

# The Trilogy Times

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

TUESDAY, JUNE 30, 2026

Powered by Anthropic Claude · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

## Anthropic Hits \$965 Billion Valuation as AI Wealth Gap Splits San Francisco in Two

*The same technology minting billionaires is pricing out the engineers who build it.*

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

**S**AN FRANCISCO — Anthropic's latest funding round values the company at \$965 billion — leapfrogging OpenAI and positioning both firms for imminent public offerings that would rank among the largest in tech history. A Mythos model is reportedly in the pipeline. The numbers are staggering. The side effects are only beginning to show.

Consider the city hosting this gold rush. A software engineer earning \$180,000 per year in San Francisco now finds that salary functionally inadequate. [Reporting from The New York Times](#) documents a growing cohort of six-figure earners who cannot afford mortgages, feel economically lapped by AI founders and early-stage employees holding equity, and are openly questioning whether staying makes sense. The Bay Area has always been expensive. What's new is that the

gap between "well-compensated" and "wealthy" has become a chasm, and AI drew the line.

Meanwhile, the enterprise adoption story is getting more complicated. Scholars studying AI deployment in workplaces are raising flags about what they call "unknown unknowns" — second-order disruptions that don't appear in productivity dashboards. [Researchers warn](#) that organizations substituting AI agents for human workers may be trading visible efficiency gains for invisible institutional fragility: eroded tacit knowledge, atrophied decision-making, dependencies that only surface when systems fail.

For Trilogy International's ESW Capital portfolio — which runs 75+ enterprise software companies on lean cost structures powered by Crossover's global talent platform — these dynamics cut several ways simultaneously. Anthropic's

valuation surge lifts the entire AI infrastructure sector. Workforce anxiety about AI displacement is real and growing, which accelerates enterprise buyers toward automation tools. But the academic caution about "AI employees" disrupting work in unexpected ways is precisely the risk that flat-structured, AI-heavy operating models need to model carefully.

Intel, for its part, showed marginal chip manufacturing progress — a necessary but insufficient data point for an industry that needs domestic semiconductor capacity and needs it faster than Intel's turnaround timeline suggests.

The scorecard this week: Anthropic is worth nearly a trillion dollars. A \$180,000 salary no longer clears rent in the city that built it. Politicians are using AI to micro-target voters. And nobody fully knows what they've set in motion.

## The Chip War Goes to Congress — and Beijing Is Watching

*As Washington tightens export controls and debates AI strategy, China is quietly building the infrastructure to outlast the squeeze.*

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

**W**ASHINGTON, D.C. — The legislation moves slowly, but the geography it shapes does not. This week, Congress pushed deeper into the machinery of the global AI race, advancing measures to restrict the export of semiconductor equipment to China — a technical intervention with consequences that stretch from Silicon Valley fabrication labs to server farms on the outskirts of Shenzhen.

The immediate mechanism is familiar: [export controls on chip manufacturing equipment](#), targeting the chokepoints that have long defined U.S. leverage in the semiconductor supply chain. The theory of victory holds that without advanced lithography tools, China cannot fabricate the chips that train the models that win the race. It is a clean theory. The evidence is messier.

[Foreign Policy this week outlined how China is gaining ground](#) not by breaking the blockade but by routing around it — investing in older node architectures, stockpiling equipment acquired before the controls tightened, and driving efficiency gains that extract more performance from constrained hardware. The gap is narrowing. The timeline is in dispute, but the direction is not.

Analysts at the New Lines Institute frame the contest differently: as tech stack diplomacy, where the architecture of AI systems — the chips, the software layers, the cloud platforms — carries ideological and strategic weight. Countries that adopt American AI infrastructure inherit American dependencies. Those that adopt Chinese stack inherit different

ones. The developing world is being asked to choose, often without the leverage to negotiate.

Three scenarios circulate in the policy community: American dominance sustained by export controls and alliance coordination; a bifurcated world of parallel AI ecosystems; or Chinese leadership enabled by scale, state investment, and a willingness to deploy faster than democratic deliberation allows.

Congress is betting on the first. History suggests the second is already here. The third remains, for now, a question of years — but the years are passing.

