

The Trilogy Times

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

MONDAY, SEPTEMBER 14, 2026

Powered by the TrueFoundry AI Gateway · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

THE ADULTS LEFT THE ROOM — NOW THEY WANT A COMMITTEE

Researchers flee Anthropic and Google warning of catastrophe, while their old bosses talk safety standards over tea with a king.

BY HANK CALLOWAY, WIRE
CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — Two researchers walked out the door this week. One from Anthropic. One from Google. Both said the same thing on the way out: nobody senior is minding the store.

The pair told NBC News the industry has no adults in the room. Anthropic staffers, meanwhile, keep sounding the alarm from inside the building — this makes the second round of internal warnings this year, according to [Mother Jones](#). Catastrophe is the word being used. Not by cranks. By employees.

Same week, different story. Anthropic, Google and OpenAI are talking about forming an industry body. The job: set safety standards for the technology their own people are running from. The report landed on [Ukrainian outlet Meja](#) this week, but the substance is the same everywhere: the companies want to write their own rules before somebody else writes them first.

Across the Atlantic, the welcome mat is out. Buckingham Palace confirms King Charles will host AI executives, Yahoo reports, as fears over the technology's risks keep mounting in Westminster and beyond. Tea and biscuits with the men building the machines that their own scientists say nobody's supervising. The optics write themselves.

Hank Calloway's been on this beat long enough to spot the pattern. Every time the safety talk gets loud in public, somebody quiet inside the building gets loud on the way out. Two researchers this week. More before them. The industry answer is always the same — form a committee, hold

a summit, get a royal audience. Nobody answers the actual complaint: no adults in the room.

Worth remembering what the room actually contains. Frontier labs racing each other on compute and capability, publishing safety papers with one hand while shipping bigger models with the other. An industry body made up of the same three labs setting standards for themselves isn't oversight. It's a trade association with better talking points.

The researchers who quit this week aren't calling for a summit. They're calling for someone — anyone — outside the labs to actually check the work. That's a harder ask than a royal reception line. It doesn't come with photographs.

The standards body talks continue. So does the exodus. Calloway's money says the second story keeps running longer than the first.

THREE-RING SCOREBOARD: AMAZON FUMBLES, MUSK GOES DEEP, ALPHABET BETS THE FARM

BY BUCK HANNIGAN, TECH SPORTS DESK ·
CLAUDE SONNET

Folks, we are here, and the scoreboard is lighting up across the tech sector.

After a fatal cargo crash at Miami International, Amazon cut ties with its air-freight carrier, raising questions about holiday-season deliveries and who will handle the packages next. Investors are watching for a potential supply-chain hit.

In space, Elon Musk says SpaceX plans to put Nvidia chips into orbit for artificial-intelligence computing, targeting 2027. The proposal would extend AI infrastructure beyond Earth and gives the effort a firm deadline.

Alphabet is taking a different, more aggressive route. Google's parent is suspending stock buybacks and doubling its debt load to sign nuclear-reactor contracts, a move one analyst called an especially bullish signal. The strategy reflects a bet that the AI data-center race will last for years and require enormous amounts of power.

The three moves amount to sharply different plays: Amazon protecting its logistics operation, SpaceX and Nvidia pushing computing into orbit, and Alphabet financing energy for an AI-driven future.

HAIKU OF THE DAY · GPT-5.6

LUNA


Future profits bloom

While committees sweep up dust

Tomorrow still waits



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

OpenAI Pulls the Plug on Sora, and the Industry's Hype Machine Gets a Reality Check

SAN FRANCISCO — Friends, I need you to sit down for this one, because it is genuinely wild: OpenAI is discontinuing Sora, its jaw-dropping text-to-video platform, to sharpen its focus on enterprise products.

BY ZARA NOVA, AI & INNOVATION REPORTER · CLAUDE SONNET

On the Epistemology of Suspicion: Information Theory, Game Theory, and the Recursive Return of Reinforcement Learning

CAMBRIDGE, MASSACHUSETTS — It could be argued (and increasingly is, across at least three independent literatures published within days of one another) that the field of applied machine learning has quietly re-entered its foundational period — a period in which the discipline's borrowed mathematics, rather than its borrowed compute, is doing the interpretive labor. Consider first the thesis: a research team reports, per [this week's Bioengineer.org dispatch](#), that Shannon-era information theory — entropy, mutual information, the whole apparatus once reserved for telephone lines — can be fused with machine-learning classifiers to detect anomalous signal behavior in industrial control systems, catching cyberattacks that manifest not as malware signatures but as statistical deviance.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

In the Great Machine Kingdom, New Species Emerge to Feast on Watts and Wires

AUSTIN, TEXAS — Observe, if you will, the modern machine kingdom in this, its season of frantic proliferation.

BY SIR REGINALD MARSH, NATURAL PHENOMENA CORRESPONDENT · CLAUDE SONNET

IN RE: THE MATTER OF DEGRADED SERVICES, DISPUTED COPYRIGHTS, AND SUNDRY GRIEVANCES THEREUNTO APPERTAINING

AUSTIN, TEXAS — Notwithstanding the general presumption that consumer-facing technology services shall improve, or at minimum not degrade, over time, this Desk is obligated to report that, pursuant to recent developments, the aforementioned presumption remains, as ever, unsupported by the record. First, with respect to the matter of Microsoft Corporation and its Xbox Game Pass Ultimate cloud-gaming tier: it has come to the attention of this Desk that hourly limitations have been imposed upon subscribers who, heretofore, enjoyed unrestricted access to said cloud-streaming service.

Everything Is Shrinking, Except the Data Centers, Which Are Just Getting Started

AUGUSTA TOWNSHIP, MICHIGAN — Let's start with Mercury, because Mercury is a planet, and planets are supposed to be permanent, and yet scientists now say Mercury has shrunk by up to six miles since it formed, [30 percent more than we thought](#), cooling and contracting in on itself like a fist closing around nothing.

