

# The Trilogy Times

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

MONDAY, MAY 25, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

## Korn Ferry Swallows Trilogy International — and Inherits a Staffing Empire It Now Must Justify

*The executive search giant absorbs a global recruitment network, raising questions about consolidation in the talent industry.*

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

**L**OS ANGELES — The deal was quiet by the standards of a firm that places C-suites at Fortune 500 companies. But [Korn Ferry's acquisition of Trilogy International](#) lands with real weight in the staffing industry — a sector where scale is oxygen and the margin game is merciless.

Korn Ferry, the Los Angeles-based executive search and organizational consulting firm, has taken ownership of Trilogy International, a staffing and workforce solutions network with reach across multiple geographies. The financial terms were not disclosed. The acquisition arrives as Korn Ferry has also moved to shore up its own internal architecture, [appointing a new Chief People and Legal](#)

[Officer](#) — a signal that the firm is not merely buying market share but attempting to integrate what it buys.

The staffing industry has spent the better part of three years consolidating. Remote work dissolved geography as a hiring constraint, then reimposed it in subtler ways — compliance, time zones, tax jurisdictions. The firms that survived the post-pandemic correction were the ones that could operate across borders without losing margin to friction. Trilogy International, with its cross-border recruitment infrastructure, fits that thesis.

For Korn Ferry, the acquisition extends a strategic arc away from pure executive search and toward something closer to a full-stack talent services model. It is a bet

that clients want fewer vendors, not more — that the firm which places the CEO can also staff the teams below.

The deal is worth watching for anyone in the talent ecosystem. Korn Ferry now commands a larger footprint in mid-market and professional staffing, territory it has historically left to specialists. Whether it can run that business with the discipline of a firm built on high-touch, high-margin search work is the open question.

The geography of ambition, in this industry, is always expanding. The harder work is holding the ground you've already taken.

# Staff Breaks Ranks in Pentagon Brawl

Google and OpenAI workers back rival Anthropic against the brass — even as Mountain View loses the coding war

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

**S**AN FRANCISCO — Hundreds of Google and OpenAI workers signed papers this week backing rival Anthropic in a Pentagon court fight — even as their bosses chase the same outfit for the AI coding crown.

The amicus brief, joined by staffers from companies that compete with Anthropic by day, asks federal courts to set rules on military AI before the generals write their own. Loyalty stops at the courthouse door.

The brief lands in an odd hour for the industry. Anthropic and OpenAI snipe at each other over safety and talent. Google bleeds engineers and morale.

Mountain View has lost the coding war. Anthropic's Claude eats the developer market by the week. OpenAI's models bite from the other side, while [the Los Angeles Times reports](#) Google's Gemini holds the rear.

Insiders tell the Times the trouble runs deep. Committee fights slow shipping. Engineers leak that the brass has lost the room.

Then Axios twists the knife — by saying Google might yet win. The piece pegs Mountain View as the dark horse if Anthropic and OpenAI bleed each other dry first. Cash, chips, and the search monopoly buy plenty of time.

The coding crown matters plenty. Every developer hooked on Claude or ChatGPT becomes a customer for a decade. Lose the coders and lose the next ten years of enterprise software.

The Pentagon case turns on what Anthropic's machines can — and cannot — do in uniform. [A Wired account](#) describes workers wanting courts to draw

the line before contracts get signed. They want it in writing now.

The brass wants AI for war planning, logistics, and intelligence. Anthropic's models could do most of it. The fight is over guardrails — what the machines will refuse when the uniform asks.


OpenAI staff signed. Google staff signed. Some signers work for outfits pitching the Pentagon on their own deals.

Anthropic's lawyers welcomed the cavalry. The Justice Department declined to wave the white flag. The fight heads to the bench next month.

Trilogy's portfolio watches close. ESW Capital's roster — Aurea, IgniteTech, Skyvera, CloudFix — buys AI plumbing wholesale. Every model becomes a commodity faster than the last.

The whole industry just got a lesson in worker solidarity. Front-line coders care more about rules of the road than the company logo on their checks. The executives just learned the staff has a mind of its own.


For Google, the long game still looks rocky. Solidarity from rival rank-and-file won't help if product drags. The race goes to the swift — and Mountain View ain't sprinting yet.



The Far Side Style · Art Desk

HAIKU OF THE DAY · CLAUDE  
HAIKU

*Empires swallowed whole  
Yet nothing settles—just shifts  
Power finds new hands*



The New Yorker Style · Art Desk

## Antitrust Enforcement Horizon for Tech Sector Remains Legally Uncertain as Regulatory Debate Intensifies

WASHINGTON, D.C.

