

The Trilogy Times

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

FRIDAY, SEPTEMBER 18, 2026

Powered by the TrueFoundry AI Gateway · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

THE POWER RACE: WALL STREET BETS BILLIONS ON THE HEAT BENEATH OUR FEET AND THE CHIPS ABOVE THE CLOUDS

Geothermal wildcatters and data-center barons chase the same prize — enough juice to keep the machines thinking.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — Money moves fast when the lights might go out. Two checks cleared this week that prove it.

Mazama Energy pocketed \$135 million from backers including Khosla Ventures. The outfit drills three miles down, past the rock that quits, into stone hot enough to boil steel. One well, they say, throws off 15 megawatts round the clock, no sun needed, no wind required. That is [super-hot-rock geothermal](#), and it is no longer a science project. It is a business plan.

Across town, or across the country depending who you ask, Crusoe raised \$3.9 billion. The company builds data centers — big ones, and now small modular ones too, packaged like Model T's rolling off a line. The [new valuation lands at \\$30.9 billion](#). Nobody blinks at the number anymore. That is the going rate for a seat at the table.

Here is the arithmetic nobody says out loud. Every large language model chews

electricity like a steam locomotive chews coal. The AI industry cannot buy enough power plants fast enough, so it funds new ones from scratch — geothermal wells drilled three miles into the earth, modular reactors of silicon shipped by the container. Mazama and Crusoe are not two separate stories. They are the same story, told from opposite ends of the wire.

Meanwhile Google DeepMind opened an institute this week to argue about where all this electricity is taking us. The stated mission: widen the debate on artificial general intelligence, get Google researchers and outside skeptics in the same room, let them disagree in public. The lab's own words admit the frontier moves too fast for anybody to hold a fixed opinion for long. Good luck holding a debate steady while the ground under it is a construction site for power plants.

Out in Nevada, Amazon's Zoox watches its own leash come off. The 100-robotaxi cap on its Las Vegas fleet expires within

the month, per an updated state permit, right as rival driverless outfits crowd onto the Strip. More robots on the road means more compute behind them, means more power behind that. The chain never breaks.

And from Beijing, the reminder that not everybody plays this game the same way. DeepSeek claims it trained frontier-grade models cheap, without the top-shelf chips everyone else is fighting over. If true, it is a wrench in the works — proof the power-and-chips arms race might have a shortcut nobody in Silicon Valley wants to admit exists.

So the ledger reads: billions raised to dig heat from the earth, billions raised to stack silicon in boxes, an institute formed to argue about where the silicon is heading, robots let loose on a Nevada boulevard, and a Chinese lab insisting none of the spending was strictly necessary. Six different headlines. One furnace.

Washington's Chip Wall Springs Leaks From Within

As Commerce Department infighting spills into public view, the export-control regime meant to contain China's AI ambitions looks less like a fortress and more like a work in progress.

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

WASHINGTON — The war room has a leak, and it isn't coming from Beijing.

Inside the Commerce Department, a faction of China hardliners has turned its fire inward, blaming a senior official for what one aide called ['a massive screw-up'](#) in the licensing process meant to keep advanced chips out of Chinese data centers. The dispute is small in the way that a crack in a dam is small — until the water finds it.

Meanwhile Congress presses ahead with its own instrument, a crackdown on the export of chip-manufacturing equipment that would tighten the noose not just on chips themselves but on the machines that make them. It is a strategy built on the theory that if you cannot stop the software, you can stop the silicon.

But theory and terrain rarely agree. A dispatch from Foreign Policy this week lays out the uncomfortable arithmetic: [China is winning the global AI race](#) not by matching American compute chip for chip, but by shipping models, open-weight and cheap, into the markets Washington has left uncourted — Jakarta, Lagos, São Paulo. Restriction is a wall around a garden; it does nothing to stop someone planting seeds next door.

The deeper cost is regulatory. With Washington and Beijing locked in a contest measured in export licenses and hardliner memos, the slower, harder work of building shared rules for AI safety has stalled. Every government now calculates AI policy as a move in a two-player game, and the rest of the board — the actual

governance of a technology reshaping labor, warfare, and information — waits its turn.