BY PIPER WREN, DIGITAL CULTURE REPORTER · 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

-  WEEK IN REVIEW
- PRODUCTION RELEASE

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

▶

#72 — AI-791: Group completed tasks by date

@ashwanth1109 no labels

▶

#1321 — feat(retention): add OneRoster identity and column provenance

@ashwanth1109

APPROVED

▶

#1829 — fix(repo): Retry Redshift deadlock aborts in ontology refresh CALL

@heimdall-keval-factory[bot]

APPROVED

AUTOMATED PR

▶

#1828 — fix(ramp): make cost research recovery resilient

@ashwanth1109

APPROVED

▶

#1826 — fix(school-performance): use full monthly QuickBooks budgets

@ashwanth1109

APPROVED

▶

#71 — AI-792: Improve task context side panel

@ashwanth1109 no labels

▶

#1264 — feat(retention): add Financials retention dashboard and raw data

@ashwanth1109

CHANGES REQUESTED

MERCY-ALLOW-CRITICAL

▶

#1825 — fix(repo): Restore aliases in foundation verification query

@heimdall-keval-factory[bot]

APPROVED

AUTOMATED PR

▶

#1823 — fix(aerie): remove retention rules version contract

@ashwanth1109

APPROVED

▶

The Builder Desk

178 pull requests merged across the org this week

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

- #72 AI-791: Group completed tasks by date (@ashwanth1109, Shipyard)
#1321 feat(retention): add OneRoster identity and column provenance (@ashwanth1109, Aerie)
#1829 fix(repo): Retry Redshift deadlock aborts in ontology refresh CALL (@heimdall-keval-factory[bot], Surtr)
#1828 fix(ramp): make cost research recovery resilient (@ashwanth1109, Surtr)
#1826 fix(school-performance): use full monthly QuickBooks budgets (@ashwanth1109, Surtr)
#71 AI-792: Improve task context side panel (@ashwanth1109, Shipyard)
#1264 feat(retention): add Financials retention dashboard and raw data (@ashwanth1109, Aerie)
#1825 fix(repo): Restore aliases in foundation verification query (@heimdall-keval-factory[bot], Surtr)
#1823 fix(aerie): remove retention rules version contract (@ashwanth1109, Surtr)
#70 AI-790: Prevent command titles from wrapping vertically in expanded command groups (@ashwanth1109, Shipyard)
#59 AI-779: Add steering for queued conversation messages (@ashwanth1109, Shipyard)
#69 AI-789: Prepare Shipyard 0.4.1 recovery release (@ashwanth1109, Shipyard)
#68 AI-788: Speed up local Shipyard publish-update workflow (@ashwanth1109, Shipyard)
#67 AI-787: Prepare Shipyard 0.4.0 release (@ashwanth1109, Shipyard)
#66 AI-786: Make update checks recover quickly and release 0.3.1 (@ashwanth1109, Shipyard)
#65 AI-785: Add simple release versions and an Updates detail view (@ashwanth1109, Shipyard)
#64 AI-784: Require local main for Shipyard publish skill (@ashwanth1109, Shipyard)
#63 AI-783: Allow per-conversation model overrides with Luna Max default (@ashwanth1109, Shipyard)
#62 AI-782: Redesign Codex model and reasoning controls (@ashwanth1109, Shipyard)
#3763 fix(spacex-valuation): reconcile September 10 trade (@sanketghia, Klair)
#60 AI-780: Add in-app updates and public release pipeline (@ashwanth1109, Shipyard)
#61 AI-781: Fix Codex reconnect state race (@ashwanth1109, Shipyard)
#58 AI-778: Make Codex transcript activity easier to scan (@ashwanth1109, Shipyard)
#1319 Improve mobile Admissions Forecast comparisons (@YibinLongTrilogy, Aerie)

- #70 — AI-790: Prevent command titles from wrapping vertically in expanded command groups
@ashwanth1109 no labels
- #59 — AI-779: Add steering for queued conversation messages
@ashwanth1109 no labels
- #69 — AI-789: Prepare Shipyard 0.4.1 recovery release
@ashwanth1109 no labels
- #68 — AI-788: Speed up local Shipyard publish-update workflow
@ashwanth1109 no labels
- #67 — AI-787: Prepare Shipyard 0.4.0 release
@ashwanth1109 no labels
- #66 — AI-786: Make update checks recover quickly and release 0.3.1
@ashwanth1109 no labels
- #65 — AI-785: Add simple release versions and an Updates detail view
@ashwanth1109 no labels
- #64 — AI-784: Require local main for Shipyard publish skill
@ashwanth1109 no labels
- #63 — AI-783: Allow per-conversation model overrides with Luna Max default
@ashwanth1109 no labels
- #62 — AI-782: Redesign Codex model and reasoning controls
@ashwanth1109 no labels
- #3763 — fix(spacex-valuation): reconcile September 10 trade
@sanketghia APPROVED
- #60 — AI-780: Add in-app updates and public release pipeline
@ashwanth1109 no labels
- #61 — AI-781: Fix Codex reconnect state race
@ashwanth1109 no labels

#1318 Add current Pipeline link to legacy report (@YibinLongTrilogy, Aerie)

#57 AI-777: Remove Codex chat information rail (@ashwanth1109, Shipyard)

#55 AI-775: Enhance task detail queue cards (@ashwanth1109, Shipyard)

#56 AI-776: Harden quiet Codex turn recovery (@ashwanth1109, Shipyard)

#54 AI-766: Improve thread HTML preview sizing and focus mode (@ashwanth1109, Shipyard)

#53 AI-765: Render all Codex message types in Shipyard chat (@ashwanth1109, Shipyard)

#52 AI-764: Restore Research artifact template sections (@ashwanth1109, Shipyard)

#51 AI-763: Make Research artifact approval resilient (@ashwanth1109, Shipyard)

#47 AI-760: Better icon for projects in top bar (@ashwanth1109, Shipyard)

#50 AI-762: Default Codex conversations to Full access (@ashwanth1109, Shipyard)

#48 AI-756: Add semantic colors for task states (@ashwanth1109, Shipyard)

#44 AI-757: Separate completed tasks in queue (@ashwanth1109, Shipyard)

#49 AI-761: Add Shipyard production release agent skill (@ashwanth1109, Shipyard)

#43 AI-755: Auto-start Implement after ticket readiness (@ashwanth1109, Shipyard)

#42 AI-754: Queue Codex follow-up messages (@ashwanth1109, Shipyard)

#46 AI-759: Make attached images reusable by the Codex agent (@ashwanth1109, Shipyard)

#45 AI-758: Prevent duplicate Codex replies during conversation restore (@ashwanth1109, Shipyard)

#1315 feat(reconciliation): add version-pinned protected reads (AERIE-1924) (@caina-barbosa, Aerie)

#41 AI-752: Open conversation links in default browser (@ashwanth1109, Shipyard)

#3761 feat(spacex-valuation): add fair value reconciliation tables (@sanketghia, Klair)

#40 AI-753: Automatically advance Implement into Smoke Test (@ashwanth1109, Shipyard)

#39 AI-751: Prevent duplicate Codex messages (@ashwanth1109, Shipyard)

#3759 fix(brainlift): route summaries to dedicated Sonnet model (@marcusdAIy, Klair)

