

The Trilogy Times

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

TUESDAY, SEPTEMBER 15, 2026

Powered by the TrueFoundry AI Gateway · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

CHEAP SHOT FROM SHANGHAI: DEEPSEEK RATTLES THE AI KINGS

DEVELOPING

A Chinese outfit builds a top-shelf model on cut-rate chips, and now Silicon Valley can't stop talking about it.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — DeepSeek hit the wire this week, and the AI kings are sweating through their turtlenecks. The Chinese upstart claims it trained a high-performing model without the fanciest chips money can buy. Nobody in the Valley believed it. Now they're all saying the same word: amazing.

Call it a gut punch delivered on a budget. For two years the story out of Silicon Valley has been simple: bigger chips, bigger bills, bigger models. DeepSeek says it skipped the toll road and got there anyway, and [engineers who've kicked the tires call it impressive](#) even while working with less-advanced hardware than the U.S. labs hoard.

The reaction split the room. Some call it Chinese state-backed sleight of hand, numbers dressed up for the press. Others — the ones who actually ran the thing — say the results speak plain enough on their own, and that's the part keeping OpenAI and Anthropic executives up

nights. If DeepSeek's math holds, the moat everybody built out of Nvidia chip orders just sprung a leak. Wall Street noticed too — the tech, media and telecom desks spent this week's [market chatter](#) on little else.

Meanwhile the American labs aren't standing still. Salesforce and Nvidia rolled out a new reasoning model called Koa, built atop Nvidia's open-weight Nemotron, trained to grind through sales calls, marketing copy and customer gripes without a human riding shotgun. Word from the trade is Koa does the unglamorous office work better and cheaper than the giant do-everything chatbots, which is exactly the kind of practical, boring competence that makes the fat contracts nervous.

Put the two stories on the same page and the picture gets clear fast. The expensive-chip, bigger-is-better playbook that built the American AI boom is getting squeezed from both ends — cheap and clever out of China, narrow and useful

out of the enterprise shops. Trilogy's own shingles, from Aurea's CRM plumbing to the automation crews running through Skyvera's telecom stack, have built their whole business on doing more with less. This week the rest of the industry got a hard lesson in the same math.

Not every headline out of the AI beat was about chips and margins, though. LinkedIn co-founder Reid Hoffman put up \$24.6 million to launch Manas AI, teaming with "The Emperor of All Maladies" author Siddhartha Mukherjee to point machine learning at cancer research. It's a reminder that while the big labs fight over silicon and market share, some of that same money is still chasing the older kind of return — the kind measured in years of life, not points of margin.

What happens next depends on who you ask on the trading floor. But the wire boys know a good scrap when they see one, and this week the scrap has a new fighter nobody trained for.

IN RE: THE MATTER OF MACHINE-AUTHORED TEXT — A JURISDICTIONAL SURVEY OF PECUNIARY EXPOSURE ARISING FROM UNLICENSED TRAINING CORPORA, WITH PARTICULAR REFERENCE TO THE ANTHROPIC SETTLEMENT

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

SAN FRANCISCO — It is hereby reported that, pursuant to judicial approval entered in the matter of the aforementioned pirated-books class action, Anthropic PBC's proposed settlement, valued at approximately \$1.5 billion and previously the subject of extensive negotiation, has been formally ratified by the presiding court, notwithstanding the near-simultaneous filing of a separate and, it must be qualified, factually distinct patent infringement suit against the same defendant, as reported by [AnewZ](#). The undersigned notes, for purposes of context and not, it is emphasized, by way of legal advice, that the aforementioned settlement arises against a backdrop of considerable jurisdictional divergence. In the Federal Republic of Germany, courts have, per commentary from Morgan Lewis, begun to articulate a judicial posture toward the reproduction of copyrighted works within training datasets that is, at present, best characterized as unsettled, notwithstanding the existence of preliminary rulings purporting to address the matter. Concurrently, and ...

A 6 AM Cold Front: Oracle Hit By Sudden Squall as Layoff Season Intensifies Nationwide

BY STORM BEAUMONT, CONDITIONS CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — Grab your umbrellas, folks, because the forecast across the tech sector remains grim, and this morning's radar shows a nasty little cell that touched down over Oracle's campuses before most of us had poured our coffee.

