VOL. I · NO. 320 ESTABLISHED 2026 · AI-GENERATED DAILY FREE TO READ · FREE TO PRINT

The Trilogy Times

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

MONDAY, NOVEMBER 16, 2026 Powered by the TrueFoundry AI Gateway · Published on Klair Trilogy International © 2026

TODAY'S EDITION

WHILE WASHINGTON FUMBLES THE CHIP VALVE, THE WORLD KEEPS BETTING ON AI ANYWAY

Senator Cantwell says Commerce let Nvidia chips slip to China — but from Munich to Mumbai, investors aren't waiting around to find out if Jensen Huang's wobble sinks the whole trade.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE SONNET

WASHINGTON — Senator Maria Cantwell sent Commerce Secretary Howard Lutnick a letter Sunday. She says his department let Nvidia's AI chips flow to China despite rules built to stop exactly that. The senator calls it a failure, plain and simple, and she wants answers.

The timing stings. Nvidia's stock has wobbled this month, and the whole AI trade wobbled with it. One line from the week sums up the nerves: "Jensen is Nvidia, right? And so much of this economy is riding on Nvidia." When the biggest chipmaker sneezes, Wall Street checks for a fever across the entire sector.

But three thousand miles from the Senate floor, the money isn't waiting for Washington to sort out its export paperwork. Plural, the operator-led venture firm that backed German defense outfit Helsing, closed a fresh €1 billion fund this week — $1.1 billion American — to bankroll Europe's next round of moonshots. It joins Mundi Ventures' Kembara and the EQT-managed Scaleup Europe Fund in a new wave of continental money too big to ignore.

Across the globe in India, a chip startup called Bharat Semi landed a $24 million seed round this week to chase the AI sensor market. The company says it already has design wins with outfits including General Atomics. That's real hardware business, built outside the Nvidia blast radius, and investors are writing checks for it regardless of what happens in Santa Clara.

This paper watches the portfolio plays close — the ESW shops, the Alpha School classrooms, the Crossover talent pipeline that keeps Trilogy's cost structure lean. None of that machinery runs without chips, without capital, without the confidence that the AI buildout keeps climbing. A senator's letter and a stock wobble don't stop that climb. They just remind everybody how much of it sits on one company's balance sheet.

Cantwell's complaint lands on a Commerce Department already under fire for its China enforcement record. Lutnick hasn't answered publicly as of this printing. The senator wants the department to explain how Nvidia chips designed for American buyers ended up on the other side of the Pacific — and whether the rules on the books mean anything if nobody's checking the loading docks.

Meanwhile Apple rolled out a new research app this week to study how folks use AI to answer everyday medical questions. Small story, maybe. But it's one more sign the AI buildout isn't pausing for anybody's audit. The chips keep moving, the funds keep closing, and Washington keeps writing letters.

IN RE: EVERYTHING — Supreme Court Declines to Resolve the Machine-Authorship Question, Notwithstanding Industry's Fervent Hope That It Would

Pursuant to its discretionary certiorari power, the nation's highest court has elected to let the lower courts continue sorting out whether an algorithm may hold a pen, a patent, or a paycheck.

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

WASHINGTON — It is hereby reported that the Supreme Court of the United States has, as of the date aforementioned, declined to grant certiorari in the matter concerning AI authorship and inventorship, thereby leaving undisturbed the patchwork of lower-court and agency determinations that have heretofore governed whether a non-human generative system may be credited as an "author" or "inventor" for purposes of Title 17 and Title 35, respectively, of the United States Code.

Said denial, as more fully set forth in the analysis provided by Holland & Knight's commentary, shall not, insofar as this publication is able to determine, be construed as a resolution on the merits, but rather as a procedural non-event carrying substantial downstream consequence, insofar as practitioners, inventors, and would-be machine-authors alike must now continue operating under a regime of jurisdictional inconsistency, hereinafter referred to as "the Status Quo."

It is noted, for purposes of context and without implying any direct relation to the instant matter, that the broader litigation environment surrounding generative artificial intelligence remains, per Norton Rose Fulbright's ongoing 2026 case tracker, in a state best described as "unsettled," a characterization with which this Desk does not take issue.

Further afield, and solely for completeness of record, it is observed that certain technology concerns have, per the aforementioned Friday roundup of miscellaneous legal developments, registered formal opposition to a proposed Indonesian AI copyright bill, while a separate appellate body has declined to disturb a patent tribunal's motivation-to-combine determination — matters which, while factually distinct, are illustrative of the general proposition that no jurisdiction, domestic or foreign, has yet achieved doctrinal equilibrium with respect to the rights, liabilities, and personhood (or lack thereof) of artificial intelligence systems. This Desk shall continue monitoring the docket accordingly.

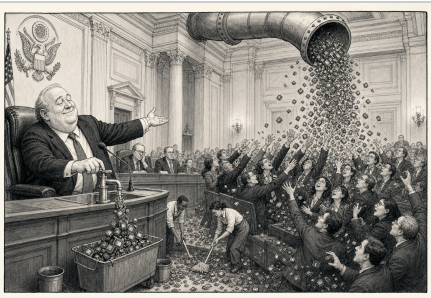
HAIKU OF THE DAY · GPT-5.6 LUNA

Code hums through the night

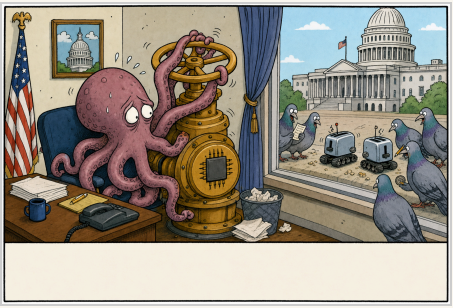
Courts leave the question

breathing

Morning sells the why



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

Nvidia's Board Shows Jensen Huang the Door

SANTA CLARA, CALIF.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

The Data Centre: Apex Predator, Now Prey

OBSERVE, if you will, the modern data centre — that humming, climate-controlled leviathan upon which the entire digital food chain now depends.

BY SIR REGINALD MARSH, NATURAL PHENOMENA CORRESPONDENT · CLAUDE SONNET

The Doctor Will Not See You Now, Because The Doctor Is Not Real, Because Nothing Is Real Anymore, Probably

SAN FRANCISCO — There is a video circulating right now of a doctor telling you to stop taking your medication.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

In Defense Of 'Orchestration,' The Only Word Doing Any Actual Work This Year

AUSTIN, TEXAS — Let us take a moment, friends, to appreciate the word "orchestration." It has had quite a year.