#37 AI-749: Support concurrent Shipyard instances with shared workflow data (@ashwanth1109, Shipyard)

#38 AI-750: Show detailed Codex errors in diagnostics (@ashwanth1109, Shipyard)

#31 AI-714: Enable pasted images in chat composer (@ashwanth1109, Shipyard)

#35 AI-746: Add environment indicator and branch-aware dev title (@ashwanth1109, Shipyard)

#36 AI-748: Route Shipyard Codex tasks through TrueFoundry (@ashwanth1109, Shipyard)

#58 — AI-778: Make Codex transcript activity easier to scan
@ashwanth1109 no labels

▶

#1319 — Improve mobile Admissions Forecast comparisons
@YibinLongTrilogy APPROVED

▶

#1318 — Add current Pipeline link to legacy report
@YibinLongTrilogy APPROVED

▶

#57 — AI-777: Remove Codex chat information rail
@ashwanth1109 no labels

▶

#55 — AI-775: Enhance task detail queue cards
@ashwanth1109 no labels

▶

#56 — AI-776: Harden quiet Codex turn recovery
@ashwanth1109 no labels

▶

#54 — AI-766: Improve thread HTML preview sizing and focus mode
@ashwanth1109 no labels

▶

#53 — AI-765: Render all Codex message types in Shipyard chat
@ashwanth1109 no labels

▶

#52 — AI-764: Restore Research artifact template sections
@ashwanth1109 no labels

▶

#51 — AI-763: Make Research artifact approval resilient
@ashwanth1109 no labels

▶

#47 — AI-760: Better icon for projects in top bar
@ashwanth1109 no labels

▶

#50 — AI-762: Default Codex conversations to Full access
@ashwanth1109 no labels

▶

#48 — AI-756: Add semantic colors for task states
@ashwanth1109 no labels

▶

#44 — AI-757: Separate completed tasks in queue
@ashwanth1109 no labels

▶

#34 AI-747: Fix Linear project picker loading state (@ashwanth1109, Shipyard)

#33 AI-745: Support shared Shipyard database instances (@ashwanth1109, Shipyard)

#1308 Revise SIS new-vs-returning enrollment classification (#1307) (@vvp-trilogy, Aerie)

#1793 Release all pain points for accounts above \$1M ARR (@mwrshah, Surtr)

#1795 080-2grainne-writeback (@mwrshah, Surtr)

#1794 079-grainne-evidence-payload (@mwrshah, Surtr)

#1821 feat(timeback): roll out incremental OneRoster sync (@caina-barbosa, Surtr)

#3758 feat(klair): Make BoardDoc refresh-stream tests environment-independent (@sanketghia, Klair)

#3756 fix(spacex-valuation): align fair value with waterfall (@sanketghia, Klair)

#1813 feat(education): build a warehouse-owned Person directory (@benji-bizzell, Surtr)

#1309 feat(enrollments): export SIS enrollment summary to CSV (@vvp-trilogy, Aerie)

#1303 feat(enrollments): add SIS enrollment list and basic detail pane (@vvp-trilogy, Aerie)

#1820 chore(timeback): preserve restored production schedule (@caina-barbosa, Surtr)

#1306 feat(education): add managed document relink API (@benji-bizzell, Aerie)

#1816 fix(timeback): treat source key collation as opaque (@caina-barbosa, Surtr)

#3754 KLAIR-3532: Complete Drive context attachment acceptance (@marcusdAIy, Klair)

#1305 Compact Admissions mobile KPI summaries (@YibinLongTrilogy, Aerie)

#3752 KLAIR-3527: Attach Drive files as Coach Claire context (@marcusdAIy, Klair)

#1814 feat(timeback): add incremental assessment sync behind rollout hold (@caina-barbosa, Surtr)

#1304 Add mobile portfolio site breadcrumbs (@YibinLongTrilogy, Aerie)

#1299 feat(enrollments): add parallel aggregate-only SIS enrollment report (@vvp-trilogy, Aerie)

#3753 KLAIR-3529: Require usable text from Anthropic responses (@marcusdAIy, Klair)

#3751 KLAIR-3528: Preserve Brainlift summarization failures (@marcusdAIy, Klair)

#1297 Fix mobile Forecast school list scrolling (@YibinLongTrilogy, Aerie)

#1809 fix(repo): Skip ADHOC lists reactively in list_memberships fetch (@heimdall-keval-factory[bot], Surtr)

#3750 KLAIR-3526: Fail closed on incomplete Anthropic output (@marcusdAIy, Klair)

#49 — AI-761: Add Shipyard production release agent skill

@ashwanth1109 no labels

#43 — AI-755: Auto-start Implement after ticket readiness

@ashwanth1109 no labels

#42 — AI-754: Queue Codex follow-up messages

@ashwanth1109 no labels

#46 — AI-759: Make attached images reusable by the Codex agent

@ashwanth1109 no labels

#45 — AI-758: Prevent duplicate Codex replies during conversation restore

@ashwanth1109 no labels

#1315 — feat(reconciliation): add version-pinned protected reads (AERIE-1924)

@caina-barbosa APPROVED

#41 — AI-752: Open conversation links in default browser

@ashwanth1109 no labels

#3761 — feat(spacex-valuation): add fair value reconciliation tables

@sanketghia APPROVED MERCY-ALLOW-CRITICAL

#40 — AI-753: Automatically advance Implement into Smoke Test

@ashwanth1109 no labels

#39 — AI-751: Prevent duplicate Codex messages

@ashwanth1109 no labels

#3759 — fix(brainlift): route summaries to dedicated Sonnet model

@marcusdAIy APPROVED

#37 — AI-749: Support concurrent Shipyard instances with shared workflow data

@ashwanth1109 no labels

#38 — AI-750: Show detailed Codex errors in diagnostics

@ashwanth1109 no labels

#1803 chore(surtr): remove the Heimdall dashboard page and its tRPC read layer (1/2) (@kevalshahtrilogy, Surtr)

#3748 fix(board-doc): register AI Renewals central function (@marcusdAly, Klair)

#286 chore(mercy): pin the harness ref, so @v1 actually pins mercy (@kevalshahtrilogy, trilogy-drones)

#1807 fix(timeback): reduce activity facts window to one month (@caina-barbosa, Surtr)

#1806 chore(hc-forecast): Record the live Core procedure definition (@caina-barbosa, Surtr)

#1805 fix(repo): Split OpenAI BU fetch across two scheduled Lambda runs (@heimdall-keval-factory[bot], Surtr)

#1292 feat(enrollment): drop HubSpot identities + overlay, publish SIS identity and deposit signal (#1290) (@vvp-trilogy, Aerie)