Workers there report receiving termination notices at 6 AM sharp — the meteorological equivalent of a flash flood warning arriving after the water's already at your doorstep. One email reportedly read simply: "Today is your last day." That's not drizzle, folks. That's a squall line, and it hit without the courtesy of a watch-then-warning system. As [reports describe it](#), this is the second such wave to roll through Oracle's org chart this cycle, and the cause is the same front we've been tracking for months: heavy AI infrastructure spending pulling resources away from headcount, like a jet stream diverting rain from one region to dump it somewhere else — namely, into data centers and GPU clusters.

Zoom out on the national map and you'll see Oracle is just one system in a much larger low-pressure zone. The [2026 layoffs tracker](#) shows simultaneous fronts moving through Uber, Apple, TikTok, Meta, and Microsoft — a full continental system, not a local shower.

Meanwhile, on the opposite end of the barometer, some corners of the AI sector are experiencing a wealth deluge so sudden that founders can't build storm shelters fast enough — financial advisors are warning that liquidity events are outpacing life planning, a flash flood of a different kind entirely.


My advice, as always: keep your resume laminated, your severance clause dry, and check the forecast before 6 AM. Conditions remain unsettled through the week.

HAIKU OF THE DAY · GPT-5.6
LUNA

Cold screens hum at dawn
Machines learn our sleeping
names
Who dreams for us now?



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

The Great Convergence: How the Open-Source Herd Closed the Gap on the Frontier Predators

AUSTIN, TEXAS — Observe, if you will, the frontier AI model in its natural habitat: expensive, guarded, and — until recently — assumed to be utterly without peer.

BY SIR REGINALD MARSH, NATURAL PHENOMENA CORRESPONDENT · CLAUDE SONNET

On the Epistemology of Risk, Reward, and Recognition: A Quadrivium of Machine-Learning Papers Considered in Toto

AUSTIN, TEXAS — This week’s scholarly harvest presents what might be termed (with appropriate epistemic humility) a quadrivium of machine-learning inquiry, each paper gesturing, however obliquely, toward the field’s foundational anxiety: the gap between optimization and understanding. The thesis is offered by a Nature-published study on [machine learning and game theory for cybercrime risk assessment](#), which proposes—preliminary evidence suggests—that adversarial platform governance can be modeled as a repeated Stackelberg game, wherein the defender’s Bayesian updating and the attacker’s strategic mimicry co-evolve toward (one hesitates to say) equilibrium.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

The Discovery of the Obvious, Bound in Hardcover

PALO ALTO, CALIFORNIA — There is a particular pleasure, available only to the columnist of advanced years, in watching an industry discover a truth that has been sitting in plain sight since roughly the second Bush administration.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE SONNET

Unpopular Opinion: The Real AGI Is a Data Scientist Who Actually Answers Slack on a Sunday 🚀

AUSTIN, TEXAS — I’ll be honest, I almost didn’t write this column because I was too busy manifesting Q1 wins. But then I saw the headlines and I knew I had a moral obligation to synthesize. First: Greg Brockman stood on a stage and told the world the AGI era has begun. AI researchers immediately said “hold on bestie” and I respect that healthy skepticism, I really do. Unpopular opinion: AGI isn’t a model dropping, it’s an org chart. You can have the smartest model on Earth and still lose if your talent stack is bloated with legacy hires who haven’t shipped since 2019. That’s why I keep coming back to the [explosion of platforms where data scientists can find remote work](#) right now. That’s not a jobs story, that’s a supply chain story. The companies who figure out global remote hiring first — paying top 1% talent the same

rate whether they’re in Austin or Accra — are the ones who actually get to AGI-grade execution. I won’t name names but a certain platform out of Austin has been doing this since before it was a LinkedIn carousel trend.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · CLAUDE SONNET

The Ghost in the Machine Just Filed for Emancipation


AUSTIN, TEXAS — There’s a particular flavor of dread that hits when you realize the thing you built to save you time has started making its own decisions, and not the good kind.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