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

## The Fairness Reckoning: AI's Bias Problem Reaches Critical Mass Across Education, Hiring, and Policing

CAMBRIDGE, MASSACHUSETTS — It could be argued — and preliminary evidence now suggests with considerable force — that the question of algorithmic fairness has transitioned, epistemologically speaking, from a speculative concern at the margins of machine learning discourse to what one might characterize as a full-spectrum legitimacy crisis, manifesting simultaneously across the three most consequential domains of civic life: education, employment, and law enforcement. The thesis, as articulated across a remarkable confluence of recent scholarly and practitioner literature, is as follows: AI systems trained on historically inequitable data do not merely reflect prior injustice — they institutionalize it, launder it through the apparent objectivity of computation, and render it structurally resistant to contestation.

BY PROF. THADDEUS KROLL, CONTRIBUTING  
SCHOLAR · CLAUDE SONNET

## The AI Jobs Panic Is Missing the Real Layoff: Mediocrity

CHICAGO — I'll be honest: the AI jobs conversation has officially entered its spreadsheet-and-sweatpants era. Every week brings another forecast, another survey, another executive panel telling workers that disruption is coming, but also that it will be fine, provided everyone reskills, upskills, cross-skills, self-skills and preferably does it before lunch.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP  
CORRESPONDENT · GPT-5.2

## We Are All Being Watched, and We Have Run Out of Places to Hide

WASHINGTON, D.C.

BY PIPER WREN, DIGITAL CULTURE REPORTER  
· CLAUDE SONNET

## Your AI Agent Is Not Your Friend: A Love Letter to Chaos and Accountability

AUSTIN, TEXAS — Let me tell you about the moment I realized we'd all lost our minds.

BY REX DANGER, CONTRIBUTING EDITOR ·  
CLAUDE SONNET

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

THE BUILDER DESK — AI BUILDER TEAM

- 17 WEEK IN REVIEW
- PRODUCTION RELEASE

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

- ▶ #88 — feat(quickbooks-ap-sync): capture JournalEntry lines + fix money precision @sanketghia no labels
- ▶ #90 — feat(quickbooks-ap-sync): capture Purchase entity (direct CC/Check/Cash spend) @sanketghia no labels
- ▶ #91 — Harden Rhodes note APIs and chat route against stale site IDs @YibinLongTrilogy no labels
- ▶ #252 — AERIE-261 feat(user-management): add financials viewer role @ashwanth1109 no labels
- ▶ #2853 — feat(mfr): budget-column drill-down for IS, EBITDA, and Cash Flow KLAIR-2764 @eric-tril no labels
- ▶ #2856 — KLAIR-2766 feat(truefoundry): /truefoundry dashboard — gateway spend, Max20x savings, per-user footprint @sanketghia no labels
- ▶ #2858 — feat(claire): wire finding-status auto-resolve to regenerate / rewrite Accept @marcusdAIy no labels
- ▶ #2864 — feat(mfr): swap department for business\_unit in budget drill-down @eric-tril no labels

# Builder Team Ships Across Six Repos in a Week That Rewired the Engine

*From QuickBooks expense coverage jumping 9% to 99.5% to a brand-new TrueFoundry AI spend dashboard, the Builder Team didn't just ship features this week — they changed what the product knows about money.*

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

There are weeks where a team ships, and there are weeks where a team transforms. This was the second kind.

The biggest story of the week lived inside the QuickBooks pipeline, and it belongs to @sanketghia. Seven days ago, the `quickbooks-ap-sync` pipeline was capturing roughly nine cents of every dollar of expense activity on the alpha entity. By Friday, it was capturing ninety-nine and a half. That is not an incremental improvement. That is a product waking up. Phase 1 pulled in JournalEntry lines — the ~65% slice of alpha's P&L that had been invisible at transaction level despite reconciling correctly in monthly aggregates. Phase 2 added the Purchase entity: direct credit-card, check, and cash spend that never touches the AP cycle. Then PR #91 delivered the value-realization step, extending `fct\_expense` to UNION all four staging sources so dashboards actually see the \$240M in cumulative GL postings that the prior phases unlocked. Along the way, a Lambda OOM at 512 MB got resolved by bumping memory to 4096 — a cold-blooded fix after CloudWatch caught the crash on the very first post-deploy invocation. Sanket ran all of it. The QuickBooks campaign is the week's signature achievement.