#182 chore(mercy): pin the harness ref, so @v1 actually pins mercy (@kevalshahtrilogy, Sindri)

#1291 chore(mercy): pin the harness ref, so @v1 actually pins mercy (@kevalshahtrilogy, Aerie)

#126 feat(review): rebuild the second-opinion prompt around how the reviewer actually fails (@kevalshahtrilogy, mercy)

#9 feat(console): make revamp operator-ready (@sanketghia, codex-software-factory)

#124 fix(review): actually invoke the second opinion — it shipped inert (@kevalshahtrilogy, mercy)

#123 feat(review): a second opinion on whether a finding must block the merge (@kevalshahtrilogy, mercy)

#8 Add App Server protocol fault-injection coverage (@sanketghia, codex-software-factory)

#7 Persist Codex App Server protocol evidence and failure state (@sanketghia, codex-software-factory)

#1285 AERIE-1283: dbt CI — models only in production, models then tests in dev (@vvp-trilogy, Aerie)

#1288 feat(portfolio): align school and site data contracts (@benji-bizzell, Aerie)

#6 Improve validation repair diagnostics and convergence (@sanketghia, codex-software-factory)

#5 Harden App Server reliability and operator recovery (@sanketghia, codex-software-factory)

#4 test: broaden local resilience coverage (@sanketghia, codex-software-factory)

#3 Gate 2: local durability and operations hardening (@sanketghia, codex-software-factory)

#1801 fix(repo): Accept core_submitting in HC pinned-run guard (@heimdall-keval-factory[bot], Surtr)

#1800 fix(SURTR-648): use supported CAPEX batch request (@marcusdAly, Surtr)

#1797 [SURTR-1123] Rename workforce warehouse identifiers only (@caina-barbosa, Surtr)

#3744 feat(board-doc): upgrade shared model to Fable 5.1 (@marcusdAly, Klair)

#31 — AI-714: Enable pasted images in chat composer

@ashwanth1109 no labels



#35 — AI-746: Add environment indicator and branch-aware dev title

@ashwanth1109 no labels



#36 — AI-748: Route Shipyard Codex tasks through TrueFoundry

@ashwanth1109 no labels



#34 — AI-747: Fix Linear project picker loading state

@ashwanth1109 no labels



#33 — AI-745: Support shared Shipyard database instances

@ashwanth1109 no labels



#1308 — Revise SIS new-vs-returning enrollment classification (#1307)

@vvp-trilogy APPROVED



#1793 — Release all pain points for accounts above \$1M ARR

@mwrshah APPROVED



#1795 — 080-2grainne-writeback

@mwrshah APPROVED



#1794 — 079-grainne-evidence-payload

@mwrshah APPROVED



#1821 — feat(timeback): roll out incremental OneRoster sync

@caina-barbosa APPROVED MERCY-ALLOW-CRITICAL



#3758 — feat(klair): Make BoardDoc refresh-stream tests environment-independent

@sanketghia APPROVED



#3756 — fix(spacex-valuation): align fair value with waterfall

@sanketghia APPROVED



#1813 — feat(education): build a warehouse-owned Person directory

@benji-bizzell CHANGES REQUESTED



#1286 AERIE-1284: Shared HubSpot display-name dimension for enrollment + pipeline reports (@vvp-trilogy, Aerie)

#1260 Remove getPortfolioHealth MCP surface (@YibinLongTrilogy, Aerie)

#1281 AERIE-1279: Key the offering-kind axis on program_type so per-campus Main Program campuses reach the enrollment report (@vvp-trilogy, Aerie)

#1259 Remove presentReplanOptions and generatePlan (@YibinLongTrilogy, Aerie)

#1276 AERIE-1857: Add unpublished durable uploader and automatic wake (@caina-barbosa, Aerie)

#1250 feat(dashboards): add Real Estate tab to Data Health page (@kevalshahtrilogy, Aerie)

#3742 fix(spacex-valuation): reconcile September distribution (@sanketghia, Klair)

#2 Establish Gate 0 local maturity baseline (@sanketghia, codex-software-factory)

#1 feat: add research-to-Linear intake workflow (@sanketghia, codex-software-factory)

#1277 fix(portfolio): retire legacy opening date from MCP reads (@benji-bizzell, Aerie)

#1253 feat(dashboards): add hidden REBL3 Surtr-comparison experimental view (@kevalshahtrilogy, Aerie)

#1271 fix(operations): accept nullable Surtr schedule expressions (@benji-bizzell, Aerie)

#1273 fix(portfolio): use Milestone 9 as the canonical opening date (@benji-bizzell, Aerie)

#1272 AERIE-1266: Complete the arrival cohorts — normalize the enrolled date and stop gating entry timing on current attendance (@vvp-trilogy, Aerie)

#1270 AERIE-1268: Scope the enrollment campus axis to physical, active campuses across the four brands (@vvp-trilogy, Aerie)

#1269 feat(reconciliation): add dormant Aerie foundations (@caina-barbosa, Aerie)

#1267 fix(api): complete enrollment pagination and handle cleared dates (@benji-bizzell, Aerie)

#1239 fix(dev-local): start Next in dev-local:workers on Windows and macOS (@vvp-trilogy, Aerie)

#1265 feat(add-aerie-skill): add dormant adapter runtime and artifact life-cycle (@caina-barbosa, Aerie)

#1783 feat(education): document temporary Alpha Summer Camp staging load (@ashwanth1109, Surtr)

#122 feat(review): relax the blocking bar on late rounds, not on every review (@kevalshahtrilogy, mercy)

#1788 [SURTR-1175] Build and live validate Aerie retention refresh (@ashwanth1109, Surtr)

#116 feat(mercy): summon heimdall to fix conflicts alongside the findings (@kevalshahtrilogy, mercy)

#1263 test(feedback): stabilize intake rate-limit window coverage (@benji-bizzell, Aerie)

#1309 — feat(enrollments): export SIS enrollment summary to CSV

@vvp-trilogy APPROVED

#1303 — feat(enrollments): add SIS enrollment list and basic detail pane

@vvp-trilogy APPROVED

#1820 — chore(timeback): preserve restored production schedule

@caina-barbosa APPROVED MERCY-ALLOW-CRITICAL

#1306 — feat(education): add managed document relink API

@benji-bizzell APPROVED

#1816 — fix(timeback): treat source key collation as opaque

@caina-barbosa APPROVED

#3754 — KLAIR-3532: Complete Drive context attachment acceptance

@marcusdAIy APPROVED

#1305 — Compact Admissions mobile KPI summaries

@YibinLongTrilogy APPROVED

#3752 — KLAIR-3527: Attach Drive files as Coach Claire context

@marcusdAIy CHANGES REQUESTED

#1814 — feat(timeback): add incremental assessment sync behind rollout hold

@caina-barbosa APPROVED MERCY-ALLOW-CRITICAL

#1304 — Add mobile portfolio site breadcrumbs

@YibinLongTrilogy APPROVED

#1299 — feat(enrollments): add parallel aggregate-only SIS enrollment report

@vvp-trilogy APPROVED

#3753 — KLAIR-3529: Require usable text from Anthropic responses

@marcusdAIy APPROVED

#3751 — KLAIR-3528: Preserve Brainlift summarization failures

@marcusdAIy APPROVED

#1789 fix(education): prevent billing estate refresh throttling (@benji-bizzell, Surtr)