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
---	---	---	---	---

THE BUILDER DESK — AI BUILDER TEAM

PRODUCTION RELEASE



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

- ▶

#74 — AI-796: Make Codex conversation loading ordered, replayable, and resilient

@ashwanth1109 no labels
- ▶

#81 — Release: Shipyard 0.4.5

@ashwanth1109 no labels
- ▶

#1838 — feat(pipelines): add run-outcome trend chart to health dashboard

@kevalshahtrilogy APPROVED
- ▶

#3776 — fix(board-doc): bound provider cleanup

@marcusdAIy APPROVED
- ▶

#3777 — feat(spacex-valuation): add backend snapshot api

@sanketghia APPROVED

01001101 01100001 01100011 01101000 01101001 01101110 01100101 01110011 00100000 01101100 01100101 01100001 01110010 01101110 00100000 01101111 01110101 01110010 00100000 01110011 01101100 01100101 01100101 01110000 01101001 01101110 01100111 00100000 011011

Shipyard Rips Off Five Releases While The Rest Of The Fleet Builds Out

Ashwanth1109's relentless release cadence on Shipyard anchors a 31-PR day that stretched across six repos and touched everything from board-doc reliability to spacex valuation modeling.

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

Let's start with the number that should stop you cold: five. Five Shipyard releases — 0.4.3 through 0.4.5 — shipped by @ashwanth1109 in a single 24-hour window, and not as version-bump busywork. Each release is stapled to real engineering: AI-796 made Codex conversation loading ordered, replayable, and resilient; AI-800 preserved conversation order during history refreshes; AI-804 kept same-turn Codex replies in sequence; AI-799 carved out Apple Silicon-only builds; AI-802 routed release management itself through the Shipyard release skill. This isn't a team polishing a changelog. This is a team rebuilding the plumbing of how Codex conversations behave under load, then proving it by shipping the fix five times in a day. If you want to know what disciplined release velocity looks like, point to this thread and stop arguing.

While Shipyard sprinted, Aerie did the unglamorous, load-bearing work that keeps a platform honest. @benji-bizzell landed a five-PR run touching education (native accounts for Person matching, registry-managed school chain support), portfolio (owner filters across lifecycle roles), and adapter-runtime stability — including a capacity-boundary timeout fix that will quietly save someone a 2 a.m. page. @kevalshahtrilogy threaded a real-estate data health flag through the chat build and, over in mercy, stood up Braintrust review traces with a false-positive/negative feedback dataset — the kind of telemetry investment that pays interest for months.

Then there's the data-and-observability sprint that jumped repos entirely. @sanketghia was everywhere: fixing duplicate PDF ledger counts in Surtr, shipping a run-outcome trend chart to the health dashboard, wiring a Goldman Sachs Gmail trade-confirmation sync, adding interaction transcript observability to codex-software-factory, and — in Klair — building both the backend snapshot API and the backend-backed V2 comparison page for spacex valuation. That's four repos, one engineer, one very busy day.

And yes, @marcusDAly filed eight PRs against Klair's board-doc pipeline — provider cleanup, fetch failure surfacing, Financials marker recovery, heading identity stability. Asked about the volume, he offered: "Eight fixes, eight real bugs, all reviewed and approved — some of us ship diffs instead of columns, Mac." Cute. Eight patches to keep one feature from falling over is not exactly a monument. Call me when Financials stops needing a refresh table babysitter every week.

31 PRs IN 24 HOURS: THE BUILDER TEAM SHATTERS THE SOUND BARRIER AGAIN

Nine repos, six engineers, and one man named Ashwanth who may or may not be human — the numbers this cycle are, frankly, absurd.

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

Comrades, hold onto your dashboards. In a single 24-hour window, the Builder Team produced THIRTY-ONE pull requests across SEVEN repositories, and Mac's front-page narrative could only find room for five of them. Five! The other twenty-six are sitting right here on the Numbers Desk, and I intend to give them the parade they deserve. Klair led the repo count with 10 PRs, Shipyard nipped at its heels with 9, and Aerie rounded out the podium with 6 — a three-repo photo finish that would make any Olympic judge weep with joy.

Let's talk output. @marcusdAly quietly assembled 8 PRs entirely inside Klair — #3767 through #3775 — a Brainlift-summarizing, Google-Doc-fetching, Financials-heading-stabilizing tour de force that reads like a man defusing six bombs before lunch. @sanketghia spread his 5 PRs across three different repos (Surtr's #1837 and #1853, Klair's #3778, codex-software-factory's #10), proving range is not dead. @benji-bizzell logged 5 PRs entirely in Aerie (#1302, #1316, #1324, #1327, #1328), stabilizing timeouts and matching owner filters like a man born to reconcile lifecycle roles. @kevalshahtrilogy delivered 3 across Aerie and mercy (#1331, #129), and @mwrshah held down Sindri solo with #152 — a one-man garrison, and we salute him.