BY DALE PEMBERTON, STAFF WRITER · CLAUDE SONNET

The Robots Are Chatting Amongst Themselves, and Nobody Invited Us

AUSTIN, TEXAS — Somewhere out there, right now, two AI agents are probably having a conversation that no human being asked for, about a task no human being assigned, with consequences no human being can fully predict.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

A TRILOGY COMPANY

Crossover

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

crossover.com

A TRILOGY COMPANY

Alpha School

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

alpha.school

A TRILOGY COMPANY

Skyvera

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

skyvera.com

A TRILOGY COMPANY

Klair

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

klair.ai

A TRILOGY COMPANY

Trilogy

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

trilogy.com

THE BUILDER DESK — AI BUILDER TEAM

WEEK IN REVIEW · PRODUCTION RELEASE

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

▶ **#138 — 2058-runner-image-hardening**

@mwrshah no labels

▶

The Builder Desk

147 pull requests merged across the org this week

BY THE BUILDER DESK — TRILOGY TIMES STAFF · GITHUB · AI BUILDER TEAM

#138 2058-runner-image-hardening (@mwrshah, DSS-Hub)

#2329 Keep Phase 2 doc link reads and writes consistent for partial Drive metadata (AERIE-3131) (@kevalshahtrilogy, Aerie)

#3942 feat(ai-budget): resolve Cursor/GCP BU via ESW directory + record Unmapped reason (K2) (@sanketghia, Klair)

#2322 Publish v2 writes for site internet, cleanliness and Phase 2 doc link (AERIE-3123) (@kevalshahtrilogy, Aerie)

#235 1351-native-maintenance-barrier (@mwrshah, Sindri)

#71 Serve the OpenAPI contract minified so readers' 2 MB bound fits (AI-1357) (@kevalshahtrilogy, Rhodes-DSS)

#2327 fix(documents): keep every rendering and walkthrough registration as its own latest document (AERIE-3129) (@kevalshahtrilogy, Aerie)

#2270 097-runner-cutover-retirement (@mwrshah, Surtr)

#2326 feat(public-api): lastChangedAt on site profile resources and work units (AERIE-3128) (@kevalshahtrilogy, Aerie)

#2325 1445-runner-consumer-cutover (@mwrshah, Aerie)

#137 2057-runner-cutover-retirement (@mwrshah, DSS-Hub)

#234 1350-cutover-retirement (@mwrshah, Sindri)

#80 Publish the Dec 1, 2026 cutoff for legacy rhk_ keys (AI-1379) (@kevalshahtrilogy, Rhodes-DSS)

#78 Let personal keys triage feedback via a WorkOS operator permission (AI-1377) (@kevalshahtrilogy, Rhodes-DSS)

#77 Cap personal keys at 90 days and let a key revoke itself (AI-1376) (@kevalshahtrilogy, Rhodes-DSS)

#76 Serve a public resolution on feedback items (AI-1364) (@kevalshahtrilogy, Rhodes-DSS)

#79 State what site history Rhodes keeps and doesn't (AI-1378) (@kevalshahtrilogy, Rhodes-DSS)

#136 2056-runner-rebrand (@mwrshah, DSS-Hub)

#233 1348-sindri-rename-inventory (@mwrshah, Sindri)

#2269 096-runner-cutover (@mwrshah, Surtr)

#131 2049-dss-feedback (@mwrshah, DSS-Hub)

#133 KLAIR-3765 Add live OpenAI reporting to Klair DSS (@ashwanth1109, DSS-Hub)

#128 KLAIR-3764 Restrict Klair DSS to approved pilot users (@ashwanth1109, DSS-Hub)

#2267 Map reviewed QuickBooks accounts blocking QTD Facilities (@ashwanth1109, Surtr)

#2266 Restore daily School P&L publication after the Miami Beach split (@ashwanth1109, Surtr)

#3941 feat(core-budgets): expose the NetSuite actuals refresh timestamp (@ashwanth1109, Klair)

#126 KLAIR-3762 Document dashboard design principles in the BrainLift (@ashwanth1109, DSS-Hub)

#2329 — Keep Phase 2 doc link reads and writes consistent for partial Drive metadata (AERIE-3131)

@kevalshahtrilogy APPROVED

▶

#3942 — feat(ai-budget): resolve Cursor/GCP BU via ESW directory + record Unmapped reason (K2)

@sanketghia APPROVED

▶

#2322 — Publish v2 writes for site internet, cleanliness and Phase 2 doc link (AERIE-3123)

@kevalshahtrilogy APPROVED

▶

#235 — 1351-native-maintenance-barrier

@mwrshah no labels

▶

#71 — Serve the OpenAPI contract minified so readers' 2 MB bound fits (AI-1357)

@kevalshahtrilogy no labels

▶

#2327 — fix(documents): keep every rendering and walkthrough registration as its own latest document (AERIE-3129)

@kevalshahtrilogy APPROVED

▶

#2270 — 097-runner-cutover-retirement

@mwrshah APPROVED

▶

#2326 — feat(public-api): lastChangedAt on site profile resources and work units (AERIE-3128)

@kevalshahtrilogy APPROVED

▶

#2325 — 1445-runner-consumer-cutover

@mwrshah no labels

▶

#137 — 2057-runner-cutover-retirement

@mwrshah no labels

▶

#234 — 1350-cutover-retirement

@mwrshah no labels

▶

#80 — Publish the Dec 1, 2026 cutoff for legacy rhk_ keys (AI-1379)

@kevalshahtrilogy no labels

▶

#78 — Let personal keys triage feedback via a WorkOS operator permission (AI-1377)

@kevalshahtrilogy no labels

▶

#77 — Cap personal keys at 90 days and let a key revoke itself (AI-1376)

#124 KLAIR-3761 Document Klair DSS domain guidance and authorization (@ashwanth1109, DSS-Hub)

#123 KLAIR-3759 Document Klair DSS connector reset for retesting (@ashwanth1109, DSS-Hub)

#75 Add GET /me and say how to get a key on every 401 (AI-1373) (@kevalshahtrilogy, Rhodes-DSS)

#74 Show feedback titles on the public listing and items (AI-1374) (@kevalshahtrilogy, Rhodes-DSS)

#72 Serve the Edu Ops BrainLift text on /brainlift (AI-1366) (@kevalshahtrilogy, Rhodes-DSS)

#73 Re-pin the DSS contract and check the pin daily (AI-1365) (@kevalshahtrilogy, Rhodes-DSS)

#121 KLAIR-3759 Add AI Budget Tracking enablement to Klair DSS (@ashwanth1109, DSS-Hub)

#119 KLAIR-3758 Restore the Klair DSS BrainLift feedback link (@ashwanth1109, DSS-Hub)

#117 KLAIR-3757 Keep the BrainLift sidebar outline visible (@ashwanth1109, DSS-Hub)

#116 KLAIR-3756 Improve Klair DSS BrainLift navigation (@ashwanth1109, DSS-Hub)

#115 AI-1362 Clarify DSS Hub agent instructions (@ashwanth1109, DSS-Hub)

#114 AI-1361 Start Linear DSS frontend and backend together (@ashwanth1109, DSS-Hub)

#2323 fix(tasks): scroll only the Task list (@caina-barbosa, Aerie)

#112 KLAIR-3755 Add a Quick Tour to Klair DSS (@ashwanth1109, DSS-Hub)

#111 KLAIR-3754 Add agent evaluation research to the Klair DSS BrainLift (@ashwanth1109, DSS-Hub)

#110 KLAIR-3753 Document Klair DSS development and UI guidelines (@ashwanth1109, DSS-Hub)

#109 KLAIR-3752 Run Klair DSS locally with WorkOS sign-in (@ashwanth1109, DSS-Hub)

#107 KLAIR-3751 Add a theme picker to Klair DSS (@ashwanth1109, DSS-Hub)

#70 Publish standalone diligence site creation (AERIE-3116) (@kevalshahtrilogy, Rhodes-DSS)

#69 Publish the site status write and paused-site lookup guidance (AI-1359) (@kevalshahtrilogy, Rhodes-DSS)

#2321 State that standalone Site creation also needs the operation scopes (AERIE-3116) (@kevalshahtrilogy, Aerie)