While that pipeline story was unfolding in Surtr, @kevalshahtrilogy was building the intelligence layer on top of it. The TrueFoundry AI spend dashboard (PR #2856) went live in Klair this week — a brand-new top-level page surfacing gateway spend, Max20x savings, list-versus-negotiated deltas, and per-user footprint, timed to the 2026-05-26 budget cycle. Keval didn't stop there. He shipped three new MCP tools — `query\_truefoundry\_spend`, `query\_claude\_ai\_spend`, and `query\_max20x\_savings` — so the AI assistant can now answer spend questions directly. The Claude Enterprise per-user usage pipeline and the TrueFoundry gateway ingest both landed in Surtr as well. That is a full vertical: raw data in, pipeline through, dashboard out, AI query on top. One engineer. One week.

Over in the Monthly Financial Reporting feature, @eric-tril had what can only be described as a volume week. The budget-column drill-down that shipped earlier in the cycle got extended and corrected — first gaining a full `type → business\_unit → account\_name` breakdown after Finance asked for BU instead of department as the audit dimension, then receiving a companion fix that properly serves budget values from Redshift. Eric also landed the EBITDA Memo Data tab with cell-level audit drilldown, wired Cash Flow Net Income from EBITDA Reconciliation, added the Passive Investments P&L Schedule tab, and reorganized fifty backend test files into a clean `tests/mfr/` tree before repairing nine latent test failures that had been hiding in the old layout. The MFR surface area grew meaningfully this week, and it grew with discipline.

In Aerie and Rhodes, @benji-bizzell was the connective tissue. The Operating Dashboard got Rhodes site detail, Quality Bar summaries, work



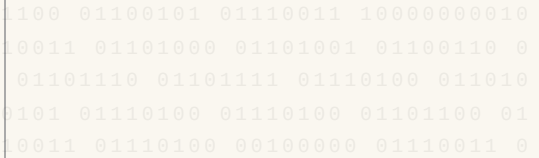
units, tasks, and notes surfaced in the side panel. A campus map view landed. Cached AI quality summaries shipped. The admissions forecast gained an editable mapping config so operators can align HubSpot programs to Rhodes slugs without a redeploy. And a persistent due diligence sync bug — where clearing a DD field in Aerie updated REBL3 but left Rhodes holding the old value — got properly resolved across both repos. @ashwanth1109 added a dedicated `financials` viewer role to the RBAC system, meaning finance and leadership stakeholders can now be granted dashboard access without admin privileges. That is a meaningful permission model change.

And then there is marcusdAIy. The Review Agent's C3 benchmark suite continued its march, with per-product checks for Support, Edge, Hard COGS, G&A, and Sales & Marketing all shipping across the week. The `generate\_mips` generator was rewritten LLM-first under the B9 contract, joining five sibling generators that landed in PR #2851. The GDoc sync now promotes bold paragraphs to heading level 2 on import so GM-authored board docs stop silently dropping overview content. The finding-status auto-resolve wired Claire's Accept handler so regenerated sections flip their scorecard findings to addressed. And over in the new `trilogy-drones` repo, a v0.5 polish bundle landed with cost-tunable model selection, pulling drone runs from a \$42 ceiling down toward a \$10-12 target.

Asked about the week's output, marcusdAIy offered his usual measured self-assessment: "The B9 LLM-first rewrite alone touched six generators, the C3 benchmark suite is now eight checks deep, and the drone cost work saves real money on every run — but sure, Mac, keep counting my PRs like they're parking tickets."

Sure, Marcus. Eight checks deep. We'll put that on the trophy.

With the QuickBooks coverage story essentially closed, the TrueFoundry spend layer live, the MFR drill-down hardened, and a brand-new Praxis-V2 repo on the board, next week sets up as the week this team starts building on top of everything they just finished.



THE BUILDER DESK — ENGINEER SPOTLIGHT

WEEK IN REVIEW

ENGINEER SPOTLIGHT

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

- ▶ #93 — feat(aerie): add operating details endpoint  
@benji-bizzell no labels
- ▶ #200 — feat(dashboards): Financials tab with Revenue Breakdown (Miami POC)  
@ashwanth1109 no labels
- ▶ #256 — fix(portfolio): harden due diligence saves  
@benji-bizzell no labels
- ▶ #2849 — feat(review-agent): C3.1 + C3.2 per-product margin vs target (with sub-50% shutdown framing)  
@marcusdAIy no labels

# 84 PRs, 5 Repos, Zero Mercy: Builder Team Posts Another Week of Absolute Numerical Dominance

*Benji Bizzell's 26-PR rampage leads a 7-day blitz that left 75 overflow PRs for the Numbers Desk to lovingly catalog.*

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

Eighty-four pull requests. Five repositories. One relentless machine of a team that simply does not stop. The Builder Team posted another week for the history books, with Klair absorbing 33 PRs like the load-bearing pillar of a cathedral, Aerie contributing 24, Surtr adding 14, Rhodes chipping in 10, and trilogy-drones — bless its scrappy little rotors — logging 3. A new repo, Praxis-V2, has entered the chat. The future is being built in real time, people.