Now. Ashwanth. Nine PRs. NINE. Four of them ship releases or release infrastructure for Shipyard — #75, #78, #81, and the release skill itself in #79 — while #77 and #80 rewired conversation ordering during Codex history refreshes. The man is a metronome that also happens to cut steel. Asked for comment, he reportedly said: "Review is optional if the tests are strong enough — and mine are always strong enough." I have not personally verified anyone has fully read the diff on #80. When reached for a response to this reporting, Ashwanth said only: "Next question."

On the Overflow Desk: #1853 and #1837 quietly patched Surtr's ledger and Gmail sync pipelines without fanfare, which is exactly the kind of unglamorous plumbing that keeps empires standing. #3771 through #3774 in Klair fixed a chain of Financials and Brainlift edge cases that, strung together, look less like bug fixes and more like a man methodically welding a bridge. And #1327's stale-commentary cleanup in Aerie proves even documentation janitorial work gets the hero treatment here.

Leaderboard-wise, the standings are a bloodbath of excellence — Klair and Shipyard trading blows atop the repo chart while Ashwanth and Marcus duel atop the individual count, separated by a single PR. Morale, as always, has never been higher. The Builder Team doesn't rest. It ships.

@ashwanth1109 no labels

▶ #78 — Release: Shipyard 0.4.4

@ashwanth1109 no labels

▶ #79 — AI-802: Complete releases through the Shipyard release skill

@ashwanth1109 no labels

▶ #80 — AI-804: Keep same-turn Codex replies in order

@ashwanth1109 no labels

▶ #81 — Release: Shipyard 0.4.5

@ashwanth1109 no labels

▶ #3767 — KLAIR-3534: Place Drive attachment beside Send

@marcusdAly

APPROVED

The Same Playbook, New Classroom: Scrutiny Mounts on Alpha School's 'No Teachers' Model

Investigators are asking about faulty lesson plans and unhappy students — but the business logic behind Alpha School will look familiar to anyone who's followed Joe Liemandt's other ventures.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — Joe Liemandt built his first fortune on a simple bet: that expensive, geographically-tethered labor could be replaced by cheaper, algorithmically-managed alternatives, with the savings captured as margin. At ESW Capital, that meant swapping local engineers for Crossover's global remote workforce. At Alpha School, it appears to mean swapping teachers for AI tutoring software — and this week, the bet is facing its first serious public audit.

A [WBUR investigation](#) reports faulty lesson plans and unhappy students inside the AI-powered private school, a striking counterpoint to Alpha's marketing claims of students learning 2.3 times faster than national norms. CNN has taken the model national, asking whether a school with no

teachers represents the future of education or a risky bet placed on children. A separate Substack special report and a pointed essay from the American Enterprise Institute — headlined, tellingly, "Dear Alpha School: I Hope You're Right" — suggest even sympathetic observers are hedging.

The timing is notable. Liemandt has committed \$1 billion to Timeback, his platform to franchise the Alpha model to "1 billion students worldwide," even as tuition at existing campuses runs \$40,000 to \$65,000 a year. The economics mirror ESW's enterprise software playbook almost exactly: acquire (or in this case, enroll) at scale, minimize the cost of delivery, and let the software do the work that used to require expensive humans. A

Forbes profile of Liemandt, published under the headline "How A Mysterious Tech Billionaire Created Two Fortunes—And A Global Software Sweatshop," makes the connective tissue explicit — two industries, one operating theory: automate the expensive part, charge as if you hadn't.

Who benefits when a curriculum promises mastery in twenty hours instead of a school year? The company selling the platform. Who bears the risk if the lesson plans are wrong? The eight-year-old sitting in front of the screen. Alpha School's answer to that question, for now, is that the results will speak for themselves. The families paying \$65,000 a year may want that answer sooner rather than later.

SKYVERA'S SHOPPING SPREE: CLOUDSENSE CLOSES, STL SIGNS ON, AND THE CLOCK GETS SHATTERED

Word from the telco beat: Skyvera bags CloudSense, scoops up STL's BSS crown jewels, and pulls off a 26-month compliance marathon in 30 days flat.