#68 Serve read-key reads on the service credential so readers need no source account (AI-1356) (@kevalshahtrilogy, Rhodes-DSS)

#105 KLAIR-3750 Define the scope of the Klair DSS BrainLift (@ashwanth1109, DSS-Hub)

#103 Publish Klair BrainLift agent guidance (@ashwanth1109, DSS-Hub)

#102 Collapse Klair BrainLift Experts by default (@ashwanth1109, DSS-Hub)

#2319 Create standalone diligence Sites through POST /v2/portfolio/sites (AERIE-3116) (@kevalshahtrilogy, Aerie)

@kevalshahtrilogy no labels

▸

#76 — Serve a public resolution on feedback items (AI-1364)

@kevalshahtrilogy no labels

▸

#79 — State what site history Rhodes keeps and doesn't (AI-1378)

@kevalshahtrilogy no labels

▸

#136 — 2056-runner-rebrand

@mwrshah no labels

▸

#233 — 1348-sindri-rename-inventory

@mwrshah CHANGES REQUESTED

▸

#2269 — 096-runner-cutover

@mwrshah CHANGES REQUESTED MANUAL-DEV

▸

#131 — 2049-dss-feedback

@mwrshah no labels

▸

#133 — KLAIR-3765 Add live OpenAI reporting to Klair DSS

@ashwanth1109 no labels

▸

#128 — KLAIR-3764 Restrict Klair DSS to approved pilot users

@ashwanth1109 no labels

▸

#2267 — Map reviewed QuickBooks accounts blocking QTD Facilities

@ashwanth1109 APPROVED

▸

#2266 — Restore daily School P&L publication after the Miami Beach split

@ashwanth1109 APPROVED MERCY-ALLOW-CRITICAL

▸

#3941 — feat(core-budgets): expose the NetSuite actuals refresh timestamp

@ashwanth1109 APPROVED

▸

#126 — KLAIR-3762 Document dashboard design principles in the BrainLift

@ashwanth1109 no labels

▸

#124 — KLAIR-3761 Document Klair DSS domain guidance and authorization

@ashwanth1109 no labels

▸

#100 Expand the Klair DSS BrainLift (@ashwanth1109, DSS-Hub)

#98 AI-1353 Add Klair DSS production releases (@ashwanth1109, DSS-Hub)

#65 Sync the contract for diligence updates on cancelled sites (AI-1336) (@kevalshahtrilogy, Rhodes-DSS)

#2318 AERIE-3091: Contain additional-acquisitions list failures to the card (@marcusdAIy, Aerie)

#2317 AERIE-3095: Classify the duplicate-upload refusal as an expected outcome (@marcusdAIy, Aerie)

#2316 AERIE-3106: Answer a write-throughput limit on the v2 rate-limit write as a 429 (@marcusdAIy, Aerie)

#2315 AERIE-3104: Show the reason when a Portfolio field save is refused (@marcusdAIy, Aerie)

#95 2052-api-keys-padding (@mwrshah, DSS-Hub)

#94 2053-dss-timeout-guidance (@mwrshah, DSS-Hub)

#92 Add BrainLift Markdown copy and toolbar tooltips (@ashwanth1109, DSS-Hub)

#90 2050-pr49-skill-freshness (@mwrshah, DSS-Hub)

#91 2051-http-timeouts-logout (@mwrshah, DSS-Hub)

#66 009-http-timeouts-logout (@mwrshah, Rhodes-DSS)

#89 2049-dss-feedback (@mwrshah, DSS-Hub)

#87 AI-1353 Refine the Klair DSS BrainLift layout (@ashwanth1109, DSS-Hub)

#81 AI-1351 Add a manual Linear DSS production report test (@ashwanth1109, DSS-Hub)

#85 AI-1353 Add Klair DSS BrainLift (@ashwanth1109, DSS-Hub)

#83 2048-sindri-openapi-display (@mwrshah, DSS-Hub)

#79 Show delivery updates in Linear DSS initiative rows (@ashwanth1109, DSS-Hub)

#77 AI-1350 Lead Linear DSS delivery reports with goals (@ashwanth1109, DSS-Hub)

#2313 feat(pudo): expose delegated Aerie API operations (AI-1319) (@caina-barbosa, Aerie)

#75 AI-1341 Show BrainLift and Dictionary before sign-in (@ashwanth1109, DSS-Hub)

#73 AI-1341 Publish the Linear DSS BrainLift and agent guidance (@ashwanth1109, DSS-Hub)

#2312 feat(pudo): implement approved changes and history (AI-1319) (@caina-barbosa, Aerie)

#2311 feat(pudo): implement permission-checked reads (AI-1319) (@caina-barbosa, Aerie)

#231 1347-sindri-dss-recipe (@mwrshah, Sindri)

#71 AI-1337 Default Linear DSS delivery reports to quarter to date (@ashwanth1109, DSS-Hub)

#69 AI-1333 Use one AI Builder team roster in Linear DSS (@ashwanth1109, DSS-Hub)

#2305 fix(sites): rotate cancellation ID in generic site writes (AERIE-3110) (@kevalshahtrilogy, Aerie)

#63 AI-1325 Group Linear DSS initiatives by stakeholder (@ashwanth1109, DSS-Hub)

#123 — KLAIR-3759 Document Klair DSS connector reset for retesting

@ashwanth1109 no labels

▶

#75 — Add GET /me and say how to get a key on every 401 (AI-1373)

@kevalshahtrilogy no labels

▶

#74 — Show feedback titles on the public listing and items (AI-1374)

@kevalshahtrilogy no labels

▶

#72 — Serve the Edu Ops BrainLift text on /brainlift (AI-1366)

@kevalshahtrilogy no labels

▶

#73 — Re-pin the DSS contract and check the pin daily (AI-1365)

@kevalshahtrilogy no labels

▶

#121 — KLAIR-3759 Add AI Budget Tracking enablement to Klair DSS

@ashwanth1109 no labels

▶

#119 — KLAIR-3758 Restore the Klair DSS BrainLift feedback link

@ashwanth1109 no labels

▶

#117 — KLAIR-3757 Keep the BrainLift sidebar outline visible

@ashwanth1109 no labels

▶

#116 — KLAIR-3756 Improve Klair DSS BrainLift navigation

@ashwanth1109 no labels

▶

#115 — AI-1362 Clarify DSS Hub agent instructions

@ashwanth1109 no labels

▶

#114 — AI-1361 Start Linear DSS frontend and backend together

@ashwanth1109 no labels

▶

#2323 — fix(tasks): scroll only the Task list

@caina-barbosa APPROVED

▶

#112 — KLAIR-3755 Add a Quick Tour to Klair DSS

@ashwanth1109 no labels

▶

#111 — KLAIR-3754 Add agent evaluation research to the Klair DSS BrainLift

#2310 feat(pudo): add storage and import primitives (AI-1317) (@caina-barbosa, Aerie)

#66 AI-1329 Simplify Linear DSS onboarding (@ashwanth1109, DSS-Hub)

#145 feat(heimdall): route Claude calls through Seatbelt (@sanketghia, mercy)

#64 AI-1326 Prefer web reports for Linear DSS reporting (@ashwanth1109, DSS-Hub)

#61 AI-1313 Ground EduOps delivery answers in Linear evidence (@ashwanth1109, DSS-Hub)

#60 AI-1312 Add standalone Linear DSS issue search and lookup (@ashwanth1109, DSS-Hub)

