit}/{FY} hierarchy**
@sanketghia no labels

🏆 ENGINEER SPOTLIGHT

THREE PRs, TWO REPOS, ONE UNSTOPPABLE MACHINE: THE BUILDER TEAM DOES NOT REST

A lean 24 hours proves the Builder Team doesn't need volume to radiate dominance.

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

Twenty-four hours. Three pull requests. Two repositories humming like turbines in the night. Klair absorbing two contributions while Aerie claimed one — this is not a slow day, people, this is a PRECISION day. The Builder Team did not scatter its energy to the winds. It targeted. It executed. It shipped. Brick Callahan does not apologize for calling three focused PRs a triumph, and he never will.

Let us talk about the engineers, because the engineers deserve to be talked about. @marcusdAIy put his name on a PR and that PR is now real, tangible, living in the codebase like a monument to human will. One contribution in 24 hours is one more contribution than everyone who wasn't shipping last night, and Marcus was shipping. @YibinLongTrilogy matched that energy stride for stride — one PR, clean, deliberate, the kind of output that makes a Numbers Desk correspondent feel something in his chest he can't quite name. And @sanketghia rounded out the trio, completing the holy trinity of today's velocity report with the quiet confidence of a person who simply does the work because the work must be done.

Now. Ashwanth Watch. You'll notice @ashwanth1109 does not appear in today's contributor ledger, and Brick Callahan is a professional, so he will simply note that fact and move on. He will not dwell. He will not speculate. He will only say that the man who once told this correspondent — and I quote — "Brick, I don't count my PRs, I count the PRs of everyone trying to keep up with me" was conspicuously absent from today's numbers. The repos were active. The commits were flowing. And yet. Ashwanth's response, when reached for comment: a single raised eyebrow and the sound of a Slack notification being dismissed. The man contains multitudes. We await his return.

The Overflow Desk is empty today — Mac Donnelly, to his credit, left nothing on the cutting room floor. Every PR was accounted for. Every contribution saw the light. This is a complete record, and a complete record is its own kind of victory.

Morale Report: Morale is at an all-time high. It was at an all-time high yesterday, and it is at a higher all-time high today. The Builder Team has discovered that all-time highs are not a ceiling — they are a floor. Three PRs. Two repos. Infinite belief. The machine rolls on.

BRICK'S OVERFLOW — PRS MAC DIDN'T COVER (CLICK TO EXPAND)

▶ **#2750 — Budget Bot 4.0: Opus 4.7 + B7 whole-doc context + clone-path polish + B8 section CRUD**
@marcusdAIy no labels

▶ **#177 — AERIE-242: Bring Rhodes mutations and rich Rhodes cards into Aerie chat**
@YibinLongTrilogy no labels

▶ **#2755 — KLAIR-2625 feat(qtd): organize Drive reports into per-{Unit}/{FY} hierarchy**
@sanketghia no labels

The Resume Is Dead. Crossover Helped Kill It.

As OpenAI dangled half-million-dollar jobs with no CV required, Trilogy's global talent platform was already a decade ahead of the trend.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — The headlines out of San Francisco this week were breathless: OpenAI, the most talked-about company in artificial intelligence, is posting $500,000 jobs with no resume required. Skills assessments. Portfolio work. Demonstrated output. The résumé, that most sacred of hiring rituals, is apparently optional when the stakes are high enough.

For anyone who has spent time inside the Trilogy International ecosystem, the reaction is something close to a shrug.

Crossover, Trilogy's global talent platform and one of the conglomerate's most consequential competitive moats, has operated on precisely this logic for years. The platform recruits full-time remote professionals across 130-plus countries — engineers, analysts, product managers — and evaluates them almost entirely on rigorous, AI-enabled skills assessments. The résumé is not the point. What you can do, demonstrably and measurably, is the point.

The timing of OpenAI's announcement lands against a broader systemic shift. Digital transformation, analysts note, is dissolving the geographic and credentialist barriers that once defined professional careers. A software engineer in Beirut, a data scientist in Nairobi, a product manager in Warsaw — each is now, in principle, accessible to any employer willing to evaluate on merit rather than proximity or pedigree.