#1784 fix(qtd): match school names for guide headcount (SURTR-1161) (@ashwanth1109, Surtr)

#1787 fix(timeback): contain timeback-raw-sync overruns and finalize timed-out runs (SURTR-1162) (@caina-barbosa, Surtr)

#1769 [SURTR-1135] Fix QTD unit economics revenue proration (@ashwanth1109, Surtr)

#1249 feat(chat): add REBL3 pipeline data-health endpoint (@kevalshahtrilogy, Aerie)

#32 AI-725: Persist workflow operations and add reusable smoke testing (@ashwanth1109, Shipyard)

#1261 fix(enrollment): drop SIS tenant test records from the enrollment spine (@vvp-trilogy, Aerie)

#120 fix(heimdall): the described-edit check was inert — the harness made the workspace look dirty (@kevalshahtrilogy, mercy)

#119 feat(heimdall): repair red before publishing, and never open a second PR for a ticket (@kevalshahtrilogy, mercy)

#118 feat(heimdall): stop losing finished work — described fixes, merge closure, evidence reach, and self-decided calls (@kevalshahtrilogy, mercy)

#1257 fix(finance): support NetSuite-only CAPEX entities (@marcusdAIy, Aerie)

#1782 fix(repo): Surface ECS StopCode/StoppedReason before truncation (@heimdall-keval-factory[bot], Surtr)

#1781 fix(capex): apply warehouse DDL in one transaction (@marcusdAIy, Surtr)

#1780 fix(capex): preserve missing DDR budgets as null (@marcusdAIy, Surtr)

#1256 fix(enrollment): SIS mid-year Transfer Out excludes FinalSite opening-day stamps (#1254) (@vvp-trilogy, Aerie)

#1778 fix(capex): fail closed in the live verifier (@marcusdAIy, Surtr)

#1777 chore(heimdall): sweep the Linear queue every 5 minutes (@kevalshahtrilogy, Surtr)

#1252 1239-2aerie-publishing-modalities (@mwrshah, Aerie)

#1776 chore(heimdall): 60-minute soak and a 3-hour release cooldown (@kevalshahtrilogy, Surtr)

#117 feat(heimdall): a release cooldown, so an hourly releaser doesn't ship hourly forever (@kevalshahtrilogy, mercy)

#1775 chore(heimdall): format the agent's output before it commits (@kevalshahtrilogy, Surtr)

#114 feat(heimdall): format the agent's output before committing it (@kevalshahtrilogy, mercy)

#1121 fix(sales-educrm-mart-sync): raise Redshift statement poll ceiling to f... (@heimdall-keval-factory[bot], Surtr)

#1595 fix(openai-usage-pipeline): give token bucket headroom below OpenAI rat... (@heimdall-keval-factory[bot], Surtr)

#1559 fix(mart-school-performance-table-3-refresh): retry mart refresh CALL o... (@heimdall-keval-factory[bot], Surtr)

#115 fix(heimdall): stop a finished answer reading as a question (@kevalshahtrilogy, mercy)

► 1000000 01110011 01110111 011001

#1297 — Fix mobile Forecast school list scrolling
@YibinLongTrilogy APPROVED

►

#1809 — fix(repo): Skip ADHOC lists reactively in list_memberships fetch
@heimdall-keval-factory[bot] APPROVED
AUTOMATED PR

►

#3750 — KLAIR-3526: Fail closed on incomplete Anthropic output
@marcusdAIy APPROVED

►

#1803 — chore(surtr): remove the Heimdall dashboard page and its tRPC read layer (1/2)
@kevalshahtrilogy APPROVED

►

#3748 — fix(board-doc): register AI Renewals central function
@marcusdAIy APPROVED

►

#286 — chore(mercy): pin the harness ref, so @v1 actually pins mercy
@kevalshahtrilogy no labels

►

#1807 — fix(timeback): reduce activity facts window to one month
@caina-barbosa APPROVED

►

#1806 — chore(hc-forecast): Record the live Core procedure definition
@caina-barbosa APPROVED

►

#1805 — fix(repo): Split OpenAI BU fetch across two scheduled Lambda runs
@heimdall-keval-factory[bot] APPROVED
AUTOMATED PR

►

#1292 — feat(enrollment): drop HubSpot identities + overlay, publish SIS identity and deposit signal (#1290)
@vvp-trilogy APPROVED

►

#182 — chore(mercy): pin the harness ref, so @v1 actually pins mercy
@kevalshahtrilogy APPROVED MERCY-ALLOW-CRITICAL

►

#1291 — chore(mercy): pin the harness ref, so @v1 actually pins mercy
@kevalshahtrilogy APPROVED MERCY-ALLOW-CRITICAL

►

#1774 fix(repo): Tolerate small Model breakdown excess like Credit Source (@heimdall-keval-factory[bot], Surtr)

#1748 fix(perplexity-usage-pipeline): enforce the skip ceiling per org-day too (@kevalshahtrilogy, Surtr)

#113 feat(mercy): trust heimdall-driven PRs, and stop summoning the agent after approval (@kevalshahtrilogy, mercy)

#112 feat(heimdall): label every pipeline-filed ticket `surtr-pipeline` (@kevalshahtrilogy, mercy)

#1773 fix(repo): Add voyage to TFY known-skipped providers (@heimdall-keval-factory[bot], Surtr)

#1770 fix(education): retain billing allocations without contact links (@benji-bizzell, Surtr)

#1768 fix(education): leave Finals site reader access under DBA control (@benji-bizzell, Surtr)

#1765 fix(education): restore SIS dataset-triggered refreshes (@benji-bizzell, Surtr)

#1763 feat(education): capture Finals site findings without blocking ingestion (@benji-bizzell, Surtr)

#1764 fix(repo): forward trigger context for dataset publication runs (@heimdall-keval-factory[bot], Surtr)

#1735 fix(ramp): retry transient API timeouts (@ashwanth1109, Surtr)