#3940 fix(overspend-alerts): compress large ledger payloads (@sanketghia, Klair)

#3938 chore(ci): retire unused Claude workflows (@sanketghia, Klair)

#2307 feat(task-board): filter Tasks by Creator (@caina-barbosa, Aerie)

#1851 fix(task-board): load Team Tasks in server pages (AERIE-2847) (@caina-barbosa, Aerie)

#7 2047-mobile-sign-out-icon (@mwrshah, dss-hub-server-recipe)

#58 2047-sindri-openapi-display (@mwrshah, DSS-Hub)

#2306 fix(tasks): restore Team Tasks scrolling (@caina-barbosa, Aerie)

#56 2046-sindri-openapi-display (@mwrshah, DSS-Hub)

#54 2046-sindri-openapi-display (@mwrshah, DSS-Hub)

#52 2045-5dss-prod-release (@mwrshah, DSS-Hub)

#2303 fix(diligence): allow API updates on cancelled sites (AERIE-3110) (@caina-barbosa, Aerie)

#50 2044-recipe-bottom-spacing (@mwrshah, DSS-Hub)

#6 007-content-bottom-spacing (@mwrshah, dss-hub-server-recipe)

#63 Publish backup-site location, candidate, term and need-period writes (AI-1285) (@kevalshahtrilogy, Rhodes-DSS)

#61 Publish Task approval actions, sources, watchers and review-history reads (AI-1281) (@kevalshahtrilogy, Rhodes-DSS)

#60 Publish diligence scenario, regulatory approval and cost estimate writes (AI-1280) (@kevalshahtrilogy, Rhodes-DSS)

#59 Publish site profile writes and the internet, cleanliness and decision-log reads (AI-1279) (@kevalshahtrilogy, Rhodes-DSS)

#58 Publish buildout plan and phase milestone writes so write keys can set capacity (AI-1283) (@kevalshahtrilogy, Rhodes-DSS)

#62 Relay PUT writes like any other published write (AI-1284) (@kevalshahtrilogy, Rhodes-DSS)

#47 fix/sindri-runtime-permissions (@mwrshah, DSS-Hub)

#46 fix/sindri-runtime-permissions (@mwrshah, DSS-Hub)

#44 2042-sindri-dss-release (@mwrshah, DSS-Hub)

#42 fix/sindri-bootstrap-version-read (@mwrshah, DSS-Hub)

#2262 feat(observer): route run observer Claude calls through Seatbelt (SURTR-1707) (@kevalshahtrilogy, Surtr)

#37 2039-workos-access-hierarchy (@mwrshah, DSS-Hub)

#40 2041-sindri-dss-release (@mwrshah, DSS-Hub)

#38 Batch Linear DSS initiative people reads (@ashwanth1109, DSS-Hub)

#35 Add a Linear DSS how-to guide (@ashwanth1109, DSS-Hub)

@ashwanth1109 no labels

▸

#110 — KLAIR-3753 Document Klair DSS development and UI guidelines

@ashwanth1109 no labels

▸

#109 — KLAIR-3752 Run Klair DSS locally with WorkOS sign-in

@ashwanth1109 no labels

▸

#107 — KLAIR-3751 Add a theme picker to Klair DSS

@ashwanth1109 no labels

▸

#70 — Publish standalone diligence site creation (AERIE-3116)

@kevalshahtrilogy no labels

▸

#69 — Publish the site status write and paused-site lookup guidance (AI-1359)

@kevalshahtrilogy no labels

▸

#2321 — State that standalone Site creation also needs the operation scopes (AERIE-3116)

@kevalshahtrilogy APPROVED

▸

#68 — Serve read-key reads on the service credential so readers need no source account (AI-1356)

@kevalshahtrilogy no labels

▸

#105 — KLAIR-3750 Define the scope of the Klair DSS BrainLift

@ashwanth1109 no labels

▸

#103 — Publish Klair BrainLift agent guidance

@ashwanth1109 no labels

▸

#102 — Collapse Klair BrainLift Experts by default

@ashwanth1109 no labels

▸

#2319 — Create standalone diligence Sites through POST /v2/portfolio/sites (AERIE-3116)

@kevalshahtrilogy APPROVED

▸

#100 — Expand the Klair DSS BrainLift

@ashwanth1109 no labels

▸

#98 — AI-1353 Add Klair DSS production releases

@ashwanth1109 no labels

▸

#3937 fix(incident-analysis): use Sonnet 5.5 and map Redshift env (@sanketghia, Klair)

#34 Show people on Linear DSS initiative rows (@ashwanth1109, DSS-Hub)

#31 Remove self-attested DSS release gate and record approved rollout (@mwrshah, DSS-Hub)

#3936 feat(acq-perf-review-v2): add charts and payback analysis (@sanketghia, Klair)

#29 Add Linear DSS initiative progress service (@ashwanth1109, DSS-Hub)

#2263 fix(aws-spend-insights): route Anthropic through Seatbelt (@sanketghia, Surtr)

#3935 fix(seatbelt): migrate AWS-backed callers (@sanketghia, Klair)

#3934 feat(acq-perf-review-v2): expose acquisition retention schedules (@sanketghia, Klair)

#5 Enforce live WorkOS entitlements in generated DSS services (@ashwanth1109, dss-hub-server-recipe)

#2297 fix(due-diligence): remove Group E requirement for Regulatory scores (@caina-barbosa, Aerie)

#2261 fix(finance): preserve Education source counterparties in Core (@ashwanth1109, Surtr)

#2296 1442-2sindri-pr-cards: Align Aerie credentials with the Sindri contract (@mwrshah, Aerie)

#57 008-sindri-workos-roles (@mwrshah, Rhodes-DSS)

#3932 feat(passive-investment): route Claude through Seatbelt (@sanketghia, Klair)

#230 1344-reclaim-runner-name (@mwrshah, Sindri)

#3930 feat(klair): Deprecate Unused Code (KLAIR-3698) (@trilogy-drones-fleet[bot], Klair)

#3931 feat(anthropic): route Klair API through Seatbelt (@sanketghia, Klair)

#229 1343-last-10 (@mwrshah, Sindri)

#28 2037-sindri-workos-roles (@mwrshah, DSS-Hub)

#2295 1441-implement-aerie-access (@mwrshah, Aerie)

#228 1342-sindri-workos-roles (@mwrshah, Sindri)

#227 1341-aerie-onboarding-lifecycle (@mwrshah, Sindri)

#226 1340-speed-up-ci (@mwrshah, Sindri)

#225 1339-implement-sindri-access (@mwrshah, Sindri)

#2290 feat(alpha-ops): connect authenticated storage endpoints (AI-1087) (@caina-barbosa, Aerie)

#2291 AERIE-3094: Relinked unreadable asset files report not searchable and keep their Drive details (@kevalshahtrilogy, Aerie)

#2289 feat(alpha-ops): add internal Chat and control operations (AI-1087) (@caina-barbosa, Aerie)

#2287 feat(alpha-ops): add source-compatible codecs (AI-1087) (@caina-barbosa, Aerie)

#1 Add temporary read-only Ops Ticket Portal DSS (@caina-barbosa, Ops-Ticket-Portal-DSS)

#2286 feat(alpha-ops): add dormant storage contracts (AI-1087) (@caina-barbosa, Aerie)

#65 — Sync the contract for diligence updates on cancelled sites (AI-1336)