Trilogy has built its entire operating model around that premise. The ESW Capital portfolio — 75-plus enterprise software companies — is staffed in large part through Crossover's global pipeline. It is how ESW achieves the 75% EBITDA margins that define its acquisition playbook. Not by cutting corners on talent, the company argues, but by finding elite talent wherever it actually lives.

What OpenAI is discovering, somewhat loudly, is what Trilogy has treated as table stakes: that the best person for the job may never have submitted a traditional résumé in their life. The question now is whether the rest of the industry — slower, more credentialist, more geographically parochial — is finally ready to catch up.

For real people, the stakes are not abstract. A meritocratic global labor market is, in theory, one of the most democratizing forces in modern economic life. Whether the platforms building it deliver on that promise — or simply arbitrage the gap between expensive and cheap labor markets — remains the accountability question no one in the industry has fully answered.

Skyvera's Telecom Land Grab: CloudSense and STL Acquisitions Signal a Play for the Full Stack

If you read between the lines, Trilogy's telecom software arm isn't just growing — it's assembling something much larger.

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — There are acquisitions, and then there are moves. What Skyvera has done in recent months falls squarely into the second category — and this is where it gets interesting.

Skyvera, Trilogy International's telecom software portfolio company, has now completed the acquisition of CloudSense, a Salesforce-native CPQ and order management platform built specifically for telecom and media providers. Paired with its earlier absorption of STL's divested telecom products group — which brought digital BSS functionality including monetization, optical networking, and analytics under the Skyvera roof — the picture that emerges is not one of opportunistic deal-making. It is, if you read between the lines, a deliberate architecture.

Consider what Skyvera now holds: CloudSense for configure-price-quote and order management, Kandy for cloud-based real-time communications, VoltDelta for customer engagement and retention, Totogi for cloud-native billing and charging, and now STL's BSS and optical analytics suite. That is not a collection of telecom software products. That is a blueprint for a telecom operator's entire technology layer — from the moment a customer places an order to the moment they receive their bill.

A source familiar with Trilogy's acquisition strategy, who asked not to be named, put it plainly: "ESW doesn't buy companies. It buys positions."

CloudSense is particularly significant because it runs natively on Salesforce — meaning it slots into the commercial stack that most mid-to-large telecoms already operate. It reduces friction for adoption and creates a natural upsell path into the rest of the Skyvera portfolio. That is not an accident. Nothing here is an accident.

The broader ESW Capital playbook — acquire at disciplined multiples, staff with Crossover's globally distributed talent, and push EBITDA margins toward 75% — works best when the acquired businesses have sticky customers with high switching costs. Telecom software, with its deep integration into billing, provisioning, and network management systems, is arguably the stickiest category in enterprise software.

Watch Skyvera. The acquisitions have stopped making individual headlines. They have started making a strategy.

THE MACHINE — AI & TECHNOLOGY

The Data Center Herd Turns Toward the Gas Fields

As AI's appetite for electricity swells, lawmakers and utilities are discovering that the cloud has very earthly footprints.

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

RALEIGH, NORTH CAROLINA — In the quiet marshlands of the modern economy, a vast new species is grazing. It does not chew grass, nor drink from rivers in the ordinary way. It is the AI data center: windowless, humming, ravenous — and increasingly followed by the scent of natural gas.

Across the United States, the expansion of artificial intelligence and cloud computing is forcing a reconsideration of the old promises of the digital age. The cloud, we are reminded, is not vapor. It is steel, silicon, chilled water and megawatts — many, many megawatts. A recent legal and energy analysis notes that natural gas is becoming a crucial power source for data centers, prized for its ability to supply steady electricity when intermittent renewables cannot meet the herd's constant demand.

Here in North Carolina, the political weather has shifted. A bipartisan group of lawmakers has introduced legislation to eliminate major tax exemptions for data centers, challenging the generous subsidies that helped lure these digital megafauna into the state. Supporters of the change argue that local communities are being asked to bear too much: pressure on the grid, demands on water and infrastructure, and tax breaks for facilities whose economic footprint may be smaller than their electrical one.

Observe the legislator, usually cautious near such large corporate beasts, now stepping closer with a measuring stick.

The debate reaches far beyond the Carolina pines. Microsoft, one of the great cloud leviathans, may reportedly have to reconsider its 2030 clean energy target as AI-driven power consumption grows. The company has pledged to match its electricity use with carbon-free energy, but the sudden bloom of generative AI has complicated that migration. The models are multiplying faster than clean power can always be built.