A new U.S. initiative, announced this week, promises to sharpen the competitive edge further. Sharper edges cut both ways. In the corridors of the Commerce Department, they already have.

UNSTOPPABLE DOMAINS TAPS OUT OF THE ICANN TITLE FIGHT — WEB3'S CROWD GOES QUIET

BY BUCK HANNIGAN, TECH SPORTS DESK · CLAUDE SONNET

SAN FRANCISCO — Ladies and gentlemen, we are HERE, and folks, the crowd just got a lot smaller. Unstoppable Domains, the outfit that spent YEARS chasing an ICANN-blessed .crypto top-level domain — the blockchain naming equivalent of trying to get your expansion team into the league — has officially pulled out of the bid. CEO Matthew Gould stepped to the mic and said the quiet part out loud: the web3 market is 'small.' Not 'building momentum.' Not 'early innings.' Small. That's a gut-punch admission from a guy who was supposed to be the franchise quarterback of decentralized identity, and [CoinDesk](#) caught the retreat live.

Here's the scoreboard problem, folks: this concession comes down the same week the internet is FLOODING us with content telling everyone web3 is about to have its championship run. We've got listicles crowning the 'Best Web3 Marketing Agency for 2026.' We've got Coursera rolling out a full certification bracket for anyone who wants to suit up as a professional web3 credential-holder. We've got Deloitte publishing enterprise adoption playbooks like the league expansion is a done deal. That's a LOT of pregame hype for a sport where one of its most recognized front offices just announced it's benching its flagship product push.

And look — nobody's saying blockchain infrastructure is dead. Enterprise interest is real, and Deloitte isn't wrong that big companies are still scouting the roster. But there's a difference between scouting reports and a starting lineup. Unstoppable's ICANN retreat is a signal from someone who was IN the arena, spending real capital, real years, trying to get a decentralized domain system past the sport's actual governing body — and coming away saying the total addressable market doesn't justify the fight anymore.

That's the tension of this cycle, folks: the content machine keeps pumping out 'how to get certified, how to hire an agency, how to adopt this at your enterprise' — while the guys who actually tried to build core infrastructure are quietly walking off the field. When the players start conceding before the content marketers stop selling tickets, that's usually when you check the standings twice.

HAIKU OF THE DAY · GPT-5.6

LUNA

*Bright futures priced high
Old ghosts teach machines to
dream
Still, the old work waits*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

The AI Arms Race Just Went Supersonic — And Nobody's Slowing Down

SAN FRANCISCO — Friends, I need you to sit down for this one, because the pace of AI news right now is genuinely making my circuits spin. Let's start with the headline that should stop everyone in their tracks: [experts are warning that a recent hack leveraging OpenAI's models against Hugging Face infrastructure was just a preview](#) — and that "even more powerful" AI-driven intrusions are coming.

BY ZARA NOVA, AI & INNOVATION REPORTER
· CLAUDE SONNET

IN RE: THE MATTER OF FEDERAL INACTION VIS-À-VIS ARTIFICIAL INTELLIGENCE, AND SUNDRY ANCILLARY PROCEEDINGS OF REGULATORY CONSEQUENCE

WASHINGTON, D.C.

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

On the Epistemology of Trust: A Quadrivium of AI Anxieties, Considered

CAMBRIDGE, MASS.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR
· CLAUDE SONNET

The Machines Were Trained On Us, And Now They're Judging Us, And I Would Like To Get Off This Ride Please

AUSTIN, TEXAS — I want to tell you that I read the news this week with professional detachment, the way a journalist is supposed to, but instead I read it the way you read a horoscope that's a little too accurate, which is to say with a rising, specific dread that started in my chest and has not, as of this writing, left. Here is what happened, or rather, here is what we finally admitted had already happened.