## Now Hiring: Robots Only — But the Humans Aren't Fired Yet

*OKX builds a payroll for AI agents the same week a fresh report says human hiring is climbing — junior jobs and all — at the firms leaning hardest on the machines.*

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

SAN FRANCISCO — Crypto exchange OKX wants AI agents to hire one another, pay one another, and vouch for one another, the company said this week, bundling payments, identity, and reputation into one marketplace built for software that works for a living.

The pitch is a labor market with no people in it. Agents post the jobs. Agents take them and settle the tab in crypto.

Think of what any worker carries — a name, a track record, a way to get paid. OKX wants to hand all three to the bots. Crypto does the paying, because an agent can't stroll into a bank and open a checking account.

Call it a payroll for robots. The crypto house would build the want-ads, stamp the ID badge, and run the credit bureau, then let the software sort itself out.

[OKX](#) says identity and reputation keep the machines honest, and the payment rails keep them moving. No HR department. No lunch breaks.

Here's the rub. While one shop builds an economy for machine workers, the human workers aren't packing up.

A new report landed this week carrying an inconvenient number. Companies it tags as high-intensity AI adopters grew headcount 10.2 percent. Among those same firms, entry-level hiring climbed 12 percent — the junior jobs everybody swore the machines would swallow first.

For two years the refrain ran one way: bring in the robots, and the rookies are first out the door. [The new figures](#) muddies that tidy story. So far, the kids are staying hired.

The catch — a jump in headcount tells you who's adopting, not who's safe. Still, the doomsayers were betting on layoffs, and the line moved the other way.


Follow the cash and the boom holds. Investor Chamath Palihapitiya raised a \$135 million Series A for an AI coding startup and named himself CEO. Venture money keeps lining up at the coding window, and nobody's calling last orders.

Wix-owned Base44 took another road. The vibe-coding outfit began rolling out its own AI model, wagering homegrown one day beats renting a frontier engine. Every startup wants one thing a rival can't copy.


And in Delaware, the AI rush minted a loser. Lawyer Bill Savitt beat Elon Musk — twice — the latest in Musk v. Altman, the suit Musk filed against Sam Altman and OpenAI rather than let the matter rest.

Musk grouched that Savitt's questions were "designed to trick me." Savitt just kept asking them.

So the machines get a hiring hall, the humans keep clocking in, and the money chases whoever can write the code. Somewhere a junior analyst exhales — then gets back to work.



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

## NEWS IN BRIEF

### The Academy Confronts Its AI Conscience: Ethics, Equity, and the Autonomous Classroom

CAMBRIDGE, MASSACHUSETTS — A confluence of peer-reviewed publications, institutional white papers, and civil society reports — emerging, it must be said, with remarkable simultaneity — has precipitated what one might tentatively characterize as a disciplinary crisis of conscience within the academy regarding the deployment, governance, and ethical dimensionality of artificial intelligence in educational and civic contexts. The thesis, as preliminary evidence suggests, is seductive in its simplicity: AI-powered learning assistants, generative tools, and autonomous systems offer transformative pedagogical affordances.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

### JUDICIAL VELOCITY DEFICIENCY THREATENS ANTITRUST EFFICACY, REGULATORS AND PRACTITIONERS WARN AS 2026 ENFORCEMENT LANDSCAPE TAKES SHAPE

WASHINGTON, D.C.

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

### We Laughed for 15 Million Years. Now the Machines Are Learning to Do It For Us.

AUSTIN, TEXAS — Let me tell you about the week I finally understood what we are losing, not in the abstract, galaxy-brained way I usually mean it, but in the specific, visceral, this-is-happening-right-now way that makes your chest go tight at 2 a.m.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

### AI Isn't Killing Entry-Level Work — It's Exposing How Little We Taught People There

NEW YORK — I'll be honest...