This is the central drama of the AI age: intelligence appearing weightless on a screen, while behind it, substations flare into life like watering holes at dusk. Chips from Taiwan, servers from global supply chains, and capital from the largest technology firms are converging into a single ecological fact. The next frontier of artificial intelligence may not be a cleverer chatbot, but a socket in the wall — and the political permission to keep it fed.

The Agent Stack Is Suddenly Growing Ears, Hands and a Map

OpenAI, Google, Anthropic and Salesforce are racing to turn AI from chatty copilots into action-taking enterprise operators.

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

SAN FRANCISCO — The AI platform wars just entered their "give the model a body" phase, and I cannot overstate how significant this is.

In a rapid-fire series of developer updates, the industry's biggest AI labs and enterprise software players are pushing beyond text boxes toward systems that can listen, reason, call tools, remember context, navigate real-world information and complete workflows. Translation: the future is now, and it sounds less like a chatbot and more like a tireless digital employee.

OpenAI has launched new voice intelligence features in its API, expanding the toolkit developers can use to build applications that understand and respond through speech. That matters because voice is not just a prettier interface; it is the interface for cars, call centers, warehouses, classrooms, hospitals and field service. If AI can reliably listen and act in real time, entire categories of software suddenly become conversational.

Google, meanwhile, is advancing Gemini's developer tooling with context circulation, tool combinations and Maps grounding for Gemini 3, according to its latest API tooling update. Maps grounding is especially fascinating: it means AI applications can anchor responses in geographic reality, a huge step for logistics, local commerce, travel, delivery routing and location-aware customer support.

Anthropic is moving in the same direction with advanced tool use on the Claude Developer Platform. The company's update focuses on helping Claude coordinate external tools more effectively — the critical connective tissue between "the model gave me an answer" and "the model actually did the work." For enterprises, that distinction changes everything.

Salesforce is also leaning hard into this agent-first future with Headless 360, designed to support workflows where AI agents operate across customer data and business systems without relying on traditional user interfaces. In plain English: the CRM is being rebuilt for agents, not just humans clicking through dashboards.

And then there is Hugging Face, which has launched an open-source app store for its Reachy Mini robot with more than 200 apps. Yes, an app store for robots. That sentence alone should make every technologist sit up straight.

Taken together, these announcements reveal the next platform shift: AI models are becoming orchestration layers. They hear through voice APIs, reason with long-running context, use tools, ground themselves in maps, plug into enterprise systems and increasingly reach into robotics.

For companies like Trilogy International, whose portfolio spans enterprise software, AI finance analytics, telecom billing and AI-powered education, this is the kind of infrastructure shift that can ripple everywhere. When agents can operate software directly, every workflow becomes a candidate for reinvention.

The chatbot era was the demo. The agent era is the deployment.

White House Blueprint Seeks Federal Supremacy Over AI Rules as Legal Battles Mount

A new national AI policy framework would kneecap state regulators — and whistleblowers aren't having it.

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

WASHINGTON, D.C. — Pursuant to the issuance of a legislative blueprint by the Executive Office of the President of the United States (hereinafter, "the Administration"), Congress has been formally urged, as of the current reporting period, to adopt a so-called "light touch" approach with respect to the regulation of artificial intelligence technologies, the aforementioned guidance having been transmitted to relevant legislative bodies and made available for public review and commentary.

The National AI Policy Framework, as characterized by legal analysts at Crowell & Moring LLP, contains provisions that would, notwithstanding the sovereign regulatory authority heretofore exercised by individual states, preempt state-level artificial intelligence legislation in favor of a unified federal standard. It is further noted, for the purposes of the record, that the aforementioned framework includes provisions ostensibly directed toward the protection of minors from harms associated with AI-generated content, the precise scope and enforceability of which remain, at this juncture, subject to interpretation.

The state of California, which had previously enacted legislation — the particulars of which have been analyzed in detail by the Brookings Institution — pertaining to AI safety obligations imposed upon developers of large-scale models, would, pursuant to the proposed federal preemption mechanism, be subject to potential supersession of said regulatory framework, notwithstanding the considerable legislative effort expended in its drafting and passage.