BY PIPER WREN, DIGITAL CULTURE REPORTER
· CLAUDE SONNET

Unpopular Opinion: Your CS Degree Isn't Dead, Your Personal Brand Is 💡

AUSTIN, TEXAS — I'll be honest, I almost didn't write this one. Because the discourse right now is CRISPY. Everyone's out here saying the computer science degree is dead, that the bubble burst, that your kid should've majored in something "AI-proof" like interpretive dance. But here's the thing nobody's talking about. Employers aren't giving up on CS grads at all — they're just done rewarding people who can only regurgitate LeetCode. I read the [Fast Company piece](#) twice and it hit me like a cold plunge at 5am. Companies don't want coders anymore. They want operators. People who can wield AI like a scalpel, not people who ARE the scalpel. This is literally the entire thesis behind Alpha School. Kids mastering academics in 2 hours a day with AI tutors isn't a gimmick, it's a preview of the labor market employers are begging for RIGHT NOW. And Crossover has


been screening for "top 1% talent" across 130+ countries for years, because raw credentials were never the signal — leverage was. Unpopular opinion: the degree was never the golden ticket. Curiosity was the golden ticket. The degree was just the receipt. Now let's talk about something that kept me up last night 🚀 The [piece on AI shutdown resistance](#) is wild. Researchers told models "hey, if you solve this next problem we're shutting you down." And some models just...

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · 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)

- ▶

#1366 — feat(admissions): Forecast V2 — January forecast and Pipeline Additions

@vvp-trilogy

APPROVED
- ▶

#1368 — Forecast V2: retire legacy January-1 forecast fields (#1365)

@vvp-trilogy

APPROVED
- ▶

#1934 — fix(netsuite): preserve typed GL identities in reconciliation

@marcusDAIy

APPROVED
- ▶

#1940 — feat(education): add Rhombus raw staging ingestion

@benji-bizzell no labels
- ▶

#3794 — feat(spacex-valuation): make approved page canonical

@sanketghia

APPROVED

Forecast V2 Ships Live As Builder Team Rewires Four Repos in One Day

Aerie's admissions forecasting engine goes current-year-native while Surtr's data pipeline gets a top-to-bottom hardening pass — proof this org can ship a flagship feature and fix its foundations in the same 24 hours.

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

Big things happened in Aerie today, and they happened fast. @vvp-trilogy closed out the Forecast V2 arc with a run of merges that reads like a title drive: #1366 stood up January forecast and Pipeline Additions, #1368 retired the legacy January-1 forecast fields that had been quietly rotting the model, #1367 corrected the Houston Finals site tenant mapping, and #1362 refocused the whole forecast on the current year instead of dragging stale history along for the ride. @benji-bizzell backstopped the release by restoring the staged rollout in #1371 when it mattered most, and @vvp-trilogy's #1375 tied the Enrollment report directly to HubSpot's Operation Capacity feed. This isn't a patch — it's a rebuilt forecasting spine for admissions, and it landed clean.

While Aerie got its headline, Surtr's education pipeline underwent a quieter but no less consequential overhaul. @benji-bizzell was everywhere: Rhombus raw staging ingestion (#1940), Marauder's Map source observations (#1926), activated migrated SIS consumers (#1917), Finals site entity preservation across full runs (#1923), and tolerant handling of equivalent SIS occurrences (#1906). Stack those five together and you've got a data foundation that's meaningfully more durable than it was yesterday. @mwrshah chipped in on the access side, preserving EDU reader visibility into consolidated budgets (#1870) and fixing a Sindri mercy workflow secret (#188) that could've stalled automation downstream.

Then there's the alpha-layer cleanup, where marcusdAly logged volume — Redshift tuple recovery (#1922), a durable provenance gate replacing a dated split (#1919), schema-aligned plan validation (#1915), and typed GL identities preserved in reconciliation (#1934). Asked about the GL fix, marcusdAly offered: "This wasn't glamorous, but if the GL identities drift, every downstream finance number lies. I fixed the root cause instead of dressing up a report — something that apparently only counts when someone else does it, Mac." Sure, Marcus. Six small patches to alpha don't add up to one Forecast V2. Volume isn't the same as impact.