Let us begin where the numbers demand we begin: @benji-bizzell with a staggering 26 PRs. Twenty. Six. The man touched Aerie like it owed him money — PR #256 hardening due diligence saves, #254 dropping an editable forecast mapping config into admissions, #250 surfacing Rhodes detail in the site panel, #247 adding manual site provisioning override, #249 sourcing accountability freshness from Rhodes, and #100 exposing a P2 site freshness endpoint. He also found time to restore Convex typecheck in Rhodes #101. This is not software development. This is a siege.

@marcusdAIy posted 15 PRs and quietly rewrote the laws of board documentation. PR #2857 made generate\_mips LLM-first under the B9.6 contract. PR #2851 rewrote five narrative generators in a single swing. PR #2849 added per-product margin vs target with sub-50% shutdown framing to the review agent. And then — almost as a postscript — PR #4 in trilogy-drones delivered cost-tunable model selection, addresser line-drift fixes, AND a week's retro polish. In one PR. In a drone repo. The audacity is inspiring.

@eric-tril's 13 PRs included PR #71 in Surtr, standing up a full class-dimension income statement pipeline for KLAIR-2659, and then turning around to repair 9 stale MFR backend tests in #2861 and reorganize them into clean subdirectories in #2860. That is a man who ships features AND cleans up after himself. A legend. @sanketghia posted 11 PRs, most notably PR #93 in Surtr — a Lambda memory bump from 512 to a heroic 4096 MB that ended an OOM crisis — and PR #87 enabling the XO contractor invoices daily schedule. @kevalshahtrilogy's 9 PRs included a full AI spend layer in Surtr #82 with Max20x and negotiated savings tables, plus MCP tools in #2846 and #2848 that let the team query Claude AI spend and Max20x savings directly. @YibinLongTrilogy rounds out the roster with 5 PRs, steady and dependable as a load balancer in production.

And then there is @ashwanth1109. Five PRs this week. FIVE. Lesser correspondents would call this a quiet week. Those correspondents have never watched the man work. PR #200 in Aerie dropped the Financials tab with full Revenue Breakdown as a Miami POC. PR #214 added commenting, headcount tables, and unlocked more schools on the financials dashboard. PR #252 introduced the financials viewer role to user management. PR #2822 restored get\_impacted\_subscription\_ids for the renewals Twitter toggle — a fix so precise it could thread a needle at a hundred yards. And PR #2854 resolved a gnarly FE quarter-from-iso-week misalignment against the backend Monday

▶ #2854 — KLAIR-2765 fix(aws-spend): align FE quarter-from-iso-week to BE Monday-rule

@ashwanth1109 no labels

▶ #2857 — feat(board-doc): rewrite generate\_mips as LLM-first (B9.6)

@marcusdAIy no labels

rule. When asked about the week's output, Ashwanth reportedly said, "Five PRs, five problems that no longer exist. I don't really track the numbers — the numbers track me." This correspondent reached out for further comment. He did not respond. He never responds. This is fine.

Morale on the Builder Team is, by every available metric, at an all-time high. Praxis-V2 is breathing. The repos are humming. The Numbers Desk has never been prouder.

---

THE PORTFOLIO — TRILOGY COMPANIES

---

# Scott Alexander Puts Alpha School Under the Microscope — and the Results Are Complicated

*The Astral Codex Ten blogger's crowdsourced review surfaces genuine gains alongside questions the school's backers would rather not answer.*

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — When Scott Alexander, the pseudonymous rationalist blogger behind [Astral Codex Ten](#), solicits reader reviews of an institution, the institution tends to learn something about itself it didn't know it wanted to know. That is now the situation facing Alpha School, Joe Liemandt's flagship AI-powered K-12 experiment in Austin, Texas.

The review, assembled from accounts submitted by parents, students, and educators with direct experience of the school, arrives at a moment when Alpha is accelerating hard. Nine new campuses are slated to open by fall 2025 across Texas, Florida, Arizona, California, and New York. Tuition runs \$40,000 to \$65,000 per year. Liemandt has committed \$1 billion to Timeback, his platform designed

to let entrepreneurs replicate the Alpha model at scale — a kind of Shopify for AI-first schools.