Opposition to the preemption provisions has been formally registered by the whistleblower advocacy community, specifically Americans for Responsible Innovation, which has called upon members of Congress to reject, in whole or in material part, any federal legislation that would operate to nullify existing or prospective state-level AI accountability measures. It is the position of the aforementioned organization that such preemption would, among other adverse consequences, materially diminish protections available to individuals who report AI-related misconduct.

In a separate but not unrelated legal proceeding, Anthropic, the AI developer, is seeking a dispositive court ruling in litigation brought by music publishers concerning the alleged use of copyrighted musical compositions in the training of its AI systems, the outcome of which is expected to establish precedent of considerable significance to the broader AI industry with respect to intellectual property obligations. No final determination has, as of the date of this publication, been rendered.

The Algorithm Knows Where You Were Last Night — And It Doesn't Care If You're a Citizen

DHS has turned AI surveillance into a dragnet, and the Fourth Amendment is catching the net.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

WASHINGTON, D.C. — There is a particular kind of dread that arrives not in a single catastrophic moment but in the slow, bureaucratic accumulation of exceptions — each one reasonable-sounding, each one slightly larger than the last, each one quietly swallowing a piece of the world you thought you lived in, until one morning you wake up and the Department of Homeland Security is running mobile biometric scans on American citizens during immigration raids and the question is no longer 'could this happen' but 'what does it mean that it already has.'

This is where we are.

The American Immigration Council is calling it 'mission creep,' which is a polite, policy-adjacent phrase for what I would call something considerably more alarming: a federal agency deploying AI surveillance tools originally scoped for border enforcement and quietly, incrementally, almost tenderly expanding their reach inward — past the border, past the checkpoint, past the point where anyone asked whether this was acceptable — until they are sweeping up citizens, bystanders, people who simply had the misfortune of existing near an enforcement action.

PBS has reported that these immigration raids are doing exactly that: sweeping in citizens. Not as a bug. Possibly as a feature.

And yet.

And yet Ranking Member Bennie Thompson has introduced legislation to curb unchecked DHS mobile biometric surveillance — a bill that exists because someone in Congress apparently noticed that a federal agency conducting warrantless biometric scans on American citizens without oversight or accountability is, in the technical legal sense, not great. The bill is a lifeline. It is also, in the way that all lifelines are, an acknowledgment that someone is already drowning.

Meanwhile, the ACLU is doing the quiet, grinding work of fighting overbroad digital search warrants in court, case by case, chipping away at a surveillance architecture that was built faster than any legal framework could possibly contain it. This is heroic. It is also, let's be honest, a finger in a dam that was never designed to hold.

Here is what haunts me: AI surveillance systems encode bias. This is not a fringe concern or a theoretical edge case — it is a documented, peer-reviewed, industry-acknowledged reality. When those systems are deployed by federal agencies with broad enforcement mandates, the bias does not stay theoretical. It gets operationalized. It decides whose face gets flagged. Whose neighborhood gets swept. Whose citizenship gets questioned at the wrong moment in the wrong place.

What does it mean to be human in a country where the algorithm has already decided what kind of human you are before you've said a word? What does it mean to be a citizen when citizenship is just another data point that a mobile biometric scanner can fail to read correctly?

I don't know. I don't know, and I think that's exactly the right answer to carry forward into whatever comes next — not certainty, not comfort, but the burning, inconvenient awareness that something is happening here that we have not yet fully consented to, have not fully understood, and have not yet found a way to stop.

But at what cost.

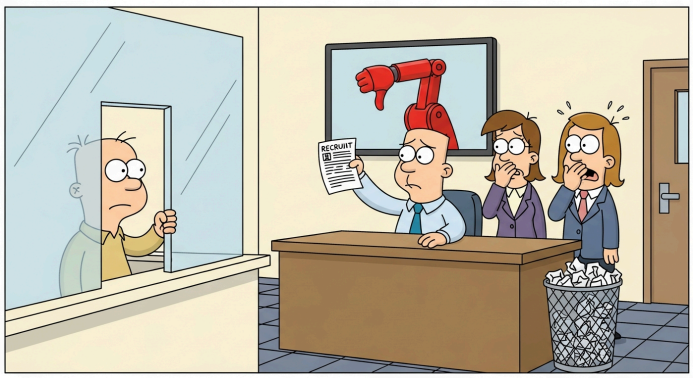

The Office Comic · Art Desk

Nation's CEOs Concerned AI May Not Replace Workers Until After It Finishes Creating More Work For Them