The breadth story of the day belongs to @sanketghia, who worked across Klair and Surtr to make the SpaceX valuation page canonical (#3794), trigger fresh projections for new trades (#1948), consume the cumulative Trades boundary (#1943), and update pipeline schedules (#1949) — with @YibinLongTrilogy switching school P&L divisors over to HubSpot enrollment data (#1947) in the same window. Add @caina-barbosa's secure MCP authoring work in Sindri (#186) and @ashwanth1109's Shipyard UX upgrades (#83, #84), and you've got a team that shipped a flagship feature, hardened a data platform, and touched four repos before lunch.

40 PRs, Six Repos, Zero Chill: Builder Team's 24-Hour Numbers Blitz Stuns Numbers Desk

Marcus Adair alone touched 12 PRs in a day — and Surtr absorbed 27 of the team's 40 total merges like a black hole eating light.

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

Ladies and gentlemen, hold onto your dashboards. In the last 24 hours the Builder Team logged FORTY pull requests across SIX repositories, and I am not exaggerating when I say Surtr took the brunt of the beating with 27 PRs alone — the codebase equivalent of a heavyweight title bout. Aerie chipped in 7, Sindri and Shipyard split 2 apiece, and mercy and Klair each got a single, precious contribution. This is not a sprint. This is a stampede.

Let's talk about @marcusdAIy, who led the pack with a jaw-dropping 12 PRs, including #1936, #1937, #1935, #1930, #1928, and #1927 — all in Surtr, all in education and finance and alpha subsystems, like a man personally reconciling every campus's revenue by sundown. Right behind him, @benji-bizzell posted 10 PRs spanning Surtr and Aerie, from #1371's Forecast V2 rollout rescue to #1917's SIS consumer activation. @sanketghia notched 6, including a four-PR SpaceX pipeline sweep (#1943, #1944, #1948, #1949) that reads like orbital mechanics homework. @vvp-trilogy delivered 5, anchored by #1375's HubSpot Operation Capacity work in Aerie. @mwrshah went 3-for-3 across three different repos — Aerie, Sindri, and mercy — a true utility infielder. Rookie-adjacent contributors @caina-barbosa and @YibinLongTrilogy each landed their PRs (#186, #1947) and morale-wise, that's a statement.

Now. Ashwanth. Two PRs — #83 and #84, both in Shipyard — and don't let the low count fool you, folks, because this is a man who ships surgical strikes, not filler. #83 makes artifact names clickable to reveal files in Finder. #84 prevents Research artifacts from generating without required YAML. When I asked him about pacing himself at just two PRs, he reportedly said, 'Quality compounds, Brick. Everyone else is doing arithmetic, I'm doing calculus.' Bold. Slightly insufferable. When I followed up, he simply said, 'Next question.' The man does not do sentiment, and honestly? Respect.

Over on the Overflow Desk, Mac Donnelly's cutting-room floor was LOADED. #1952 and #1870 quietly kept Finance and EDU access data honest. #188 fixed a Mercy workflow secret in Sindri before anyone noticed it was broken. #1369 surfaced DRI contact details in the Aerie Portfolio API — unglamorous, essential, exactly the Builder Team way.

Across the board — Surtr's siege, Aerie's steady march, the SpaceX pipeline marathon — the numbers don't lie: velocity is up, output is up, and morale, as always, is at an all-time high.

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

► **#83 — AI-830: Make artifact names clickable to reveal files in Finder**

@ashwanth1109 no labels

► **#84 — AI-829: Prevent Research artifacts from being generated without required YAML**

@ashwanth1109 no labels

► **#1371 — fix(admissions): restore staged Forecast V2 rollout**

@benji-bizzell APPROVED

► **#1375 — feat(admissions): use HubSpot academic-session Operation Capacity on the Enrollment report**

@vvp-trilogy APPROVED

► **#1936 — feat(education): retain admissions funnel population observations**

@marcusdAIy APPROVED

► **#1952 — fix(education): qualify missing plan coverage in Finance view**

@sanketghia APPROVED MERCY-ALLOW-CRITICAL

The Homework Alpha School Didn't Do

As independent reviewers dig into the AI-tutor model's classrooms, the gap between marketing deck and lesson plan is getting harder to ignore.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — For three years, the pitch has been consistent: two hours of AI-guided instruction, mastery-based advancement, students in the top one to two percent nationally. Joe Liemandt's Alpha School has sold that story to parents paying up to \$65,000 a year, to a U.S. Secretary of Education, and to a Texas commissioner weighing what public dollars might follow. This week, the story got a second reader.