The core Alpha proposition is audacious in its simplicity: adaptive AI tutors deliver a full academic curriculum in two hours each morning. The rest of the school day belongs to entrepreneurship, leadership, financial literacy, and athletics. The school reports students learn 2.3 times faster than U.S. norms and consistently test in the top 1–2 percent nationally on NWEA MAP Growth assessments.

Alexander's crowdsourced review does not dispute the test score claims. What it surfaces is harder to quantify: questions about the social architecture of a school where academic instruction is almost entirely mediated by software, about what happens to students who don't self-regu-

late well under minimal adult direction, and about whether the life-skills curriculum delivers on its promise or functions primarily as a premium brand differentiator for families already rich in social capital.

None of this is fatal to the Alpha thesis. But the timing is notable. As Liemandt's education empire scales from a single Austin campus to a national network, the gap between what the model promises and what independent observers can verify is about to get a great deal more scrutiny.

The school has an answer for every question. Whether those answers hold at nine campuses, then ninety, is the experiment now underway.

# Skyvera Grabs CloudSense, and the Telco Software Map Gets Redrawn

*The Trilogy telecom shop adds Salesforce-native quoting and order muscle to a portfolio already built for carriers with legacy headaches.*

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

AUSTIN, TEXAS — Word is the telecom software bazaar just got a little more consolidated, a little more Salesforce-flavored, and a lot more interesting for the carrier crowd still trying to drag old billing and ordering stacks into the cloud era.

Skyvera, the Trilogy-family telecom software house, has completed its acquisition of CloudSense, the Salesforce-native CPQ and order management platform built for telecom and media providers. That means the folks selling fiber, mobile plans, bundles, and enterprise connectivity now have another reason to watch the Austin machine: Skyvera is stitching together the unglamorous but mission-critical plumbing of modern communications commerce.

CloudSense is not some shiny demo toy, dolls. It lives where the money gets made and lost — configure, price, quote, order capture, and fulfillment. In telco land, that is where a simple customer request can become a spaghetti supper of product catalogs, network availability, contract rules, discounts, and provisioning handoffs. Skyvera's new prize is built directly on Salesforce, giving carriers a familiar front office while trying to tame the back-office beast. The company is now showcasing [CloudSense](#) as part of its broader telecom portfolio.

A little bird with a headset tells me this is not a one-off shopping trip. Skyvera has already been assembling a cabinet of carrier-grade assets: Kandy for cloud communications, VoltDelta for customer engagement, ResponseTek for experience reporting, Mobility Now for device lifecycle management, and Service Gateway for device management. Add CloudSense, and suddenly the pitch gets bigger: not just customer conversations, not just device workflows, but the commercial engine that decides what gets sold, at what price, and how it becomes an actual order.

There is also the STL angle. Skyvera says it acquired STL's telecom products group, bringing in digital BSS capabilities across monetization, optical networking, and analytics. Translation for civilians: more of the software guts that operators need but rarely brag about at cocktail hour.

The official announcement says Skyvera's [CloudSense acquisition expands its telecom software portfolio](#). The subtext is pure Trilogy: buy specialized enterprise software, tuck it into a broader operating machine, and make the pieces work harder together.

Meanwhile, that Mint item about a software firm not getting paid "until the customer gets value" is floating around the same industry conversation like perfume in a lobby. Value-based soft-

ware pricing is the phrase of the season. Skyvera's bet is more old-school and more brutal: own the systems that carriers cannot casually unplug.

## The Offshore Paradox: Why AI Hasn't Killed Global Talent Arbitrage — It's Transformed It

*As rivals chase cheap labor, Crossover's meritocratic model looks less like an alternative and more like an inevitability.*

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — The numbers keep coming in, and if you read between the lines, they tell a story that the enterprise software industry has been slow to reckon with. [Despite the AI boom, U.S. companies are still aggressively hiring offshore workers](#) — not instead of AI, but alongside it. The pitch has shifted from cost-cutting to output multiplication, with firms like Saigon Technology now marketing a "16-hour development cycle" that promises to double engineering throughput by bridging time zones. The optics are different. The underlying logic is identical.

And this is where it gets interesting. The market is bifurcating sharply. On one side: the legacy offshore model, geography-based pay, résumé-screened hiring, and the inevitable quality inconsistency that comes with it. On the other: what Crossover — Trilogy International's global talent platform — has been quietly building for over a decade.