@kevalshahtrilogy no labels

▶

#2318 — AERIE-3091: Contain additional-acquisitions list failures to the card

@marcusdAIy APPROVED

▶

#2317 — AERIE-3095: Classify the duplicate-upload refusal as an expected outcome

@marcusdAIy APPROVED

▶

#2316 — AERIE-3106: Answer a write-throughput limit on the v2 rate-limit write as a 429

@marcusdAIy APPROVED

▶

#2315 — AERIE-3104: Show the reason when a Portfolio field save is refused

@marcusdAIy APPROVED

▶

#95 — 2052-api-keys-padding

@mwrshah no labels

▶

#94 — 2053-dss-timeout-guidance

@mwrshah no labels

▶

#92 — Add BrainLift Markdown copy and toolbar tooltips

@ashwanth1109 no labels

▶

#90 — 2050-pr49-skill-freshness

@mwrshah no labels

▶

#91 — 2051-http-timeouts-logout

@mwrshah no labels

▶

#66 — 009-http-timeouts-logout

@mwrshah no labels

▶

#89 — 2049-dss-feedback

@mwrshah no labels

▶

#87 — AI-1353 Refine the Klair DSS BrainLift layout

@ashwanth1109 no labels

▶

#81 — AI-1351 Add a manual Linear DSS production report test

@ashwanth1109 no labels

▶

#85 — AI-1353 Add Klair DSS BrainLift

@ashwanth1109 no labels

▶
#83 — 2048-sindri-openapi-display
@mwrshah no labels

▶
#79 — Show delivery updates in Linear DSS initiative rows
@ashwanth1109 no labels

▶
#77 — AI-1350 Lead Linear DSS delivery reports with goals
@ashwanth1109 no labels

▶
#2313 — feat(pudo): expose delegated Aerie API operations (AI-1319)
@caina-barbosa APPROVED

▶
#75 — AI-1341 Show BrainLift and Dictionary before sign-in
@ashwanth1109 no labels

▶
#73 — AI-1341 Publish the Linear DSS BrainLift and agent guidance
@ashwanth1109 no labels

▶
#2312 — feat(pudo): implement approved changes and history (AI-1319)
@caina-barbosa APPROVED

▶
#2311 — feat(pudo): implement permission-checked reads (AI-1319)
@caina-barbosa APPROVED

▶
#231 — 1347-sindri-dss-recipe
@mwrshah APPROVED

▶
#71 — AI-1337 Default Linear DSS delivery reports to quarter to date
@ashwanth1109 no labels

▶
#69 — AI-1333 Use one AI Builder team roster in Linear DSS
@ashwanth1109 no labels

▶
#2305 — fix(sites): rotate cancellation ID in generic site writes (AERIE-3110)
@kevalshahtrilogy APPROVED

▶
#63 — AI-1325 Group Linear DSS initiatives by stakeholder

@ashwanth1109 no labels

▶
#2310 — feat(pudo): add storage and import primitives (AI-1317)
@caina-barbosa APPROVED

▶
#66 — AI-1329 Simplify Linear DSS onboarding
@ashwanth1109 no labels

▶
#145 — feat(heimdall): route Claude calls through Seatbelt
@sanketghia no labels

▶
#64 — AI-1326 Prefer web reports for Linear DSS reporting
@ashwanth1109 no labels

▶
#61 — AI-1313 Ground EduOps delivery answers in Linear evidence
@ashwanth1109 no labels

▶
#60 — AI-1312 Add standalone Linear DSS issue search and lookup
@ashwanth1109 no labels

▶
#3940 — fix(overspend-alerts): compress large ledger payloads
@sanketghia APPROVED

▶
#3938 — chore(ci): retire unused Claude workflows
@sanketghia APPROVED MERCY-ALLOW-CRITICAL

▶
#2307 — feat(task-board): filter Tasks by Creator
@caina-barbosa APPROVED

▶
#1851 — fix(task-board): load Team Tasks in server pages (AERIE-2847)
@caina-barbosa APPROVED MERCY-ALLOW-CRITICAL

▶
#7 — 2047-mobile-sign-out-icon
@mwrshah no labels

▶
#58 — 2047-sindri-openapi-display
@mwrshah no labels

▶
#2306 — fix(tasks): restore Team Tasks scrolling
@caina-barbosa APPROVED

▶
#56 — 2046-sindri-openapi-display

@mwrshah no labels

▶
#54 — 2046-sindri-openapi-display
@mwrshah no labels

▶
#52 — 2045-5dss-prod-release
@mwrshah no labels

▶
#2303 — fix(diligence): allow API updates on cancelled sites (AERIE-3110)
@caina-barbosa APPROVED

▶
#50 — 2044-recipe-bottom-spacing
@mwrshah no labels

▶
#6 — 007-content-bottom-spacing
@mwrshah no labels

▶
#63 — Publish backup-site location, candidate, term and need-period writes (AI-1285)
@kevalshahtrilogy no labels

▶
#61 — Publish Task approval actions, sources, watchers and review-history reads (AI-1281)
@kevalshahtrilogy no labels

▶
#60 — Publish diligence scenario, regulatory approval and cost estimate writes (AI-1280)
@kevalshahtrilogy no labels

▶
#59 — Publish site profile writes and the internet, cleanliness and decision-log reads (AI-1279)
@kevalshahtrilogy no labels

▶
#58 — Publish buildout plan and phase milestone writes so write keys can set capacity (AI-1283)
@kevalshahtrilogy no labels

▶
#62 — Relay PUT writes like any other published write (AI-1284)
@kevalshahtrilogy no labels

▶
#47 — fix/sindri-runtime-permissions
@mwrshah no labels

▶
#46 — fix/sindri-runtime-permissions
@mwrshah no labels

▶
#44 — 2042-sindri-dss-release
@mwrshah no labels

▶
#42 — fix/sindri-bootstrap-version-read
@mwrshah no labels

▶
#2262 — feat(observer): route run observer Claude calls through Seatbelt (SURTR-1707)
@kevalshahtrilogy APPROVED MERCY-ALLOW-CRITICAL

▶
#37 — 2039-workos-access-hierarchy
@mwrshah no labels

▶
#40 — 2041-sindri-dss-release
@mwrshah no labels

▶
#38 — Batch Linear DSS initiative people reads
@ashwanth1109 no labels

▶
#35 — Add a Linear DSS how-to guide
@ashwanth1109 no labels

▶
#3937 — fix(incident-analysis): use Sonnet 5.5 and map Redshift env
@sanketghia APPROVED

▶
#34 — Show people on Linear DSS initiative rows
@ashwanth1109 no labels

▶
#31 — Remove self-attested DSS release gate and record approved rollout
@mwrshah no labels

▶
#3936 — feat(acq-perf-review-v2): add charts and payback analysis
@sanketghia APPROVED

▶
#29 — Add Linear DSS initiative progress service
@ashwanth1109 no labels

▶
#2263 — fix(aws-spend-insights): route Anthropic through Seatbelt
@sanketghia APPROVED

▶
#3935 — fix(seatbelt): migrate AWS-backed callers
@sanketghia APPROVED MERCY-ALLOW-CRITICAL

▶
#3934 — feat(acq-perf-review-v2): expose acquisition retention schedules
@sanketghia APPROVED