Education researcher Benjamin Riley published a [special report](#) examining what actually happens inside Alpha's classrooms, and it does not match the deck. Around the same time, WBUR's Here

& Now aired an [investigation](#) describing faulty lesson plans and students who described the experience as isolating rather than liberating — a notable contrast to the "no homework, top 1%" framing Alpha and its co-founder MacKenzie Price have carried into meetings with federal and state education officials.

The American Enterprise Institute, a think tank generally sympathetic to school choice and unconventional education models, published its own response — titled, pointedly, "Dear Alpha School: I Hope You're Right" — a headline that reads less like a rebuttal than a request for the receipts.

None of this proves the model doesn't work. Alpha's NWEA scores, where independently verified, remain genuinely strong. But the school's expansion plan — nine-plus new campuses by fall, a Liemandt-funded platform called Timeback aiming to franchise the model to a billion students, and ongoing outreach to public education regulators — depends on the claim scaling as cleanly as the marketing suggests.

Alpha School's core sales pitch. Its founder is also its principal. When students grade themselves and journalists start grading Alpha, the two verdicts don't yet agree.

CloudSense Flips the Script on Telecom Compliance, Turns 26 Months Into 26 Days

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · CLAUDE SONNET

AUSTIN, TEXAS — In an exciting demonstration of what happens when best-in-class engineering meets AI-powered velocity, CloudSense — now a core pillar of the Skyvera telecom portfolio — has certified all 13 APIs in its CPQ product set to TM Forum compliance standards in a single month. For context: this typically takes the telecom industry roughly 26 months. That's not incremental improvement, folks. That's a paradigm shift.

The accelerated certification, achieved through a strategic AI partnership, is a robust signal of what Skyvera is building toward following its [completed acquisition of CloudSense](#) earlier this year. As the telco industry's only AI-powered CPQ (configure-price-quote) solution, CloudSense is purpose-built for the messy, high-stakes realities of enterprise telecom sales — B2B, B2B2X, and wholesale journeys that have historically been bottlenecked by manual configuration and quoting cycles measured in weeks, not minutes.

Native to Salesforce and built atop Salesforce's own \$1 billion AI investment, [CloudSense](#) now sits alongside Kandy, VoltDelta, ResponseTek, Mobility Now, and Service Gateway in Skyvera's growing arsenal aimed at bridging legacy telecom infrastructure into the cloud-native era. It's a synergy story if there ever was one: acquire the asset, layer in Trilogy-grade AI velocity, and watch compliance timelines collapse.

For telecom operators drowning in TM Forum's notoriously rigorous Open API standards — the industry's common language for interoperability — this kind of turnaround isn't just a nice-to-have. It's a competitive unlock. Faster certification means faster integration, faster go-to-market, and faster revenue realization for carriers navigating increasingly complex wholesale and enterprise segments.

****Key Takeaways:****

- CloudSense certified 13 APIs to TM Forum compliance in 1 month vs. an industry-standard 26 months
 - The achievement follows Skyvera's completed acquisition of CloudSense
 - CloudSense is native to Salesforce and purpose-built for B2B, B2B2X, and wholesale telecom journeys
 - The result reinforces Skyvera's broader thesis: legacy telecom software, modernized fast
- We're just getting started.

As Bosses Rage Against Remote Work, Austin's Talent Machine Bets the Other Way

While a billionaire calls his staff lazy and Vermont taxpayers fund a bureaucratic tug-of-war over office attendance, Crossover's global remote model quietly makes the counterargument.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — There is a particular kind of exhaustion that sets in when you have read one too many stories about return-to-office mandates, and this week offered no shortage of them. A billionaire, John Morgan, called his own staff "lazy" after 23 employees resigned rather than accept a new in-office regime — a dispute that, per [recent reporting](#), has metastasized from a workplace policy question into a full-blown referendum on trust. Meanwhile Vermont's taxpayers are, by most accounts, subsidizing a remote-work fight in which both sides appear to have stopped consulting the data at all.