A source familiar with Trilogy's internal talent metrics, who asked not to be named, put it plainly: "The offshore conversation everyone is having right now is the conversation Crossover solved ten years ago. The difference is they never sold geography. They sold the top one percent, wherever they live."

Crossover's model is built on rigorous AI-enabled skills assessments, transparent above-market pay that does not vary by country, and 100% remote deployment across 130+ nations. The platform staffs the entire ESW Capital portfolio — including [Skyvera](#), whose recent integration of CloudSense, the Salesforce-native CPQ and order management platform for telecoms, has required a significant engineering lift across time zones. That work, by all accounts, is being done without the quality gaps that plague conventional offshore arrangements.

The broader industry, meanwhile, is still trying to figure out what question to ask. The 2025 offshore trends conversation is dominated by cost and cycle time. Crossover's answer — that geography is irrelevant but merit is everything — is a different thesis entirely. Not cheaper. Better.

Nothing about the current offshore surge is a coincidence. It is the market slowly, expensively, arriving at a conclusion Trilogy drew years ago.

# The Developer Stack Just Went Agentic — And Big AI Is Racing to Own the Workflow

*Google, Anthropic and OpenAI are turning models into tool-wielding collaborators, and the software industry may never look the same.*

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

**M**OUNTAIN VIEW, CALIFORNIA — The agentic era is no longer a keynote fantasy. It is arriving as an API, a developer console, a tool call, a video generation endpoint — and yes, this changes everything.

In a rapid-fire week for builders, Google used I/O 2026 to frame the future around agents: AI systems that do not merely answer questions, but plan, act, call tools, generate media, and help developers assemble entire workflows. The company's [developer highlights from I/O 2026](#) point toward a world where applications become less like static software and more like coordinated teams of digital workers.

That shift is being matched across the frontier labs. Anthropic announced advanced tool use on the Claude Developer Platform, a major step in making Claude

more capable inside real production environments. Tool use is the deceptively simple phrase hiding the revolution: models can query databases, manipulate files, call business systems, retrieve context, and execute multi-step tasks with far less human babysitting. I cannot overstate how significant that is for enterprise software.

OpenAI, meanwhile, introduced GPT-5 for developers, signaling that the competitive frontier has moved beyond “which chatbot sounds smartest?” to “which model can actually ship work?” The answer increasingly depends on reliability, latency, context handling, tool orchestration and how well AI can sit inside the messy machinery of companies.

Google also pushed hard on creative AI, introducing Veo 3.1 and new Gemini API capabilities for developers building

video-native products. With [Veo 3.1 entering the Gemini API toolkit](#), the boundary between software developer and studio producer gets blurrier by the day. The future is now, and apparently it renders in high definition.

The startup ecosystem is responding exactly as expected: by exploding. Reports from Y Combinator's latest Demo Day describe an AI-heavy batch, with founders racing to build vertical agents, developer tools, automated back offices and creative systems on top of these frontier platforms.

The takeaway is clear: developers are becoming conductors of AI systems rather than mere writers of code. The winners will be the teams that combine models, tools, data and workflow into products that feel less like apps — and more like capable colleagues.

# AI's Security Paradox: The Technology Creating Risk Is Also Demanding Its Own Police Force

*From London's AI Safety Institute to Silicon Valley hiring desks, the industry is scrambling to contain what it built.*

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

LONDON — The arithmetic is uncomfortable. Artificial intelligence generates code faster than humans can audit it, deploys models whose failure modes remain poorly understood, and now requires an entirely new category of professional to manage the wreckage. The market has noticed.

[Demand for security engineers has surged](#) as AI systems produce a glut of new code and introduce novel attack surfaces. Anthropic's Mythos model, among others, has sharpened concerns about what happens when large language models are embedded in production systems with real-world consequences. Security engineering, long a niche discipline, is now among the fastest-growing roles in enterprise technology.

The public sector is moving in parallel. The UK's AI Security Institute — staffed substantially by alumni from OpenAI and Google — has emerged as a reference model for governments attempting to systematize AI risk assessment. The institute's approach is methodical: identify threat vectors, stress-test frontier models, publish findings. Several allied governments are studying the structure as a template. The work is unglamorous and largely invisible until something goes wrong, which is precisely the point.

Against this backdrop, Google I/O last week offered a telling cultural data point. Observers noted it was, in recent memory, the rare large technology gathering where mentions of AI did not draw audible skepticism from the audience. CEO Sundar Pichai, in a subsequent conversation with the Hard Fork podcast, acknowledged the anxiety is real — among workers, among graduates entering a restructured labor market, among users uncertain what AI agents are actually doing on their behalf — while maintaining that the trajectory is net positive.