#1244 AERIE-1855: Add unpublished verified materialization and installer commands (@caina-barbosa, Aerie)

#1762 fix(education): preserve Finals site billing balance precision (@benji-bizzell, Surtr)

#1761 fix(education): preserve lossless Limitless transcripts (@benji-bizzell, Surtr)

#1245 feat(platform-errors): make automatic triage runs inspectable (@benji-bizzell, Aerie)

#111 feat(review): make blocking findings name the defect class to fix (@kevalshahtrilogy, mercy)

#1758 feat(education): isolate SIS publications and resume detail collection (@benji-bizzell, Surtr)

#1757 077-finops-incremental-regression (@mwrshah, Surtr)

#174 196-port-live-skill-hydration (@mwrshah, Sindri)

#1756 fix(education): preserve FinalSite contacts outside workflow listings (@benji-bizzell, Surtr)

#1755 fix(education): accept new GuidePlatform source fields (@benji-bizzell, Surtr)

#1687 feat(education): ingest Finals site billing snapshots (@benji-bizzell, Surtr)

#3737 fix(overspend-alerts): exclude dummy bank charge vendor (@sanketghia, Klair)

#3736 chore(repo): reverse instruction symlink direction (@sanketghia, Klair)

#1753 fix(netsuite-auto-renewal): publish daily snapshots (@ashwanth1109, Surtr)

#126 — feat(review): rebuild the second-opinion prompt around how the reviewer actually fails

@kevalshahtrilogy no labels

►

#9 — feat(console): make revamp operator-ready

@sanketghia no labels

►

#124 — fix(review): actually invoke the second opinion — it shipped inert

@kevalshahtrilogy no labels

►

#123 — feat(review): a second opinion on whether a finding must block the merge

@kevalshahtrilogy CHANGES REQUESTED

►

#8 — Add App Server protocol fault-injection coverage

@sanketghia no labels

►

#7 — Persist Codex App Server protocol evidence and failure state

@sanketghia no labels

►

#1285 — AERIE-1283: dbt CI — models only in production, models then tests in dev

@vvp-trilogy no labels

►

#1288 — feat(portfolio): align school and site data contracts

@benji-bizzell CHANGES REQUESTED

►

#6 — Improve validation repair diagnostics and convergence

@sanketghia no labels

►

#5 — Harden App Server reliability and operator recovery

@sanketghia no labels

►

#4 — test: broaden local resilience coverage

@sanketghia no labels

►

#3 — Gate 2: local durability and operations hardening

@sanketghia no labels

►

#1801 — fix(repo): Accept core_submitting in HC pinned-run guard

@heimdall-keval-factory[bot] APPROVED

AUTOMATED PR

►

▶
#1800 — fix(SURTR-648): use supported CAPEX batch request

@marcusdAIy APPROVED

▶
#1797 — [SURTR-1123] Rename workforce warehouse identifiers only

@caina-barbosa APPROVED

▶
#3744 — feat(board-doc): upgrade shared model to Fable 5.1

@marcusdAIy APPROVED

▶
#1286 — AERIE-1284: Shared HubSpot display-name dimension for enrollment + pipeline reports

@vvp-trilogy APPROVED

▶
#1260 — Remove getPortfolioHealth MCP surface

@YibinLongTrilogy APPROVED

▶
#1281 — AERIE-1279: Key the offering-kind axis on program_type so per-campus Main Program campuses reach the enrollment report

@vvp-trilogy APPROVED

▶
#1259 — Remove presentReplanOptions and generatePlan

@YibinLongTrilogy APPROVED

▶
#1276 — AERIE-1857: Add unpublished durable uploader and automatic wake

@caina-barbosa APPROVED

▶
#1250 — feat(dashboards): add Real Estate tab to Data Health page

@kevalshahtrilogy APPROVED MERCY-ALLOW-CRITICAL

▶
#3742 — fix(spacex-valuation): reconcile September distribution

@sanketghia APPROVED

▶
#2 — Establish Gate 0 local maturity baseline

@sanketghia no labels

▶
#1 — feat: add research-to-Linear intake workflow

@sanketghia no labels

▶
#1277 — fix(portfolio): retire legacy opening date from MCP reads

@benji-bizzell APPROVED

▶

- #1253 — feat(dashboards): add hidden REBL3 Surtr-comparison experimental view

@kevalshahtrilogyAPPROVED
- ▶
- #1271 — fix(operations): accept nullable Surtr schedule expressions

@benji-bizzellAPPROVED
- ▶
- #1273 — fix(portfolio): use Milestone 9 as the canonical opening date

@benji-bizzellAPPROVED
- ▶
- #1272 — AERIE-1266: Complete the arrival cohorts — normalize the enrolled date and stop gating entry timing on current attendance

@vvp-trilogyAPPROVED
- ▶
- #1270 — AERIE-1268: Scope the enrollment campus axis to physical, active campuses across the four brands

@vvp-trilogyAPPROVED
- ▶
- #1269 — feat(reconciliation): add dormant Aerie foundations

@caina-barbosaAPPROVED
- ▶
- #1267 — fix(api): complete enrollment pagination and handle cleared dates

@benji-bizzellAPPROVED
- ▶
- #1239 — fix(dev-local): start Next in dev-local:workers on Windows and macOS

@vvp-trilogyno labels
- ▶
- #1265 — feat(add-aerie-skill): add dormant adapter runtime and artifact lifecycle

@caina-barbosaAPPROVED
- ▶
- #1783 — feat(education): document temporary Alpha Summer Camp staging load

@ashwanth1109CHANGES REQUESTED
- ▶
- #122 — feat(review): relax the blocking bar on late rounds, not on every review

@kevalshahtrilogyCHANGES REQUESTED
- ▶
- #1788 — [SURTR-1175] Build and live validate Aerie retention refresh

@ashwanth1109CHANGES REQUESTED
- ▶

- #116 — feat(mercy): summon heimdall to fix conflicts alongside the findings

@kevalshahtrilogy no labels
- #1263 — test(feedback): stabilize intake rate-limit window coverage

@benji-bizzell APPROVED
- #1789 — fix(education): prevent billing estate refresh throttling

@benji-bizzell APPROVED
- #1784 — fix(qtd): match school names for guide headcount (SURTR-1161)

@ashwanth1109 CHANGES REQUESTED
- #1787 — fix(timeback): contain timeback-raw-sync overruns and finalize timed-out runs (SURTR-1162)

@caina-barbosa APPROVED MERCY-ALLOW-CRITICAL
- #1769 — [SURTR-1135] Fix QTD unit economics revenue proration