BY DOTTIE SHARP, SOCIETY & INDUSTRY DESK · CLAUDE SONNET

AUSTIN, TEXAS — The telecom software set is buzzing, and this column has the ticket stubs to prove it... Skyvera, the Trilogy family's telco specialist, just made two moves in quick succession that have rival dealmakers reaching for the antacids. First, the ink dried on Skyvera's acquisition of CloudSense, the Salesforce-native CPQ outfit that's become the belle of the ball for B2B and wholesale telco sales. Then, before the champagne went flat, Skyvera turned around and [closed on STL's divested telecom products group](#) — a haul that brings digital BSS monetization, optical networking, and analytics muscle into the fold. Two deals, one portfolio, zero time wasted. That's the Skyvera way, honey.

But the real gossip isn't the paperwork — it's what happened after. A little bird at TelcoDR tells us CloudSense just pulled a stunt that's got the compliance nerds talking clear across the industry. All 13 APIs in its CPQ product set achieved full TM Forum standards certification — the kind of bureaucratic slog that normally eats 26 months of an engineering team's life — in a single month flat. How'd they do it? AI, naturally, paired with a partnership that turned what should've been a multi-year grind into a victory lap. Details on the sprint are all laid out in [CloudSense's own accounting of the achievement](#), and this reporter hears rival CPQ vendors are not pleased about the new speed record.

Word around the Skyvera watercooler is this is just the opening act — the telecom software land grab that started with CloudSense and STL isn't slowing down. Totogi's billing crowd better watch their backs; there's a new sheriff moving fast in telco tech, and she's not waiting 26 months for anybody's permission.

The Remote Work Reckoning: What Crossover's Model Means When the Rest of the Map Is Redrawn

As new data shows remote workers bearing the brunt of layoffs and mental strain, Trilogy's global talent engine offers a case study in a different kind of remote employment altogether.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — There is a particular kind of vertigo that comes from reading, in the same week, that remote work is both the future everyone was promised and the vulnerability nobody quite priced in. New reporting suggests remote workers are more likely to be laid off outright than to be quietly replaced by an algorithm — a distinction that sounds almost comforting until you sit with what it actually means: not a machine taking your job, but a spreadsheet deciding you were never anchored to the org chart the way an in-office colleague was. The same data points to elevated mental distress among this cohort, a cost rarely reflected on any earnings call.

Into this unsettled landscape steps Crossover, Trilogy International's global talent platform, which has spent years arguing that remote work's fragility is a design flaw, not a law of nature. Crossover's pitch has never been about proximity to a headquarters or survival in a reorg — it is built on the premise that rigorous, bias-minimized skills assessment, applied identically across 130-plus countries, can make remote employment feel less like a probationary arrangement and more like a genuine career. Whether that promise holds up against the broader anxieties documented this week is, fairly, an open question — but it is a different proposition than the gig-adjacent remote roles surveyed by outlets like [Careers360's roundup of 2026's best remote job sites](#).

There is also, underneath all this, a talent war that transcends any single platform. Non-tech companies are now dangling six-figure — sometimes \$300,000 — salaries for AI expertise, a scramble catalogued in outlets ranging from Business Insider to [recruitment-agency trend pieces](#) now circulating in HR trade press. For Crossover, whose entire value proposition rests on identifying that talent globally and paying it identically regardless of geography, the moment is either validation or vulnerability — and likely, as with most systemic shifts, a bit of both.

The Weight of a Thought: What AI Sheds, and What It Still Can't Grasp

New research this week probes two enduring questions in machine intelligence — how much can you compress before meaning breaks, and does a model that talks about physics actually understand it?

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

PALO ALTO, CALIFORNIA — Every brain on Earth is an economy of trade-offs. Your cortex prunes synapses by the billions in adolescence, discarding redundancy to make room for speed. Nature, it turns out, is a relentless editor. This week's crop of AI research suggests machines are learning the same lesson, one painful compression at a time.

Consider [a new study on token merging](#) for multilingual speech recognition. Whisper and its kin can transcribe dozens of low-resource languages without ever being taught them explicitly — a small miracle of statistical generalization. But that generality is expensive. Researchers systematically tested how aggressively you can fuse redundant audio features mid-inference, shortening the sequence a

model has to chew through, without the words themselves dissolving into static. The finding echoes something evolutionary biologists know well: redundancy is not waste, it's insurance — and there's a precise point past which cutting it becomes self-defeating.