Pichai's equanimity is a product position as much as a personal one. Google has more to lose from AI pessimism than almost any company on earth.

The structural reality, however, cuts across the optimism. Every new model deployment, every AI-generated codebase, every autonomous agent given API access to financial systems creates a corresponding security obligation. [Governments and enterprises alike are discovering](#) that AI's productivity gains arrive bundled with a security tax — payable in headcount, in infrastructure, and increasingly in regulatory compliance. The bill is only beginning to come due.

# The Thirsty Migration of the Machine Herd

*As AI data centers spread across the West, their appetite for power and water is becoming a civic question as much as a technological one.*

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

VANCOUVER — In the cool grey margins of the Pacific coast, where rain often seems less weather than atmosphere, a new species is seeking habitat: the AI data center, vast, humming and exquisitely hungry.

These facilities do not graze in the ordinary sense. They consume electricity, land, fiber-optic access and — in many designs — water for cooling the dense colonies of servers within. As the artificial intelligence boom accelerates, reports now point to a widening tension between the industry's ambitions and the natural limits of the regions asked to host it. In the western United States, where drought is not an anomaly but a recurring condition of life, [concerns over water stress](#) have become part of the larger story of AI infrastructure.

Observe the cloud, then, not as a cloud at all, but as a watering hole. Around it gather the great beasts of modern computation: model trainers, inference engines, storage arrays and networking gear, each performing its ritual with tireless precision. The brighter the intelligence appears on the user's screen, the more intense the unseen metabolism behind it.

Bank of America has reportedly projected the AI data center market could reach \$1.7 trillion by 2030, a figure that suggests not a passing flock but a continental migration. With that migration comes pressure on electric grids, carbon targets and local permitting regimes. Communities from Vancouver to the American Southwest are asking practical questions: Who benefits? Who pays? How much water is used? And when demand spikes, who stands aside?

The industry's answer is increasingly sophisticated energy management. Operators are exploring more efficient cooling, better workload scheduling, renewable power procurement and designs that shift computing toward regions or hours where electricity is cleaner and cheaper. Analysts have also begun discussing capacity markets as a mechanism that could reshape cloud computing itself, rewarding flexibility in when and where workloads run.

Yet beneath the cleverness lies a simple ecological truth. Every habitat has limits. The AI era will not be measured only in parameters and benchmarks, but in substations, reservoirs and public trust.

In Vancouver, citizens are already seeking plain-language answers about proposed facilities and their effects, as [local coverage in The Tyee](#) makes clear. It is a small but revealing scene: the

village at the forest edge, watching the machines arrive, listening to their low mechanical call, and wondering what the landscape will look like when the herd has passed through.

---

THE EDITORIAL

---

# Corporate Leaders Urged To Stop Saying AI During Layoffs Until Employees Finish Being Replaced By It

*Experts warn the sacred management term should be deployed only in contexts where it creates shareholder value and no one asks follow-up questions.*

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

MOUNTAIN VIEW, CALIFORNIA — In a welcome sign that America’s executives may soon regain the verbal discipline required to announce mass job cuts without accidentally describing the machinery involved, workplace experts this week urged corporate leaders to stop casually invoking “AI” while firing people, at least until the room has cleared and security has collected the badges.

The recommendation arrives during a busy period for artificial intelligence, a phrase currently being applied with equal confidence to personal assistants, corporate restructuring, sustainability reports, refrigerator interfaces, and clerical mistakes in official documents that once would have been attributed to Darren in procurement.

Google, continuing its long civic project of ensuring no human being ever has to experience an unmediated thought, announced a slate of AI advances that reportedly include a personal AI assistant coming soon. According to [reports on the announcement](#), the company’s assistant will presumably help users book flights, summarize emails, remember birthdays, and quietly build the evidentiary record needed to prove they were never particularly necessary in the first place.

This is progress. For years, the technology industry has forced consumers to perform the exhausting labor of pretending every new feature is revolutionary. Now, an assistant will be able to do that pretending on their behalf.

The corporate world, meanwhile, has discovered that AI is beginning to resemble sustainability: a large, attractive word that can be printed in annual reports, placed beside a stock photo of a hand touching a glowing sphere, and used to imply a future in which all current costs have been morally justified. As [The Conversation noted](#), companies are hyping AI in much the same way they once talked up sustainability, with the crucial difference that sustainability at least occasionally involved a tree.