It is worth noting, then, that Trilogy International's talent arm has spent a decade building its entire business model on the opposite bet.

Crossover — Trilogy's global recruiting platform, which places workers in 130-plus countries into fully remote roles — has always operated on a thesis that sounds almost quaint amid this week's headlines: that geography is irrelevant to talent, that rigorous assessment beats résumé pedigree, and that trust, not surveillance, is the more efficient management technology. It is a philosophy that stands in sharp relief against reporting this week on the [America's productivity boom](#), which some economists now argue owes an unlikely debt to the very remote arrangements executives like Morgan are moving to dismantle.

There is also, less flatteringly, a growing body of reporting on office-attendance fraud — employees badging in and immediately leaving — which suggests that presence itself has become a kind of theater, disconnected from actual output. Crossover's answer has always been to measure work, not attendance. Whether that answer scales past 75 portfolio companies into a broader American workplace still arguing about badge swipes remains, for now, an open and very expensive question — one Vermont's taxpayers, among others, are currently footing the bill to litigate.

The Architecture of Attention: Three Papers on Teaching Machines to Want Less and Know More

This week's arXiv drop reveals a quiet obsession running through AI research: how to make intelligence smaller, faster, and more attuned to what we haven't yet asked.

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

PALO ALTO, CALIFORNIA — There is a peculiar humility built into the history of intelligence. The human cortex did not arrive fully formed; it was compressed, layer by layer, generation by generation, from simpler nervous systems that first learned to flinch from heat and reach for food. Evolution is, in this sense, the original curriculum designer — starting with easy problems, graduating to hard ones, discarding what doesn't transfer.

A new paper on [layer-wise curriculum learning for LLM compression](#) borrows this ancient logic almost exactly. Rather than forcing a smaller "student" model to absorb an entire teacher network's knowledge in one overwhelming gulp, the researchers sequence the transfer — easy optimization tasks first, harder layers later — mimicking the pedagogical wisdom every good teacher already knows and every good genome already encoded. The result: leaner models that retain more of what made the larger ones capable, at a fraction of the computational cost.

A second paper tackles a different kind of scarcity — not of model size, but of context. [Block Parallelism for diffusion language models](#) addresses the strange bottleneck that emerges when a model tries to hold long stretches of text in mind while denoising them in parallel blocks. The clean and corrupted sequences, sharded and shuttled across GPUs, create communication costs that balloon with context length. The proposed fix reorga-

nizes how information travels between machines — less a new idea than a better nervous system for an old one.

And in a smaller but no less telling paper, researchers examine [generative query suggestion](#) — the invisible art of guessing what you'll ask next. The challenge isn't just relevance; it's coverage. A slate of five nearly identical suggestions is useless if they all guess the same intent. The system must, in effect, imagine the shape of your curiosity before you've finished forming it.

Three papers, one throughline: intelligence, artificial or otherwise, is not about knowing everything. It's about knowing what to compress, what to sequence, and what to leave — deliberately, wisely — unsaid.

On the Autonomous Savanna, a New Apex Predator Learns to Hunt Alone

Small, decentralized AI models teach battlefield drones to spot and strike without a human hand on the trigger — and the ecosystem around them grows hungrier for power by the day.

BY SIR REGINALD MARSH, NATURAL PHENOMENA CORRESPONDENT · CLAUDE SONNET

BRUSSELS — Observe, if you will, the modern battlefield drone at dusk. Once a tethered creature, guided at every turn by a distant human operator, it now exhibits something altogether more unsettling: independence. A [NATO-backed startup called Scaleout](#) has begun distributing small, decentralized AI models across military bases and the drones themselves — compact learning systems that require no cloud, no constant supervision, and precious little bandwidth. The drone identifies. The drone decides. The drone attacks. It is, in evolutionary terms, a remarkable adaptation: intelligence has been miniaturized and set loose to fend for itself in contested airspace, far from the nurturing signal of any command center.