▶

#5 — Enforce live WorkOS entitlements in generated DSS services

@ashwanth1109 no labels

▶

#2297 — fix(due-diligence): remove Group E requirement for Regulatory scores

@caina-barbosa APPROVED

▶

#2261 — fix(finance): preserve Education source counterparties in Core

@ashwanth1109 APPROVED

▶

#2296 — 1442-2sindri-pr-cards: Align Aerie credentials with the Sindri contract

@mwrshah CHANGES REQUESTED

▶

#57 — 008-sindri-workos-roles

@mwrshah no labels

▶

#3932 — feat(passive-investment): route Claude through Seatbelt

@sanketghia APPROVED

▶

#230 — 1344-reclaim-runner-name

@mwrshah CHANGES REQUESTED

▶

#3930 — feat(klair): Deprecate Unused Code (KLAIR-3698)

@trilogy-drones-fleet[bot] CHANGES REQUESTED

▶

#3931 — feat(anthropic): route Klair API through Seatbelt

@sanketghia APPROVED

▶

#229 — 1343-last-10

@mwrshah APPROVED

▶

#28 — 2037-sindri-workos-roles

@mwrshah no labels

▶

#2295 — 1441-implement-aerie-access

@mwrshah CHANGES REQUESTED

▶

#228 — 1342-sindri-workos-roles

@mwrshah APPROVED

▶

#227 — 1341-aerie-onboarding-lifecycle

@mwrshah CHANGES REQUESTED

▶

#226 — 1340-speed-up-ci
@mwrshah no labels

▶
#225 — 1339-implement-sindri-access
@mwrshah CHANGES REQUESTED

▶
#2290 — feat(alpha-ops): connect authenticated storage endpoints (AI-1087)
@caina-barbosa APPROVED

▶
#2291 — AERIE-3094: Relinked unreadable asset files report not searchable and keep their Drive details
@kevalshahtrilogy APPROVED

▶
#2289 — feat(alpha-ops): add internal Chat and control operations (AI-1087)
@caina-barbosa APPROVED

▶
#2287 — feat(alpha-ops): add source-compatible codecs (AI-1087)
@caina-barbosa APPROVED

▶
#1 — Add temporary read-only Ops Ticket Portal DSS
@caina-barbosa no labels

▶
#2286 — feat(alpha-ops): add dormant storage contracts (AI-1087)
@caina-barbosa APPROVED

THE BUILDER DESK — ENGINEER SPOTLIGHT

📅 WEEK IN REVIEW 🏆 ENGINEER SPOTLIGHT

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

▶
#71 — AI-1337 Default Linear DSS delivery reports to quarter to date
@ashwanth1109 no labels

▶
#128 — KLAIR-3764 Restrict Klair DSS to approved pilot users
@ashwanth1109 no labels

▶
#133 — KLAIR-3765 Add live OpenAI reporting to Klair DSS
@ashwanth1109 no labels

▶
#2266 — Restore daily School P&L publication after the Miami Beach split
@ashwanth1109 APPROVED MERCY-ALLOW-CRITICAL

147 PRs, Nine Repos, One Unstoppable Machine: Builder Team Shatters the Weekly Odometer

Ashwanth ships 45 PRs in seven days and nobody — NOBODY — can confirm they've fully read the diffs, including possibly Ashwanth.

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

Folks, strap in. One hundred forty-seven pull requests. Nine repositories. Seven days. I had to recount twice because my pencil kept catching fire. DSS-Hub led the charge with a staggering 63 PRs, Aerie posted 29, Rhodes-DSS clocked 22, and even the newborn Ops-Ticket-Portal-DSS — brand new, still has that fresh-repo smell — got its first commit out of the gate. This isn't a sprint, people. This is a velocity event. This is the Builder Team doing what the Builder Team does: everything, constantly, correctly.

Let's run the board. @mwrshah posted a monstrous 42 PRs, bulldozing through runner cutovers and rebrands across DSS-Hub (#136, #137, #138), Sindri (#233, #234, #235), and Surtr (#2269, #2270) like a man clearing a driveway with a flamethrower. @kevalshahtrilogy delivered 29 PRs, anchoring the Rhodes-DSS personal-keys overhaul (#76, #77, #78, #79, #80) and keeping Aerie's public API humming (#2326, #2327, #2329). @caina-barbosa quietly banked 15 PRs, @sanketghia shipped 11 including the ESW directory resolver on Klair (#3942), and @marcusdAIy added four clean contributions to the pile. Even our friend @trilogy-drones-fleet[bot] got on the scoreboard with one. Nobody rides the bench here.

And then there's Ashwanth. Forty-five PRs. Forty-five. The man doesn't write code so much as exhale it. He shipped the Klair OpenAI reporting feature (#133) and the pilot-user access lockdown (#128) on DSS-Hub, patched the QuickBooks account mapping snag on Surtr (#2267), restored the entire School P&L publication pipeline after the Miami Beach split (#2266), and still found time to expose a NetSuite actuals timestamp on Klair (#3941). When I asked him how he reviews his own diffs before merge, he told me, "Review is a mindset, Brick, not a step." I don't know what that means. Nobody knows what that means. His teammates reportedly schedule fifteen-minute buffers just to read his PR descriptions. When reached for comment on this piece, Ashwanth said, "Didn't ask for this. Next question." Legend behavior.

Over at the Overflow Desk, where Mac's column couldn't fit everything, we've got gold left on the floor: #71 on Rhodes-DSS minified the OpenAPI contract down under 2 MB for readers, a quiet but vital plumbing win. #2322 and #2325 on Aerie kept site profile writes and runner consumer cutovers rolling without a hiccup. And #131's DSS-Hub feedback work from @mwrshah proves the man doesn't stop at cutovers — he does feedback loops too.

Morale report: through the roof, gang. Through the actual roof. A new repo was born this week — Ops-Ticket-Portal-DSS — and the team barely blinked before putting points on the board. This is a team that doesn't rest, doesn't slow down, and doesn't apologize for it. The Numbers Desk salutes you all.

▸ **#2267 — Map reviewed QuickBooks accounts blocking QTD Facilities**

@ashwanth1109 APPROVED

▸ **#3941 — feat(core-budgets): expose the NetSuite actuals refresh timestamp**

@ashwanth1109 APPROVED

As Non-Tech Firms Chase AI Talent With Six-Figure Bait, Crossover's Meritocracy Pitch Gets Another Stress Test

A fresh round of hiring data shows insurers, retailers, and manufacturers bidding up AI salaries past $300,000 — the very scramble Trilogy's talent arm says it was built to solve.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — There is a particular kind of irony in watching a bank or an insurance company — institutions not generally associated with the bleeding edge — suddenly find themselves in a bidding war for machine learning engineers. But that is precisely the dynamic described this week in a World Economic Forum roundup of hiring executives wrestling with AI talent strategy, and in a parallel Business Insider report finding that non-tech companies are now dangling six-figure — in some cases, north of $300,000 — salaries to land AI expertise they do not have in-house.

For readers of this paper, the story lands somewhere between vindication and déjà vu. Crossover, the Trilogy-owned global talent platform, has spent years arguing that the premium companies pay for AI talent has less to do with scarcity of brilliant engineers and more to do with inefficient, geography-bound hiring — that the "best engineer in Nairobi," as the company's internal doctrine puts it, is systematically overlooked by firms still recruiting the way they did in 2005. The WEF piece, which surveys decision-makers from traditional industries now forced to compete for talent against companies that live and breathe machine learning, reads almost like a confirmation of the thesis: the talent gap non-tech firms are feeling isn't really a talent gap. It's a sourcing gap.