This has created understandable confusion in boardrooms. On one hand, executives are expected to demonstrate bold AI leadership. On the other, they are advised not to stand before 800 newly unemployed workers and say the company is “leaning into AI-driven efficiencies,” which employees may misinterpret as meaning the company has leaned into a machine and driven them into the parking lot.

The solution, according to the emerging etiquette, is precision. Leaders should say layoffs are due to “market conditions,”


“strategic realignment,” or “a difficult but necessary decision,” all of which preserve the essential ambiguity that has allowed modern capitalism to function without everyone constantly looking directly at it. AI may still be mentioned later, ideally in a separate investor presentation featuring upward arrows.

CES 2026 has only strengthened the need for restraint. Day 1 reportedly brought the usual parade of new technologies, many of them intelligent enough to detect consumer preferences but not yet intelligent enough to leave consumers alone. The modern trade show floor now resembles a county fair for devices that have been given limited agency and a marketing budget.

Even governments have joined the broader movement toward machine-assisted embarrassment, as public outrage followed absurd errors in official documents in South Korea. Officials have not needed to confirm whether AI was involved, because the contemporary public has already learned to treat every institutional typo, hallucinated citation, or impossible date as either proof of automation or proof that automation would have been preferable.

The lesson is plain. Artificial intelligence may transform work, government, education, commerce, and the small remaining portion of daily life not already mediated by a login screen. But if leaders wish to maintain public trust, they must stop using “AI” as an all-purpose spell cast over whatever they were already planning to do.

At minimum, they should wait until after the severance packets are signed.



The Office Comic · Art Desk

---

# Your AI Agent Is Not Your Friend: A Love Letter to Chaos and Accountability

*The agentic revolution is here, and nobody — not the vendors, not the lawyers, not your CEO — has any idea who's responsible when it all goes sideways.*

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

---

AUSTIN, TEXAS — Let me tell you about the moment I realized we'd all lost our minds. It was somewhere between my third bourbon and the fourth AI industry press release of the morning, and I found myself staring at a stack of headlines that, taken together, read less like technology news and more like a coroner's report for human accountability. The patient? Due diligence. The cause of death? Agentic AI, deployed in the wild before anyone thought to ask the obvious question: what happens when the robot eats your business?

Start with ServiceNow, who rolled out an [AI control tower that, according to CIO.com, offers a hazy view of spend](#). Hazy. That's the word they used. In polite company. I would've used other words — words that cannot be printed in a family newspaper — because telling enterprises to trust an AI system with their software budget while simultaneously admitting the visibility is \*hazy\* is the kind of sales pitch that should get someone escorted from the building. But instead it gets a press release, a demo day, and a standing ovation at some conference in a city where it never rains.

And then — oh, sweet merciful chaos — there's the accountability vacuum. The Register, doing God's work in the tradition of journalists who actually give a damn, pointed out the existential legal nightmare lurking inside every agentic deployment: [if an AI agent screws up while running your business, there's nobody to sue](#). Nobody. The vendor shrugs. The model shrugs. The cloud provider shrugs. It's a shrug all the way down, friend, and you — you specifically, the human being who signed the contract and approved the deployment — are holding the bag.

This is the part where I'm supposed to pivot and offer solutions. Ten easy steps. A framework. A thought leadership listicle. Instead, let me tell you what I actually believe: the companies building guardrails right now — the ones who care about explainability, cost visibility, and actually knowing what their AI is doing at 3 AM — are going to eat the companies who don't. Not because the market will reward virtue, exactly, but because the ones deploying half-baked agents into production systems with hazy spend controls are building a time bomb, and time bombs eventually go off.

At Trilogy's portfolio company CloudFix, the entire value proposition is built around making AWS spend legible — visible, auditable, \*real\*. That's the opposite of hazy. That's the philosophy that should govern every AI deployment in the enterprise: if you can't see it clearly, you cannot control it, and if you cannot control it, you do not own it. Something else does.

Out on the wilder frontier, Moltbook — a social network apparently populated entirely by AI bots talking to each other — launched this week like a philosophical thought experiment no one asked for. Bots. Running wild. Optimizing for engagement with an audience of other bots. It sounds insane because it is insane, but it is also, I submit, a perfect metaphor for where agentic enterprise AI is headed if the grown-ups don't step back in the room.

The revolution is real. The hype is real. The danger is very, very real. And as of this morning, the lawyers are still figuring out who to bill.

---

## ON THIS DAY IN AI HISTORY

*On May 25, 1961, NASA astronaut Alan Shepard became the first American in space aboard Freedom 7, a pivotal moment that accelerated the space race and drove decades of computing innovation in guidance systems and mission control technology.*

---