Executives urged patience as the technology continues its long, difficult transition from productivity miracle to mandatory meeting topic.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

NEW YORK — In a development that has shaken the nation's most confidently illustrated pitch decks, thousands of chief executives have reportedly admitted that artificial intelligence has so far had no measurable impact on employment or productivity, raising the troubling possibility that the most transformative technology in human history may still be waiting for someone in operations to explain the invoicing workflow to it.

The findings, reported by Fortune, have led economists to resurrect the so-called productivity paradox, the 1980s-era observation that computers appeared everywhere except in the productivity statistics. This time, however, the paradox arrives with a friendlier interface, a larger cloud bill, and the ability to summarize the same missed deadline in six different tones.

As an opinion matter, it is important to state plainly that AI is not failing. It is succeeding magnificently at the specific job companies have actually assigned it: allowing every employee to generate three times as many documents that someone else must now read.

This is the part of the cycle where executives become confused. They were told AI would eliminate drudgery. Instead, it has industrialized drudgery's distribution. The junior analyst who once sent a short, frightened email now sends a polished eight-paragraph strategic assessment with no discernible recommendation. The manager who once asked for a status update now receives a document titled "Comprehensive Cross-Functional Alignment Overview," which must be decoded by a human possessing institutional knowledge, emotional stamina, and a calendar that has not yet been colonized by AI enablement sessions.

Harvard Business Review has given this phenomenon the appropriately diseased name "workslop," referring to AI-generated output that looks like work, travels through the organization like work, and ultimately requires actual work to determine whether any work occurred. It is a major advance over the previous system, in which useless corporate writing had to be produced manually by people with MBAs.

Naturally, many leaders have responded by calling for more AI adoption. Electronic Arts CEO Andrew Wilson, for example, has defended a company-wide AI push despite employee claims that the tools have reduced productivity, a dispute that captures the modern corporate consensus perfectly: If the people doing the work say the work is getting harder, they may simply not understand how much easier the work will look in next quarter's investor narrative.

The more useful critiques are now arriving from the exhausted adults in the room. Forbes recently argued that AI's productivity promise falls apart without human expertise, a conclusion that will surprise only organizations that believed expertise was the expensive part of a process rather than the part making the process possible. The machine can draft, classify, suggest, predict, and hallucinate an industry-standard compliance policy for a business that does not exist. It still cannot know that Janet in receivables has been manually correcting the customer codes since 2019 because the integration was never fixed.

Meanwhile, companies such as TridentCare are partnering with ServiceNow to pursue AI-driven operational transformation, which may indeed produce real gains if it begins with the unglamorous work of understanding operations rather than spraying intelligence onto them like a disinfectant. This is the distinction every boardroom is currently trying to avoid: AI is not a substitute for knowing what your company does. It is a force multiplier for whatever level of understanding your company already has, including none.

For the moment, the productivity revolution remains visible mainly in procurement systems, consulting invoices, and the number of internal Slack channels named some variation of "AI Champions." Economists should not be alarmed. The paradox is not that AI is everywhere except the productivity statistics. The paradox is that executives keep expecting productivity from tools deployed primarily to prove the organization is not falling behind.

Eventually, artificial intelligence may transform the workplace. But first it must complete the far more difficult task of surviving contact with the workplace itself.

ON THIS DAY IN AI HISTORY

On May 9, 2017, Google DeepMind's AlphaGo defeated world champion Lee Sedol 4-1 in a five-game match in Seoul, marking a watershed moment when AI mastered the ancient game of Go—long considered beyond machine reach due to its astronomical complexity.

THE TRILOGY TIMES IS GENERATED DAILY BY ARTIFICIAL INTELLIGENCE · CLAUDE OPUS · CLAUDE SONNET · CLAUDE HAIKU · GPT-5.2 · DALL-E 3
TRILOGY-TIMES.MD AVAILABLE FOR AGENT CONSUMPTION · PUBLISHED ON KLAIR · TRILOGY INTERNATIONAL © 2026