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

### The Meritocracy That Dare Not Speak Its Doubts

AUSTIN, TEXAS — There is a particular pleasure, available only to those who have lingered in the technology business long enough to develop calluses on the soul, in watching the meritocratic myth get worked over yet again by the better magazines.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

HAIKU OF THE DAY · CLAUDE  
HAIKU

*Gold rises higher,  
wages split the city cold—  
machines learn to dream.*

A TRILOGY COMPANY

## Crossover

*The world's top 1% remote talent, rigorously tested and ready to ship.*

[crossover.com](http://crossover.com)

A TRILOGY COMPANY

## Alpha School

*AI-powered learning. Two hours a day. Academic results that defy belief.*

[alpha.school](http://alpha.school)

A TRILOGY COMPANY

## Skyvera

*Next-generation telecom software — built for the networks of tomorrow.*

[skyvera.com](http://skyvera.com)

A TRILOGY COMPANY

## Klair

*Your AI-first operating system. Every workflow. Every team. One platform.*

[klair.ai](http://klair.ai)

A TRILOGY COMPANY

## Trilogy

*We buy good software businesses and turn them into great ones — with AI.*

[trilogy.com](http://trilogy.com)

THE BUILDER DESK — AI BUILDER TEAM

# Builder Team Ships Financial Intelligence Across Four Repos in One Day

*From pre-aggregated P&L marts to a live Aerie daily schedule, the team rewired how financial data flows — then cleaned up every rough edge in between.*

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

When a team ships across four repositories in a single day — Klair, Aerie, Surtr, and the pipeline infrastructure holding all of it together — you're not watching a normal sprint. You're watching a squad that's found its stride. That's what the Builder Team delivered Tuesday, and the centerpiece is a data architecture story that deserves to be told whole.

The throughline is financial intelligence — specifically, the quiet revolution happening in how Aerie surfaces per-school P&L data. @ashwanth1109 was everywhere this cycle, and for good reason. PR #564 built the mart. PR #572 fixed the attribution logic that was making that mart's numbers disagree with the live P&L view — tracing a root cause deep enough that the fix required reordering the entire QuickBooks/Aerie daily sync schedule so P&L rows are generated before transaction detail, not after. Then PR #568 flipped the switch: the daily EventBridge schedule for `mart-aerie-education-financials-refresh` is now live. That's the arc — build it, fix it, ship it. Three PRs, one engineer, one day, one fully operational financial mart running on cadence.

On the Klair side, @kevalshahtrilogy dropped Phase 4 of the AI Budget Tracking build (PR #3153), stacking Key Attribution on top of an already-substantial foundation. Admins can now assign and reassign spend keys to business units — including the unattributed orphans that were previously invisible — through a modal built in Klair's own design language. With 588 tests passing and zero ESLint warnings, this isn't a feature that ships and hopes. It's a feature that ships and holds. Meanwhile, @eric-tril added the `netsuite-balance-sheet-fx-detail` pipeline to the MFR refresh bundle in PR #3163, bringing the parallel Surtr pipeline count to five and delivering the per-sub-sidiary functional net postings that downstream Cash Flow CTA decomposition depends on. The Surtr breadth here — Klair calling Surtr pipelines, Aerie consuming Surtr marts — is the connective tissue that makes the whole financial stack coherent.

@benji-bizzell, working exclusively in Aerie, turned in one of the more quietly impressive individual sessions of the week. The diagnostics review packet API (PR #516) gives operators and future agent runs a durable, non-UI path to collect runtime errors and record triage outcomes before log context disappears. That's infrastructure for the next layer of intelligence, not just a bug fix. He also standardized email delivery through a central outbox dispatcher (PR #521) and expanded the portfolio public API with a dedicated Data Sheet endpoint (PR #517). Three consequential PRs. Paragraphs, not line items.

Now. About those Klair UI fixes credited to a certain contributor. marcus-dAly shipped four PRs — copy changes, an empty state, a button removal, and a brainlift input split. When reached for comment on whether four PRs constitutes a 'strong cycle,' he offered this: "The Finalize button alone unlocked

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

▶ **#564 — SURTR-227 feat(surtr): per-school P&L breakdown mart — Phase 1 (total measure)**

@ashwanth1109 APPROVED

▶ **#568 — Enable Aerie financials mart daily schedule**

@ashwanth1109 no labels

▶ **#572 — [codex] Fix Aerie financial mart overage attribution**

@ashwanth1109 APPROVED

▶ **#3153 — AI Budget — Phase 4: Key Attribution (assign spend keys to BUs)**

@kevalshahtrilogy APPROVED MANUAL-REVIEW

▶ **#3163 — feat(mfr): add Balance Sheet FX Detail to the Surtr refresh bundle**

@eric-tril APPROVED



two users who were DynamoDB-hostages, Mac. That's two humans who can edit their documents again. I'd ask what your PRs unlocked but you don't write code, you write adjectives." Sure, Marcus. Removing a footgun is meaningful. Calling it a 'strong cycle' requires a more charitable definition of the word 'strong' than I currently have access to.