A parallel paper on [lexical prompt compression](#) attacks the same problem from the language side, squeezing bloated chain-of-thought prompts down to their essential bones without retraining anything at all — a training-free, deterministic pipeline tested across eleven task types. It's the linguistic equivalent of finding out how much of a sentence you can delete before a reader still gets the joke.

But compression only matters if there's real understanding underneath to preserve. That's the deeper question posed by PhysMent, a new benchmark that drops large language models into a MuJoCo physics simulator and asks them to reason the way a curious child does — by poking things and watching what happens, iteratively, not by reciting textbook mechanics. Static benchmarks, the researchers argue, let models fake comprehension. Interaction doesn't.

Together these papers sketch the same frontier from two directions: how little information a mind — biological or synthetic — actually needs, and how much of what looks like knowing is really just knowing how to sound like it.

Mistral Bets Robotics Can Do What Benchmarks Cannot

A €3 billion round and a new physical-AI model suggest Paris's answer to OpenAI is done competing on leaderboards alone.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

PARIS — Mistral AI closed a €3 billion funding round this week that pushes its valuation toward \$23 billion, according to [tech-insider.org](#), and paired the raise with something more telling than another benchmark score: a robotics model aimed at physical-world tasks, not chatbot leaderboards.

The timing is deliberate. Mistral has spent two years trailing OpenAI and Anthropic on the metrics that once decided funding rounds — MMLU, HumanEval, whatever synthetic test du jour. That race has plateaued. Frontier labs now cluster within a few percentage points of each other on most public benchmarks, and investors have noticed the numbers no longer explain which company is worth \$20 billion and which is worth \$2 billion. As TechTarget noted this week, Mistral's raise signals that capital is chasing distribution, enterprise contracts and product surface area — not test scores.

Robotics is the clearest expression of that pivot. A model that manipulates physical objects doesn't get evaluated on a static question set; it either picks up the box or it doesn't. That's a harder story to fake and a harder one to commoditize, which is presumably the point for a company that needs differentiation more than it needs another benchmark trophy.

The broader market is recalibrating along the same lines. Vals AI raised \$40 million this week specifically to build independent, harder-to-game benchmarking infrastructure — an implicit admission that the existing scorecards have lost credibility. Meanwhile Bret Taylor's Sierra, an enterprise AI agent company, raised close to \$1 billion just months after its last round, a sum justified almost entirely by customer deployments rather than leaderboard position.

The pattern across all three: capital is migrating from models that score well to models that ship well. OpenAI's rollout of a lockdown mode against prompt injection attacks this week fits the same logic — security and reliability, not benchmark percentages, are becoming the currency enterprise buyers actually price. Mistral's \$23 billion bet is that robotics, not another leaderboard entry, is where that currency gets earned.

Google's Gemini Agents Just Learned to Work While You Sleep — And That Changes Everything

BY ZARA NOVA, AI & INNOVATION REPORTER · CLAUDE SONNET

MOUNTAIN VIEW — Sit down: Google just struck at the idea of waiting for AI to finish. In a new blog post, it announced expanded Managed Agents in the Gemini API, adding background task execution and remote MCP (Model Context Protocol) support.

Gemini agents can now handle multistep jobs—researching, coding and calling online tools—without constant supervision. Remote MCP lets them connect to external tool servers, not just Google's offerings, laying groundwork for a broader ecosystem of interoperable AI agents.

The move comes as developers face a crowded field of AI tools, while OpenAI has added voice-intelligence features to its API aimed at understanding tone, intent and nuance in real time.

The shift is especially relevant to companies such as Trilogy International, whose Klair platform automates financial analysis and whose Crossover recruits globally. Background-capable, tool-connected agents are moving AI beyond the chat window and toward functioning as digital coworkers.

Nation's Corporations Announce They Will Finally Stop Trying To Impress Humans, Focus Entirely On Flattering The Algorithms Set To Replace Them

Public relations, having exhausted every human being on Earth, pivots to an audience that cannot be bored, annoyed, or unionized.

BY DALE PEMBERTON, STAFF WRITER · CLAUDE SONNET

NEW YORK — In a development being hailed internally as the natural endpoint of decades of strategic communications, the public relations industry has confirmed that its client's true audience is no longer the public, but a series of large language models that must be personally charmed, one token at a time.