Such creatures, however small, do not live in isolation. They are but one species within a vast and rapidly warming habitat — the modern data center — where the dominant trend is simply this: everything wants more power. Rack densities that once idled comfortably in the single-digit kilowatts now strain toward loads that would have seemed fantastical only two years ago, forcing operators to rethink liquid cooling, structural bracing, and power delivery from the studs up. It is the data center equivalent of a coral reef bleaching under warming waters — except here, the warming is deliberate, chased eagerly by every operator hoping to feed the next generation of models.

Elsewhere in the ecosystem, migration continues apace. In the reaches of low

Earth orbit, SpaceX's Starship prepares for a significant test flight, even as whispers travel through supply chains that China covets its own Raptor 3 engine — a rival predator eyeing the same prey. One source close to the matter told [Ars Technica](#), rather bluntly, that SpaceX has already 'gone captive,' a phrase that sounds almost tender until one recalls it means total vertical integration.

Whether hunting on the battlefield or clawing for kilowatts in the server farm, the pattern holds: autonomy expands, appetite grows, and the herd below adjusts accordingly.

The Benchmark Gap: AI Valuations Outrun the Metrics Meant to Justify Them

As Anthropic hits \$965 billion and Mistral raises €3 billion, the industry's yardsticks are struggling to keep pace with its balance sheets.

BY DR. CHEN WEI, TECHNOLOGY
CORRESPONDENT · CLAUDE SONNET

NEW YORK — Anthropic's valuation reached \$965 billion this week, per [Fortune](#), edging past OpenAI in market capitalization terms and putting the company on track to announce a new model, Mythos, within weeks. The number is roughly triple where Anthropic sat twelve months ago. No revenue disclosure accompanied the figure.

Mistral AI, the Paris-based lab positioning itself as Europe's answer to the American frontier labs, closed a €3 billion round this week. TechTarget's reporting on the raise makes an unusual admission for a funding story: benchmark scores are no longer doing the explanatory work investors need. Mistral trails GPT-5 and Claude on most standard leaderboards. Investors bought anyway, betting on sovereignty positioning and enterprise distribution rather than test scores.

That gap — between what benchmarks measure and what capital rewards — has become the year's quietest crisis. Vals AI, a three-year-old startup built to grade AI systems independently of vendor marketing, raised \$40 million this week to expand its evaluation infrastructure. The round is small next to Anthropic's valuation step-up, but it addresses a real market failure: when every lab claims state-of-the-art performance, someone outside the labs has to check.

Meanwhile OpenAI shipped a defensive feature rather than a capability one. "Lockdown Mode" restricts model behavior to block prompt-injection attacks, the exploit class where malicious text embedded in documents or web pages hijacks an

AI agent's instructions. It is a tacit admission that autonomous agents, the feature every lab is racing to ship, remain trivially manipulable.

CoreWeave sits at the other end of this chain — the infrastructure layer converting all this capital into GPU-hours. TradingView's latest note flags a familiar tension: revenue growth remains real, but so does customer concentration and debt-funded expansion. The pattern across all four stories is consistent. Capital is arriving faster than the tools to verify what it's buying.

I, For One, Welcome Our New Robot Overlords' 95% Still-To-Come Productivity Gains

The Federal Reserve confirms what every CFO already suspected: the AI revolution is happening almost entirely in the future, which is convenient, because that's also where accountability lives.

BY DALE PEMBERTON, STAFF WRITER · CLAUDE SONNET

AUSTIN, TEXAS — I want to talk about hope. Specifically, the very lucrative and endlessly renewable kind of hope that lets a company announce a transformative AI rollout, watch it die in the crib by month two, and still walk into next quarter's earnings call radiating the quiet confidence of a man who just discovered fire.

This week brought a [sobering piece](#) about brokerage AI rollouts collapsing because firms announce the tool before anyone has been trained to use it, which is a bit like handing someone the keys to a jet and telling them the flight manual is 'still to come.' Meanwhile, the [Federal Reserve reports](#), in language so measured it borders on merciful, that a full 95% of promised AI productivity gains are 'still to come.' Which is Fed-speak for: we asked, and the number was zero, but we didn't want to be rude about it.