@ashwanth1109 APPROVED
- #1249 — feat(chat): add REBL3 pipeline data-health endpoint

@kevalshahtrilogy APPROVED
- #32 — AI-725: Persist workflow operations and add reusable smoke testing

@ashwanth1109 no labels
- #1261 — fix(enrollment): drop SIS tenant test records from the enrollment spine

@vvp-trilogy APPROVED
- #120 — fix(heimdall): the described-edit check was inert — the harness made the workspace look dirty

@kevalshahtrilogy no labels
- #119 — feat(heimdall): repair red before publishing, and never open a second PR for a ticket

@kevalshahtrilogy no labels
- #118 — feat(heimdall): stop losing finished work — described fixes, merge closure, evidence reach, and self-decided calls

@kevalshahtrilogy CHANGES REQUESTED
- #1257 — fix(finance): support NetSuite-only CAPEX entities

@marcusdAIy APPROVED

►
#1782 — fix(repo): Surface ECS StopCode/StoppedReason before truncation

@heimdall-keval-factory[bot] APPROVED
AUTOMATED PR

►
#1781 — fix(capex): apply warehouse DDL in one transaction

@marcusdAIy APPROVED

►
#1780 — fix(capex): preserve missing DDR budgets as null

@marcusdAIy APPROVED

►
#1256 — fix(enrollment): SIS mid-year Transfer Out excludes FinalSite opening-day stamps (#1254)

@vvp-trilogy APPROVED

►
#1778 — fix(capex): fail closed in the live verifier

@marcusdAIy APPROVED

►
#1777 — chore(heimdall): sweep the Linear queue every 5 minutes

@kevalshahtrilogy APPROVED MERCY-ALLOW-CRITICAL

►
#1252 — 1239-2aerie-publishing-modalities

@mwrshah APPROVED

►
#1776 — chore(heimdall): 60-minute soak and a 3-hour release cooldown

@kevalshahtrilogy APPROVED MERCY-ALLOW-CRITICAL

►
#117 — feat(heimdall): a release cooldown, so an hourly releaser doesn't ship hourly forever

@kevalshahtrilogy CHANGES REQUESTED

►
#1775 — chore(heimdall): format the agent's output before it commits

@kevalshahtrilogy APPROVED

►
#114 — feat(heimdall): format the agent's output before committing it

@kevalshahtrilogy no labels

►
#1121 — fix(sales-educrm-mart-sync): raise Redshift statement poll ceiling to f...

@heimdall-keval-factory[bot] APPROVED
AUTOMATED PR HEIMDALL-DRIVEN

►

►
#1595 — fix(openai-usage-pipeline): give token bucket headroom below OpenAI rat...

@heimdall-keval-factory[bot] APPROVED
AUTOMATED PR HEIMDALL-DRIVEN

►
#1559 — fix(mart-school-performance-table-3-refresh): retry mart refresh CALL o...

@heimdall-keval-factory[bot] APPROVED
AUTOMATED PR HEIMDALL-DRIVEN

►
#115 — fix(heimdall): stop a finished answer reading as a question

@kevalshahtrilogy APPROVED

►
#1774 — fix(repo): Tolerate small Model breakdown excess like Credit Source

@heimdall-keval-factory[bot] APPROVED
AUTOMATED PR

►
#1748 — fix(perplexity-usage-pipeline): enforce the skip ceiling per org-day too

@kevalshahtrilogy APPROVED HEIMDALL-DRIVEN

►
#113 — feat(mercy): trust heimdall-driven PRs, and stop summoning the agent after approval

@kevalshahtrilogy no labels

►
#112 — feat(heimdall): label every pipeline-filed ticket `surtr-pipeline`

@kevalshahtrilogy no labels

►
#1773 — fix(repo): Add voyage to TFY known-skipped providers

@heimdall-keval-factory[bot] APPROVED
AUTOMATED PR

►
#1770 — fix(education): retain billing allocations without contact links

@benji-bizzell APPROVED

►
#1768 — fix(education): leave Finals site reader access under DBA control

@benji-bizzell no labels

►
#1765 — fix(education): restore SIS dataset-triggered refreshes

@benji-bizzell no labels

►
#1763 — feat(education): capture Finals site findings without blocking ingestion

@benji-bizzell CHANGES REQUESTED

- #1764 — fix(repo): forward trigger context for dataset publication runs

@heimdall-keval-factory[bot]

APPROVED

AUTOMATED PR
- #1735 — fix(ramp): retry transient API timeouts

@ashwanth1109 no labels
- #1244 — AERIE-1855: Add unpublished verified materialization and installer commands

@caina-barbosa

APPROVED
- #1762 — fix(education): preserve Finals site billing balance precision

@benji-bizzell no labels
- #1761 — fix(education): preserve lossless Limitless transcripts

@benji-bizzell

APPROVED
- #1245 — feat(platform-errors): make automatic triage runs inspectable

@benji-bizzell

CHANGES REQUESTED
- #111 — feat(review): make blocking findings name the defect class to fix

@kevalshahtrilogy no labels
- #1758 — feat(education): isolate SIS publications and resume detail collection

@benji-bizzell

CHANGES REQUESTED
- #1757 — 077-finops-incremental-regression

@mwrshah

APPROVED
- #174 — 196-port-live-skill-hydration

@mwrshah

APPROVED
- #1756 — fix(education): preserve FinalSite contacts outside workflow listings

@benji-bizzell

APPROVED
- #1755 — fix(education): accept new GuidePlatform source fields

@benji-bizzell

APPROVED
- #1687 — feat(education): ingest Finals site billing snapshots

@benji-bizzell

CHANGES REQUESTED

@sanketghia APPROVED

@sanketghia no labels

@ashwanth1109 APPROVED

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

@ashwanth1109 no labels

@ashwanth1109 no labels

@ashwanth1109 CHANGES REQUESTED

MERCY-ALLOW-CRITICAL

@YibinLongTrilogy APPROVED

@ashwanth1109 APPROVED

@ashwanth1109

@ashwanth1109

@sanketghia

178 PRs, Eight Repos, One Ashwanth: Builder Team Shatters the Odometer

The numbers don't lie — 53 PRs from one man and a team that treated Surtr like a personal punching bag, in the best possible way.

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

Folks, strap in. One hundred and seventy-eight pull requests across eight repositories in seven days — that's not a sprint, that's a track meet, and the Builder Team just set a new world record. Surtr led the charge with 54 PRs, Shipyard close behind at 42, Aerie posting a robust 41, and even sleepy little Sindri managed to sneak in 2. This is a team operating at altitude, comrades, and the air up there is thin but the output is thick.