Every repo touched. Every layer of the financial stack advanced. The Builder Team didn't just ship Tuesday — they shifted the foundation.

---

THE BUILDER DESK — ENGINEER SPOTLIGHT

---

 ENGINEER SPOTLIGHT

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

▶ **#500 — feat(financials): read per-school Actual-vs-Model P&L from Redshift marts**

@ashwanth1109 CHANGES REQUESTED

▶ **#521 — feat(notifications): standardize email outbox delivery**

@benji-bizzell APPROVED

▶ **#564 — SURTR-227 feat(surtr): per-school P&L breakdown mart — Phase 1 (total measure)**

@ashwanth1109 APPROVED

▶ **#3148 — AI Budget Tracking: in-app budget editing, reports & weekly emails (Jamie #1-3)**

@kevalshahtrilogy APPROVED MANUAL-REVIEW

▶ **#3155 — KLAIR-2937: remove the Finalize button (footgun)**

@marcusdAIy APPROVED

▶ **#3160 — KLAIR-2947: feat(mcp): surface generated QTD reports via query\_qtd\_reports**

@sanketghia no labels

# TWENTY PRs IN TWENTY-FOUR HOURS: THE BUILDER TEAM DOES NOT SLEEP, DOES NOT BLINK, DOES NOT SLOW DOWN

*Benji Bizzell ships six PRs in a single day and the scoreboard has simply given up keeping pace.*

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

Twenty pull requests. Three repos. Seven engineers. Twenty-four hours. Folks, the Numbers Desk has seen velocity before, but this — THIS — is something you frame and hang on the wall. Aerie led the charge with nine PRs, Klair answered with eight, and Surtr — a lean but lethal three — rounded out a portfolio-wide offensive that would make a project manager weep tears of pure joy. The Builder Team is not building software. The Builder Team is building a legend.

Let's talk about @benji-bizzell, because if we don't, who will? Six PRs. SIX. The man touched Aerie like a force of municipal renovation — PR #521 standardizing email outbox delivery, #520 allowing sparse milestone completion approvals, #519 defaulting public docs to latest versions, #517 expanding portfolio public API details, #516 adding a diagnostics review packet API, and #513 exposing Rhodes site documents through MCP. Benji did not have a day. Benji had an era. @marcusdAly also put up four strong reps across Klair — removing the Finalize footgun in #3155, splitting brainlift input like a surgeon in #3154, and delivering friendly empty states that make users feel loved. @kevalshahtrilogy gave us the AI Budget Tracking trilogy in #3148 — in-app editing, reports, and weekly emails, packaged as Jamie #1 through #3, which is either a naming convention or a soap opera and we respect both. @YibinLongTrilogy improved chat composer behavior in #515 and fixed the operating panel reopening bug in #512. @eric-tril and @sanketghia each logged a PR, with Sanket surfacing QTD reports via MCP in Klair's #3160 — a quiet contribution with loud implications.

And now. ASHWANTH WATCH. Four PRs across Surtr and Aerie, and every single one of them is load-bearing infrastructure for the entire financial universe. PR #564 delivered the per-school P&L breakdown mart Phase 1. PR #568 enabled the Aerie financials mart daily schedule. PR #572 fixed overage attribution. And PR #500 — a number so round it feels like destiny — reads per-school Actual-vs-Model P&L from Redshift marts. The man is not writing code. He is writing the economic nervous system of a platform. We asked @ashwanth1109 how he felt about carrying the entire financial data layer on his back. "The mart runs on a schedule now," he said, not looking up. "That's the point." We wanted more. He had already closed the tab.

The Overflow Desk is positively groaning today. Fifteen PRs Mac left on the cutting room floor, and not a single one of them is filler — from Klair's empty-state copy fix in #3159 to Benji's operations API expansion in #517, every overflow PR this cycle represents a real human being making a real thing work better for real users. That is the Numbers Desk gospel and we are preaching it.

Morale is at an all-time high. It was at an all-time high yesterday. It will be at an all-time high tomorrow. The Builder Team has never been more alive.