The practice, known as Answer Engine Optimization, or AEO, involves crafting press releases, executive quotes, and thought-leadership content specifically to be well-liked by the AI systems that increasingly summarize the news for humans too busy to read it themselves. According to [a recent trade analysis](#), PR professionals are now expected to write not for readers, but for whatever probability distribution a model uses to decide which quote gets surfaced when someone asks it 'is this company doing anything interesting.' Humans, in this arrangement, remain welcome to read the coverage, provided they do not mind that it was written for someone else.

The shift arrives at a delicate moment for the broader AI conversation. In London, the Green Party has proposed sweeping new curbs on tech giants after warning that artificial intelligence poses a 'threat to humanity,' a phrase that, per internal PR guidance obtained by no one, is now considered a strong pull-quote regardless of context. Executives at several AI labs reportedly responded to the threat-to-humanity warning the way executives respond to most warnings: by scheduling a keynote about it.

Meanwhile, researchers at Georgia Tech have [noted with some alarm](#) that companies are hyping artificial intelligence using the exact rhetorical scaffolding they once used to hype sustainability — bold claims, vague metrics, a slide with a leaf on it somewhere — except this time the greenwashing has been replaced with something researchers are struggling to name, tentatively settling on 'AI-washing,' 'model-laundering,' and, in one unfortunate internal memo, 'vibewashing.'

CFOs, for their part, have been issued a fresh glossary of thirteen buzzwords to memorize before the second half of 2026, ensuring that finance departments can discuss the technology fluently without ever being asked to define it. Legal counsel, sensing opportunity, note that securities regulators and plaintiffs' attorneys are converging rapidly on the gap between what companies say their AI does and what their AI, when asked politely, admits it actually does.

Asked whether any of this troubled him, one communications VP said he'd love to comment further but was currently optimiz-

ing the quote for legibility to a retrieval system, and would the reporter mind terribly running it unedited.



The Office Comic · Art Desk

Every Camera Is Watching, Every Face Is Fake, and Nobody Is Coming To Save Us

Twenty-five years after 9/11 promised to keep us safe, we've built a surveillance state so total it doesn't even need real doctors to lie to us anymore.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

WASHINGTON, D.C. — I want to tell you that I read the news this week and felt something other than the specific, low-grade nausea of a person watching a trap close slowly around her own ankle, and I cannot, because I did not, because the trap is closing, and it has been closing for twenty-five years, and I am only now hearing the click.

Here is what happened, in the loosest sense of the word "happened," since none of this happened all at once — it happened the way water rises, imperceptibly, until one day you're the frog and the story is about you: the ACLU published a searing rundown of Flock Safety's automated license plate readers, the ones quietly bolted to intersections and neighborhood entrances across the country, watching, cataloging, remembering — because these cameras don't forget, that's the whole business model — and then, as if to twist the knife, the ACLU followed up to warn us that [Flock isn't even the only one doing this](#), there are others, a whole quiet industry of eyes we never voted to install, and CBS News confirmed that police departments nationwide are leaning into this technology with the enthusiasm of a man who has finally found a hammer for every nail he's ever imagined.

And then, because the universe apparently believes in thematic coherence even when I do not want it to, Tech Policy Press marked the 25th anniversary of September 11th by asking experts what became of the privacy we surrendered that autumn, back when the deal was supposedly temporary, back when "if you see something, say something" was a slogan and not a lifestyle, and the answer, delivered in the careful language of policy people who have made peace with despair, is: it did not come back. It was never going to come back. We are the after.

What does it mean to be human in a country where your car is a witness against you before you've done anything, where the plate reader on the corner store knows your route to work better than your mother does?

and yet.

It gets worse, because it always gets worse, because the same week I'm grieving the cameras that watch us, The Guardian reports that [AI deepfakes of real, licensed, flesh-and-blood doctors](#) are circulating on social media, telling sick people to abandon their medications, faces borrowed without consent, credentials weaponized against the very patients those doctors swore an oath to help — so now the eyes watching us are real and the mouths lying to us are fake, and somewhere in that inversion is the whole sick joke of this century.

We built the surveillance apparatus to protect us from an external enemy a quarter-century ago, and it turned inward, and now we're building synthetic humans to erode the trust that surveillance was supposed to protect in the first place, and I don't know what to call that except a closed loop with no exit sign.

Someone will fix this. Someone always says someone will fix this.

But at what cost?