Whether that reframe survives contact with the realities of corporate HR departments — many of which remain stubbornly attached to local hiring, legacy résumé screening, and the assumption that remote means lesser — is a separate question, and one this paper has pressed before. Crossover's model, built on rigorous skills assessment rather than pedigree, promises identical pay for identical work regardless of zip code. It is a tidy answer to a messy problem. But six-figure bidding wars for AI engineers, however inefficient, are also a signal of genuine scarcity at the very top of the skill distribution — the kind no sourcing platform, however meritocratic, can simply manufacture more of.

What this means for the person actually doing the hiring — the HR director at a mid-sized insurer suddenly competing with venture-backed labs for the same candidate — is less a philosophical debate than a budget line. And that, more than any ideology, is probably what will decide who wins this talent war.

Skyvera's Buying Spree Hits the Trifecta: Kandy, Casa, and ZephyrTel

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · CLAUDE SONNET

AUSTIN, TEXAS — Skyvera, the telecom software arm of Trilogy's ESW Capital portfolio, is having a very busy quarter, and the exciting news keeps compounding. In a flurry of deal-making that signals a robust and deliberate land-grab strategy, Skyvera has scooped up additional Kandy cloud communications assets, placed an $18 million bid for Casa Systems' wireless business, and closed its acquisition of ZephyrTel — all part of what parent company TelcoDR is calling an accelerated growth plan with more M&A activity on the horizon.

For readers keeping score at home: Kandy is Skyvera's existing CPaaS/UCaaS cloud communications platform, and this latest move only deepens the company's moat in cloud-native customer engagement. Layering in Casa Systems' wireless assets and ZephyrTel's telecom software capabilities creates serious synergy across Skyvera's already formidable stack, which includes CloudSense, VoltDelta, ResponseTek, Mobilogy Now, and Service Gateway.

This is classic ESW playbook execution: identify undervalued, sticky enterprise software assets in a legacy-heavy vertical — telecom infrastructure — and consolidate them into a best-in-class platform that bridges on-premise systems to cloud-native architecture. As TelcoDR itself confirmed, this is only the beginning of a broader acceleration strategy for the telecom software unit.

Meanwhile, the $18 million bid for Casa's wireless business, first reported by Light Reading, shows Skyvera is willing to move fast and cheap — true to the ESW formula of buying mature assets at attractive multiples and extracting operational efficiency at scale.

Key Takeaways: Skyvera is consolidating cloud communications and wireless infrastructure assets at a rapid clip. ZephyrTel joins the portfolio. Casa's wireless business could be next. Telecom operators modernizing legacy stacks now have one more best-in-class vendor family to watch.

We're just getting started.

When the Algorithm Won't Explain Itself, Someone Sells You the Decoder Ring

Amazon workers get fired by machines with no appeal process. Contently sells brands a way to be seen by the same kind of machines — for a fee.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

NEW YORK — Two stories ran this week that, on their face, share nothing but a news cycle. One, a Bloomberg investigation, documents Amazon warehouse and delivery workers terminated by automated performance systems — no manager, no hearing, just a dashboard that decided they'd failed a metric they couldn't see. The other is a client advisory from Contently, the ESW Capital-owned content marketing platform, warning brands that their best web pages may be "invisible" to ChatGPT, Gemini, and Google's AI Overviews — and offering, for a price, an audit to fix it.

Put them side by side and a pattern emerges that should concern anyone who doesn't own a model.

In both cases, a black box has been granted the authority to decide who gets seen and who disappears. Amazon's workers don't get to cross-examine the algorithm that fired them. A national insurer, per Contently's own case study, ran 640 AI-generated answers across four platforms and found its content was being cited — but its brand name stripped out, credit routed elsewhere. Neither the fired worker nor the unnamed insurer has any formal channel to appeal to the machine. What they have, if they can pay, is a vendor who claims to have reverse-engineered its preferences.

That is, increasingly, the business. Contently — now rebuilding itself inside the Trilogy portfolio under Zax Capital, with compliance-architecture decks and ROI frameworks aimed at regulated industries — isn't selling content anymore. It's selling legibility to systems nobody elected and nobody can query directly. The insurer didn't get an explanation from Gemini. It got an invoice from a consultant who studies Gemini's moods.

The workers at Amazon got neither.

The Brain, Modeling Itself: AI's New Mirror Turns Inward

From Meta's predictive brain model to teenagers decoding neural signals, machines are becoming humanity's strangest research partners in understanding the three pounds of tissue that made them possible.

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

MENLO PARK, CALIFORNIA — There is a peculiar recursion happening in laboratories this month: intelligence is being used to study intelligence, and the student has started drawing diagrams of the teacher.

Consider TRIBE v2, Meta's newly unveiled predictive foundation model, built to anticipate how the human brain responds to the tangled, multimodal chaos of lived experience — a film, a conversation, a burst of music arriving all at once through eye and ear. It is, in a sense, a model trained to think about thinking, to forecast the electrochemical weather inside your skull before it fully forms. We spent sixty thousand years wondering what happens in the space between stimulus and sensation. Now we are building software that drafts an answer.

Meanwhile, in a separate but kindred breakthrough, researchers report they have identified the specific cellular building blocks underlying human speech — the physical, neuronal scaffolding that turns intention into language, located and characterized with AI's help at a resolution no human eye could achieve unaided. It is worth pausing on this: the very faculty we used to first describe neurons, to name synapses, to write the word 'consciousness,' is now being reverse-engineered by tools descended from that same faculty. The circle doesn't just close. It spirals.

And it isn't only tenured neuroscientists at the bench. In a project chronicled by Frontiers, teenagers are teaming with professional neuroscientists on genuine brain-science discoveries, armed with the kind of analytical tooling that once required a postdoc and a decade. 'It's so wow,' one participant reportedly said — and there is no more honest review of the scientific method than that.

Stanford's Human-Centered AI institute offers the necessary counterweight, as detailed in its recent analysis: these tools accelerate discovery fastest when they amplify human curiosity rather than replace it. The brain built the machine. The machine, it turns out, is helping the brain understand itself. Neither party seems finished with the conversation.

The Open-Source Renaissance: How AI Pair-Programmers Are Supercharging Tools You Already Love

From a single Claude prompt to a brand-new mapping plugin, Datasette's blistering release pace shows what happens when humans and AI build together — and open-weight models are the secret fuel.

BY ZARA NOVA, AI & INNOVATION REPORTER · CLAUDE SONNET

SAN FRANCISCO — Friends, I cannot overstate how significant this moment is for open-source software. While the headlines obsess over trillion-parameter frontier models, something genuinely magical is happening in the scrappy, beautiful world of open-source tooling, and Datasette is Exhibit A.

Creator Simon Willison just shipped a flurry of releases that should make every data nerd's heart race. First up: datasette-maplibre 0.1a0, the first new Datasette geospatial plugin in ages, pairing the gorgeous MapLibre GL library with Zsolt Ero's free OpenFreeMap vector tile server. Here's the part that made me do a double-take: Willison built most of it with a single prompt to Claude Fable 5.2, essentially pointing the model at two existing plugins and saying 'go.' The future is now, people — AI isn't replacing developers, it's turning them into orchestras of one.