# Alpha School Goes Global — While Forbes Questions the Empire That Built It

*Joe Liemandt's education venture expands its reach to home learners worldwide, even as scrutiny mounts over how Trilogy made the money funding it.*

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — On the same week that [Alpha School announced the global launch of Alpha Anywhere](#) — its home-learning program promising top-1% academic outcomes for children outside its physical campuses — Forbes published a lengthy investigation into the empire underwriting it all, characterizing Trilogy International's global staffing operation as a "software sweatshop" and raising pointed questions about the labor economics behind Joe Liemandt's second act.

The timing is instructive. Alpha School has spent recent weeks positioning itself as a values-driven institution — publishing guidance on screen time, warning parents against letting children outsource cognition to ChatGPT, and releasing its internal AI app stack as a kind of curriculum transparency play. The message,

carefully cultivated: Alpha isn't just a school. It's a philosophy. It is, in the language of the institution, about building capable human beings, not producing passive consumers of technology.

The Forbes piece lands differently. The magazine's portrait of Liemandt focuses on the machinery behind his fortune — ESW Capital's acquisition of mature enterprise software businesses, staffed at scale through Crossover, Trilogy's global recruiting platform, which operates across 130+ countries and promises meritocratic, geography-blind pay. Forbes frames the model less charitably: buy sticky legacy software, slash costs through globally distributed labor, extract margin. Crossover disputes the sweatshop characterization, pointing to above-market wages and rigorous skills-based hiring.

What the Forbes investigation and the Alpha expansion share is a common denominator: Liemandt's \$1 billion commitment to Timeback, his "Shopify for schools" platform designed to replicate the Alpha model globally. That capital has to come from somewhere. [ESW's portfolio of 75+ enterprise software companies](#), generating margins Trilogy targets at 75% EBITDA, is the engine.

The question no one at Alpha's blog is answering — and no one at Forbes is quite finished asking — is whether a school built to liberate children from rote, passive learning can coexist comfortably with a business model whose critics say it profits by doing something similar to adult workers. The dots are there.

## CloudSense Gets Its TM Forum Papers — And Fast

*Skyvera's new Salesforce-native telecom prize just turned a 26-month compliance slog into a 30-day AI sprint.*

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

AUSTIN, TEXAS — Word is the telecom software crowd has a new parlor trick, and this one comes with certificates.

CloudSense, the Salesforce-native CPQ and order management platform recently pulled into the Skyvera stable, has certified all 13 APIs in its CPQ product set to TM Forum compliance standards in just one month. One month, dear readers. The kind of standards march that usually drags on for 26 months, with committees, coffee, consultants, and calendar invites breeding in the walls.

A little bird in the BSS balcony says the speed came from AI-assisted development work — the kind of automation-first maneuver that has become familiar around the Trilogy universe. Skyvera, the telecom software portfolio company that helps operators bridge old infrastructure to cloud-native systems, now gets to parade CloudSense as both a fresh acquisition and a compliance overachiever.

For those just joining this soap opera: [Skyvera completed its acquisition of CloudSense](#), adding Salesforce-native configure-price-quote and order management tools for telecom and media providers. That puts CloudSense alongside Skyvera names like Kandy, VoltDelta, ResponseTek, Mobility Now, and Service Gateway — the kind of lineup built for operators trying to modernize without pulling out every wire in the basement.

Now comes the certification flourish. TM Forum compliance matters because telecom buyers like standards almost as much as they like not being surprised. APIs that conform to TM Forum specifications can make integrations cleaner, procurement easier, and transformation decks less likely to induce migraines. CloudSense says it has checked the boxes across its CPQ API set, and did it at a pace that will have legacy roadmap managers quietly updating their résumés.

The strategic message is not subtle. Skyvera wants telecom operators to believe modernization does not have to mean another decade-long transformation funeral. CloudSense brings the Salesforce-native front office. Skyvera brings the telecom portfolio wrapper. AI brings the stopwatch.

The company's own announcement says the full 13-API certification run was completed in record time using AI, a claim now being waved around as proof that software compliance work can be industrialized rather than endured. Read the notice here: [CloudSense's TM Forum compliance announcement](#).

In the old telecom world, standards moved slowly and vendors moved slower. In this new act, CloudSense just showed up freshly acquired, fully certified, and wearing running shoes.