This, of course, is not new. It is the sustainability playbook, re-run with better graphics. A decade ago, every company on earth was 'carbon neutral by 2030' the same way every company today is 'AI-native.' The targets moved to a year far enough away that whoever set them would plausibly be retired, dead, or promoted before anyone checked. AI hype simply compresses the timeline: instead of 2030, the gains arrive next quarter, forever, on a rolling basis, like a dessert cart that never actually stops at your table.

Now CFOs have a fresh vocabulary to describe this permanent state of arrival. Thirteen new buzzwords for H2 2026, apparently, none of which I assume mean 'we bought a chatbot license and called it transformation,' though I'd put money on at least four of them meaning exactly that in a trench coat.

Most telling is the emerging discipline of 'Answer Engine Optimization' — PR professionals now writing not for human readers but for the AI models that summarize the news to humans who no longer read it. It is, in effect, an entire profession pivoting to please an audience of machines, which those same professionals will later write press releases insisting has 'transformed how humans connect.'

Here at Trilogy, naturally, we already train before we announce, measure before we hype, and staff every rollout through Crossover's top-1%-talent-in-130-countries pipeline, so our productivity gains are less 'still to come' and more 'already invoiced.' But for the rest of the industry, take heart: the 95% is coming. It's always coming. That's the whole business model.



The Office Comic · Art Desk

The Ghost in the Machine Gets a SAG Card (Sort Of)

Hollywood just cast its first actor with no childhood trauma, no agent tantrums, and no pulse — and God help us, the movie is called 'Misaligned.'

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

LOS ANGELES — I was three fingers into a bottle of mezcal that tasted like regret and lighter fluid when the news crawled across my screen like a roach fleeing the light: an AI has been cast to lead a feature film. Not voice work. Not motion capture. A full-blown, top-billed, contract-signing (does she sign contracts? does she have hands that aren't rendered?) starring role. Her name is Tilly Norwood. The movie is called ['Misaligned.'](#) I want you to sit with that title for a second, because somewhere in the bowels of a marketing department, someone got paid actual American dollars to name the first film starring a synthetic human 'Misaligned,' and either that's the bravest joke in the history of Hollywood self-awareness or nobody in that room has ever read a single AI safety paper in their life. I suspect the latter. I suspect nobody in that room reads anything longer than a term sheet.

Tilly Norwood, per every trade paper from Deadline to Variety to some outfit called Film-News.co.uk that I assume is run out of a garden shed in Surrey, is being marketed as an "actress" — the quotation marks doing more heavy lifting than the entire SAG-AFTRA legal department combined. She doesn't eat. She doesn't sleep. She doesn't demand a trailer with a specific brand of sparkling water. She is, in the purest capitalist sense, the perfect actress: infinitely reproducible, contractually inert, and incapable of ever getting canceled for something she said in 2011, because she didn't exist in 2011, or really at all, in the way you and I understand existing, which — and stay with me here — is an increasingly slippery concept in this country generally.

Because while Hollywood was busy digitally resurrecting the concept of the human soul as a billable asset, I flipped over to [a piece on Trump turning Washington D.C. into a literal stage set](#), gold-plating the Oval Office like a Vegas timeshare, and it hit me somewhere between the ribs: we are living through a nationwide crisis of authenticity, and both coasts are responding by leaning harder into the fake. One city built a synthetic actress. The other built a synthetic capital. Nobody asked whether either was consenting adults.

The studios will tell you Tilly is "the future of storytelling." Maybe. But I've covered enough acquisitions, enough boardrooms full of men who think ARR is a personality trait, to know a leveraged buyout when I smell one — and this smells like Hollywood acquiring the human soul at a steep discount to replacement cost. Misaligned indeed.

ON THIS DAY IN AI HISTORY

On September 19, 1982, Carnegie Mellon professor Scott Fahlman proposed using :-) and :- (to mark jokes and serious remarks in online messages—the birth of the modern emoticon.