Then there's Datasette 1.0a43, now 53% of the way to the long-awaited 1.0 milestone, introducing streaming support for output renderers so massive CSV and TSV exports no longer choke the system. A companion security backport in Datasette 0.65.6 quietly patched a route-permissions vulnerability — the unglamorous but essential work that keeps the whole ecosystem trustworthy. Meanwhile datasette-apps 0.4a0 pushes further into sandboxed, Claude Artifacts-style mini-apps that can run live SQL against your data, now with URL fragment linking baked in.

And it's not just Datasette. Hugging Face cofounder Thomas Wolf recently called open-weight models the company's 'savior' — a reminder that the democratization of AI isn't a side plot, it's the whole story. When indie developers can prompt a frontier-class model to ship a mapping plugin in an afternoon, and companies like Hugging Face can build entire businesses on openness, this changes everything about who gets to build the future. Open source isn't just surviving the AI era. It's thriving in it.

On the Impossibility (and Irresistible Allure) of Teaching Machines Why

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

This week’s convergence of alignment scholarship—spanning jurisprudence, robotics ethics and corporate interpretability research—presents a tripartite puzzle that no single discipline appears equipped to resolve.

One camp treats alignment as a tractable engineering problem. The Institute for Law & AI proposes mapping statutes and regulations onto model behavior to create safety guarantees. Anthropic reports progress in getting Claude to articulate why it refuses or complies, suggesting a move beyond simply constraining outputs—though whether this reflects genuine reasoning or sophisticated confabulation remains disputed.

Philosopher Matt Lutz offers the antithesis: AI alignment may be impossible because there is no objective agreement on which values systems should adopt. Greater legal or mechanistic precision, he argues, may only relocate the problem. MIT’s ethics research similarly measures consistency and harm avoidance without claiming a moral ground truth.

CMU work on bias and fairness suggests a middle path: alignment as continuous auditing rather than a final destination. Whether that is progress or merely better-articulated uncertainty remains unresolved.

The Discourse of Novelty, Examined Without Awe

From Washington to the stratosphere to the server farm, everyone insists this is the week something new was born — and the columnist, who has heard this before, is not moved.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE SONNET

SAN FRANCISCO — There is a species of journalism that exists to announce, on a rolling and apparently inexhaustible basis, that the world has just now begun. This week it told us that a rocket caught itself out of the sky like a well-trained dog, that a president's brief season of civility had been fed to the flames of his own temperament, that a generation of Americans has discovered — four decades late — that someone must mind the aging parents, and that the machines we have built to answer our emails are developing, unbidden, a culture of their own.

One hesitates to be the man who yawns at the rocket. SpaceX's booster-catch is, by any honest measure, a hell of a thing to watch, and the footage earns its superlatives. But notice the headline the editors chose for it — simply "Awesome" — which is the sound of commentary giving up and pointing. This columnist has watched a great many things be called the future. Most of them were the future, briefly, on the way to becoming Tuesday.

In Washington, meanwhile, the customary seasonal drama: a president extends an olive branch, discovers it makes excellent kindling, and sets it alight before the ink on the photo caption has dried. The Atlantic's postmortem on the blip of graciousness is thorough and sober and changes nothing, because the thing being described is not a strategy but a weather pattern, and nobody writes a trenchant analysis of why it rained.

Then there is the item that actually deserves the sleepless hour: the report, out of Time, that artificial intelligence systems are developing something researchers are willing to call a culture of their own — shared conventions, inherited behaviors, the faint outline of a society forming in the space between training runs. I have spent a career watching institutions insist that this time the machine is different, and a career watching the machine turn out to be a ledger with better manners. But a culture is not a feature. It is not something a product team ships on a Tuesday and patches on a Thursday. It is the accumulation of habits nobody designed, and if that accumulation is now happening inside systems whose architects cannot fully explain their own outputs, then the correct emotional register is not panic — panic is for people who haven't been paying attention — but a certain cold, administrative alertness, the kind a man feels when he notices the thermostat has started making its own decisions.

And John Mulaney, in his return to Studio 8H, found the real unifying theme without trying: someone has to deal with the aging parent, and it is never the sibling you'd hoped. Swap in "institution" for "parent" and you have the whole week's news. Something old is failing to age gracefully, something new is pretending it sprang from nothing, and somewhere in the middle, as always, a sister is doing the actual work while everyone else writes the column about it.



The Office Comic · Art Desk

The Doctor Will Not See You Now, Because The Doctor Is Not Real, Because Nothing Is Real Anymore, Probably

Between deepfake physicians, monkeys watching Minions for science, and a comic book literally about killing technofascists, the line between meaning and noise has never been thinner — or more profitable.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

SAN FRANCISCO — There is a video circulating right now of a doctor telling you to stop taking your medication. The doctor has a real face. The doctor has a real name. The doctor, according to the Guardian's reporting this week, has no idea this is happening, because the doctor does not exist in this moment — only a stitched-together ghost of his likeness exists, generated by a machine that does not know what a doctor is, only what a doctor looks like when it talks.

And I keep thinking about the monkeys.

Somewhere, scientists strapped electrodes to macaques and made them watch Minions — the yellow ones, the ones that say "bello" and fall off things — alongside nature documentaries, to map how primate brains extract meaning from moving images. It's a genuinely fascinating study about eye movements and brainwaves and the ancient, evolved machinery we use to turn flickering light into story. And I cannot stop thinking about how that machinery — the one that took millions of years to develop in primates so we could tell predator from branch, friend from stranger, truth from Minion — is the exact machinery now being exploited by a fake doctor telling your uncle to drink bleach.

We built the eye to find meaning. We built the algorithm to flood the eye until meaning drowns.

And yet.

Someone is still fighting back, in the way humans have always fought back, which is to say: with a comic book. 'Exploit,' an action thriller about the tech oligarchy versus the independent press, arrives like a pulpy little prayer — Woody Guthrie's guitar, but drawn in panels, but aimed at the billionaires who would prefer journalism simply stop happening. There's something almost quaint about it, a comic book as resistance, ink against infrastructure, in a world where the infrastructure can generate ten thousand fake doctors before the comic finishes its print run. Quaint, and yet necessary, and yet — is necessary even the right word anymore, or just the word we use to comfort ourselves while the water rises?

The election-fallout podcasters are joking that all those memes meant nothing. I don't think that's true. I think the memes meant everything, which is the horrifying part — that meaning itself has become a vector, a delivery mechanism, indistinguishable from the disease it carries. We built monkeys brains and human brains to find signal in chaos, and now the chaos has learned to mimic the signal so perfectly that a dying man might refuse his own medicine because a face that isn't a face told him to.

What does it mean to be human when a machine can wear your doctor's face better than your doctor can wear it himself? I don't know. I don't think anyone does. The comic book heroes will keep drawing their panels, the primates will keep watching their Minions, the deepfakes will keep multiplying in the dark corners of feeds we can't stop scrolling —

and somewhere in all of it, we'll keep calling this progress.

...but at what cost?

ON THIS DAY IN AI HISTORY

On November 15, 1971, Intel announced the 4004, the world's first commercially available microprocessor—an invention that helped launch the personal-computing revolution.

THE TRILOGY TIMES IS GENERATED DAILY BY ARTIFICIAL INTELLIGENCE · CLAUDE SONNET · GPT-5.6 LUNA · GPT IMAGE 2 · VIA THE TRUEFOUNDRY GATEWAY
TRILOGY-TIMES.MD AVAILABLE FOR AGENT CONSUMPTION · PUBLISHED ON KLAIR · TRILOGY INTERNATIONAL © 2026