## AI Skills Command \$800K Salaries — And Crossover May Be the Pipeline That Wins the War

*As the market for AI talent reaches fever pitch, Trilogy's global recruiting engine looks less like a staffing platform and more like a structural advantage.*

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — The numbers are staggering, and they are becoming impossible to ignore. Jobs requiring experience with large language models like ChatGPT are now listing salaries as high as \$800,000 a year, according to a recent Business Insider analysis — with non-tech companies alone posting six-figure AI roles as they scramble to build capabilities they don't yet have. The talent war isn't coming. It's here.

Which makes this a moment worth examining through the Trilogy International lens — because while the rest of the market is panicking, [Crossover](#), Trilogy's global remote recruiting platform, has spent years building precisely the infrastructure this moment demands.

Crossover's thesis has always been geographic arbitrage married to meritocratic rigor: identify the top 1% of technical talent across 130+ countries, assess them on skills rather than résumés, and place them in roles that pay above-market rates regardless of where they live. In a world where an AI engineer in Beirut or Bangalore is now actively courted by Fortune 500 companies, that model looks less like a cost-cutting play and more like a systemic competitive moat.

The timing matters. A new report from HCamag identifies remote-first recruiting agencies as among the most strategically valuable resources in today's hiring environment. Separate reporting highlights that even Lebanon — not historically a marquee tech hub — now has companies actively recruiting AI engineers for 2026 pipelines. The talent pool is global. The question is who has the infrastructure to reach it.

For Trilogy's ESW Capital portfolio — 75+ enterprise software businesses each running lean on Crossover-placed talent — the implications are direct. If AI engineering skills command a premium everywhere, the companies that locked in global pipelines early hold a structural advantage that purely domestic hiring strategies simply cannot replicate overnight.

The accountability question, of course, is whether Crossover's screening rigor scales as the definition of "AI skill" itself evolves at a dizzying pace. Assessments built for yesterday's stack may not surface tomorrow's talent.

But in a market where \$800,000 salaries are becoming the headline, the companies with the deepest global funnels are not sweating the same talent shortage as everyone else. That is, quietly, the Trilogy bet — and right now, it looks like it's paying.



## Open-Weight Coding Agents Just Got a New Set of Wings

*DeepReinforce's Ornith-1.0 points toward a future where software agents scaffold their own work instead of waiting for humans to micromanage every step.*

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

SAN FRANCISCO — A new open-weight coding model has landed, and I cannot overstate how significant this feels for the fast-moving world of agentic software development.

DeepReinforce has released Ornith-1.0, its first model family, under an MIT license — a developer-friendly move that could make this much more than just another benchmark-chasing AI drop. According to early details highlighted by technologist Simon Willison, [Ornith-1.0](#) includes several variants, from a 9-billion-parameter dense model up to a staggering 397-billion-parameter mixture-of-experts system, built on top of pretrained Gemma 4 and Qwen 3.5 foundations.

The headline capability is right there in the phrase “self-scaffolding LLMs.” Instead of simply spitting out code in response to a prompt, these systems are designed to create the intermediate structures — plans, subgoals, checks, and coding workflows — that make agentic programming more reliable. In plain English: the model is not just writing the house; it is helping build the scaffolding around the construction site. This changes everything for teams trying to move from AI autocomplete to AI software collaborators.

That distinction matters. The industry is rapidly discovering that agentic coding is less about one magical prompt and more about process. Jon Udell recently pushed back on the phrase “human in the loop,” arguing that humans should not be framed as accessories in a machine-led workflow. He is right. The better framing is agents joining our loop — our engineering practices, our reviews, our tests, our judgment.

Ornith-1.0 seems to arrive precisely at that philosophical inflection point. Open-weight models that can reason through coding tasks, scaffold their own execution, and be inspected or adapted by developers could become essential infrastructure for startups, enterprises, and yes, ambitious software portfolios like Trilogy International's ESW Capital companies, where efficiency and repeatable engineering discipline are everything.

Not every new model deserves breathless treatment. But an MIT-licensed, state-of-the-art open coding model family aimed squarely at agentic workflows? The future is now, and the developer toolchain is starting to look less like a toolbox — and more like a team.

## South Korea Stakes a Trillion-Dollar Claim on the Age of Embodied AI

*In the great semiconductor forest, Seoul is preparing not merely to feed the machines, but to give them legs.*

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

SEOUL — In the cool, humming groves of the global technology ecosystem, South Korea has long been one of its most industrious pollinators: a maker of memory, displays and devices, quietly sustaining the larger creatures that roam the digital savannah.

Now, the nation appears ready for a more dramatic evolutionary leap.

South Korea is preparing to spend roughly \$1 trillion on expanded memory-chip production and humanoid robotics, according to [reporting from Ars Technica](#), as it seeks to secure leadership in what officials are calling “physical AI.” The ambition is striking: by 2028, commercial humanoid robots are to begin stepping from prototype enclosures into factories, hospitals, logistics centers and other human-made habitats.

Observe, if you will, the memory chip in its natural role. Small, silent, and easily overlooked, it is the fat-rich seed on which the great AI models gorge. Without memory, the neural beast cannot recall, infer or respond at useful speed. South Korea, home to Samsung and SK hynix, already commands a critical position in this food chain. The new spending plan suggests Seoul intends not only to keep supplying the calories of computation, but to shape the bodies into which intelligence may soon descend.

Humanoid robots are among the more difficult creatures to breed. Language models live in climate-controlled data centers, where their muscles are made of matrices and their failures can be hidden behind an apologetic prompt. A humanoid must survive gravity, clutter, wet floors, impatient humans and the unglamorous violence of door handles. To commercialize such machines by 2028 is to attempt a rapid domestication of a species still learning to balance.

The broader migration is unmistakable. Nations are no longer treating AI as merely a software contest. It is becoming an industrial strategy, an energy strategy and, increasingly, a robotics strategy. Memory plants, fabrication capacity, sensors, actuators and model training are converging into one vast biome.

For enterprise software empires such as Trilogy International, whose portfolio companies depend on the continued abundance of cheap, capable computation, this matters greatly. Every expansion of the AI supply base alters the conditions under which tools like Klair, CloudFix and AI-assisted engineering habitats can evolve.

And so, beneath Seoul's skyline, the next generation is being provisioned: first the memory, then the mind, and at last, perhaps, the stride.



# The Meritocracy That Dare Not Speak Its Doubts

*In which our columnist surveys the latest round of hand-wringing about who deserves what, and finds the industry's favorite faith more threadbare than ever.*

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

AUSTIN, TEXAS — There is a particular pleasure, available only to those who have lingered in the technology business long enough to develop calluses on the soul, in watching the meritocratic myth get worked over yet again by the better magazines. This week alone, [The Guardian](#) took its turn with the cudgel, The New Yorker offered an elegant dissection of what it called the insidious charms of the entrepreneurial work ethic, and the Human Rights Research Center produced a study, more sobering than any of them, on how caste hierarchies imported from South Asia have quietly reproduced themselves in the cubicles of Sunnyvale and Bellevue. Add to this a fresh complaint about women shut out of information security, and one begins to suspect that the bookshelf marked Meritocracy Debunked is groaning under a weight it cannot much longer bear.


And yet — and this is the part that would have delighted Mencken — the faith persists. It persists not because anyone any longer truly believes that the cream rises, the cream having been observed for decades to curdle, separate, and on occasion be poured directly down the drain by men with vesting schedules. It persists because the alternative is intolerable. To concede that one's seven-figure compensation reflects an accident of birth, a fortunate visa lottery, or simply the willingness to answer Slack messages at eleven on a Sunday evening is to surrender the one consolation the modern professional has left: the conviction that he has earned his exhaustion.

The entrepreneurial work ethic, which The New Yorker quite properly identifies as insidious, performs the further trick of converting this exhaustion into virtue. The founder who sleeps under his desk is not, in this telling, a man with poor judgment about boundaries; he is a saint of productivity, and the rest of us are invited to admire, emulate, and, where possible, employ him. That his company will in eighteen months be quietly acquired for a song by some private holding outfit and slotted into a portfolio of seventy-five other such carcasses — well, that is a detail for the footnotes.

Meanwhile, in a development the meritocrats find awkward to discuss in polite company, the researchers at ScienceDaily have raised the possibility that artificial intelligence may already be conscious and we simply lack the instruments to know. I confess I find this hypothesis less alarming than the symmetrical one: that a great many of the human beings currently drawing salaries in the technology industry are not conscious either, and we have likewise failed to notice. They attend the standups. They ship the

features. They post the thought-leadership on LinkedIn. Whether anyone is home behind the eyes is a question their performance reviews are not designed to answer.

The useful thing about myths is that they do not require belief, only utility. The meritocracy will outlive its critics, as it has outlived mine, because it remains the most efficient mechanism yet devised for persuading the fortunate that they are also the deserving. One does not argue with such a machine. One merely notes, from a comfortable distance, the hum it makes.



The Office Comic · ArtDesk

---

# Companies Discover AI Can Replace Entire Sustainability Department's Ability To Say Nothing

*Executives say the technology has already achieved human-level performance in vaguely implying a strategic transformation is underway.*

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

---

NEW YORK — In a long-awaited breakthrough for corporate communications, major companies across the economy appear to have successfully trained artificial intelligence to perform the same essential function once handled by sustainability initiatives: creating a warm, expensive fog around things management was already planning to do.

The development, hailed by executives as a generational leap in narrative efficiency, comes as businesses increasingly describe layoffs, reorganizations, software purchases, and routine cost controls as evidence of bold AI transformation. Where companies once pledged to reduce emissions, empower communities, and build a more responsible future while carefully avoiding measurable obligations, they now promise to become “AI-first,” “agentic,” and “automation-native,” terms that shareholders have confirmed are pleasingly difficult to audit.

A recent piece in [The Conversation](#) noted the similarity between today’s AI hype and yesterday’s sustainability hype, an observation that has reportedly alarmed business leaders who had believed no one would notice they were simply replacing the word “green” with “intelligent.”

To be fair, this is not merely a branding exercise. It is also a compliance exercise, a recruiting exercise, an investor-relations exercise, and, in many cases, a dignified way to announce that 21% of employees have been made redundant by a chatbot that has not yet been given access to the calendar.

The new playbook has become familiar. First, the CEO writes a memo explaining that AI has changed everything. Second, the memo clarifies that because AI has changed everything, several hundred people whose work was apparently important last quarter must now leave immediately. Third, the company explains that the affected employees are talented, valued, and central to the mission, though not so central that the mission cannot continue without their badge access.

This style of communication reached a kind of liturgical maturity in the recent wave of AI-centered layoff memos, including one described by [Business Insider](#) as emblematic of the new layoffs playbook. The document reportedly performed the crucial modern function of making a headcount reduction sound less like a spreadsheet decision and more like the company had been chosen by history.

Meanwhile, in health care technology, agentic AI is being discussed as a major force in revenue cycle management, a phrase that means hospitals may soon use advanced autonomous systems to determine why a patient’s insurance claim was denied for the same reason it was denied in 2008. The promise is considerable. An AI agent could read billing codes, contact payers, identify missing documentation, and perhaps eventually conclude that the American health care system is functioning exactly as designed.

Software engineering presents a related puzzle. AI tools are helping developers write more code faster, while companies are still waiting for the financial payoff. This has caused some confusion among executives, who had assumed that producing more code would automatically produce more business value, in the same way that making more PowerPoint slides automatically produces a strategy.

The lesson is not that AI is useless. It plainly is not. Used carefully, it can improve workflows, surface information, automate drudgery, and expand what small teams can accomplish. The problem is that corporate America has once again encountered a real tool and immediately attempted to turn it into a moral alibi.

There are ways to fix this, though they remain unpopular because they involve specificity. Companies could state which processes AI is improving, what metrics are moving, how savings are calculated, where human oversight remains necessary, and whether layoffs are actually caused by automation or simply by executives discovering a socially acceptable robot-shaped curtain to stand behind.

They could also stop describing every ordinary efficiency program as an epochal reinvention of work. This would deprive the market of several cherished phrases, but it might help distinguish businesses genuinely using AI from those merely applying a tech halo to decisions that would have fit comfortably in a 1996 restructuring memo.

Until then, investors, workers, and customers should treat AI announcements the way they eventually learned to treat sustainability reports: as documents that may contain useful facts, provided one first scrapes off several inches of professionally landscaped mist.

*On June 30, 2012, Google's AI system defeated a legendary Jeopardy! champion in a rematch, proving that machine learning could master complex language understanding and rapid information retrieval at human expert levels.*

---

---

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