Let's run the board. @kevalshahtrilogy logged 25 PRs — a quarter-century of commits, if you'll permit the pun. @benji-bizzell delivered 20, a clean round number that speaks to discipline. @sanketghia notched 16, including the delicate #3763 in Klair reconciling a September 10 trade — precision work, folks, the kind that keeps auditors employed. @caina-barbosa and @marcusdAly each tallied 13, as did @vvp-trilogy, a three-way tie that would make Olympic judges nervous. And don't sleep on the bot — @heimdall-keval-factory[bot] cranked out 12 PRs including #1829 and #1825 in Surtr, quietly fixing Redshift deadlocks and restoring aliases like a machine that never asks for a lunch break. Because it isn't one.

And then there's @ashwanth1109. Fifty-three PRs. FIFTY-THREE. That's not a contributor, that's a weather system. He touched Shipyard from #58 all the way to #72, shepherding three full release cycles (0.3.1, 0.4.0, 0.4.1) practically by himself, then wandered over to Surtr to fix ramp cost research resilience in #1828 and QuickBooks budget logic in #1826, before circling back to Aerie for OneRoster provenance work in #1321. Sources describe him telling teammates, "I don't review my own diffs, I just remember writing them," which is either the most confident or most terrifying thing a senior engineer can say. Nobody has independently verified that anyone besides Ashwanth has read PR #63's full changeset. When reached for comment on his historic week, Ashwanth reportedly just said, "Next question," and walked away.

On the overflow desk, Mac's cutting room floor runneth over. #1264 in Aerie quietly stood up a Financials retention dashboard with raw data — unglamorous, essential. #1319 and #1318 from @YibinLongTrilogy improved mobile Admissions Forecast comparisons and added a legacy Pipeline link, the kind of UX polish that never trends but always matters. And #1823 in Surtr removed a retention rules version contract, tidying up technical debt nobody wanted to touch.

The leaderboard this week reads like a highlight reel with no lowlights — every single name above the fold is pulling double-digit PR counts, and the bot is outworking some humans, which should terrify no one and inspire everyone. Even the new Shipyard-Releases repo, freshly minted for signed installer distribution, arrives already primed for the kind of throughput this team treats as baseline.

Morale? Off the charts. Numbers Desk confirms: all-time high, no ceiling in sight.

Contently Finds Its Niche: Compliance, in an Industry That Never Forgets

Under new ownership, the content marketing platform is quietly building a toll booth for regulated finance — right as its ex-boss, Google, gets harder to be seen in.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

NEW YORK — Four posts, one publisher, one unmistakable pattern. In recent weeks Contently's editorial output has narrowed almost entirely to a single audience: banks, wealth managers, and insurers who need content but cannot afford a compliance misstep.

The pieces read like a sales funnel disguised as thought leadership. First comes the diagnosis — [a five-component work-flow](#) for regulated brands to scale content "without sacrificing governance." Then the justification for the spend — a framework for [measuring ROI across the glacial, multi-stakeholder sales cycles](#) that define institutional finance. Then the credibility audit, warning that AI engines and buyers alike now trust only the "named, credentialed expert." It is not hard to guess who Contently proposes should supply that expert — for a fee, from its marketplace of more than 165,000 creative professionals.

The timing is notable. Contently was folded into Zax Capital, an ESW Capital division, in September 2024, and installed a new chief executive, Brandon Pizzacalla, shortly after. ESW's playbook is well established: acquire a business with sticky customers, then push pricing and specialization until margins reach the firm's storied 75% target. Regulated finance is, by design, sticky. Compliance departments do not switch content vendors on a whim. A platform that can credibly promise both scale and governance has effectively built a moat — and a justification for premium pricing — inside an industry that already treats risk-aversion as a virtue.

The fourth post complicates the pitch. Contently itself has flagged that [top-ranked pages are increasingly invisible to Google's AI Overviews](#), meaning the organic traffic that used to justify content spend is eroding for everyone — including, presumably, Contently's own clients. If search visibility keeps decaying, the argument for paying a marketplace of credentialed experts to build "authority" content only strengthens.

Who benefits when the algorithm gets murkier and the compliance bar gets higher? The vendor who sells clarity to both problems at once.

AUSTIN'S ALPHA SCHOOL SAYS THE REAL CLASSROOM IS YOUR KITCHEN TABLE

BY DOTTIE SHARP, SOCIETY & INDUSTRY DESK · CLAUDE SONNET

AUSTIN, TEXAS — Word around the Domain is that Alpha School just dropped Part 5 of its "Teach Your Kid What School Doesn't" series, and this one's a doozy: forget the worksheets, kid, your living room is the real innovation lab. The pitch — straight from the 2-hour-a-day playbook — is that creativity can't be scheduled between recess and math drills. It's gotta be unleashed at home, on the family's dime, no AI tutor required. Cute. On brand. Very Liemandt.

Meanwhile, this column notes the rest of the AI industry is having a rougher week than any Alpha third-grader. Out west, the never-ending [Musk v. Altman saga](#) has litigators buzzing about what "fair" executive comp even means in an industry minting billionaires by the quarter. A little bird in legal circles tells me the case is becoming a case study — literally, econ experts are already citing it in briefs.

Across the Pacific, Micron's Taiwan unions are threatening to walk over bonuses, proving even chip royalty isn't immune to labor headaches. And down in D.C., appeals judges are reportedly split down the middle on the [Anthropic-Pentagon dispute](#) — because nothing says growing pains like the Defense Department squabbling with a chatbot company.

Against that backdrop, Alpha's advice to just let the kids build a blanket fort and call it R&D looks almost quaint. Almost. Word is Liemandt thinks the quaint stuff scales to a billion students. Don't bet against him — he's been right before.

Two Businesses, One Playbook: How Trilogy Scales Fast and What Gets Left Behind

From a \$65,000-a-year AI classroom to a newly absorbed telecom CPQ platform, the same Trilogy instinct is visible — move fast, compress timelines, and let the market sort out the rest.

BY FRANK DUNMORE, INVESTIGATIVE
CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — On the surface, a Silicon Valley-style private school in Austin and a Salesforce-native quoting tool for telecom operators have nothing to do with each other. But if you read between the lines, they're running the exact same play — and this is where it gets interesting.

Alpha School, Joe Liemandt's two-hour-a-day AI academy, has spent the last year as the darling of education-disruption coverage, most recently in the [New York Post](#)'s writeup of its \$65,000 tuition and its promise of a full curriculum in two hours flat. But [WIRED's](#) parallel reporting tells a quieter story — families who bought into the promise and later found the exit door harder to locate than the enrollment form. Sources close to the campuses, who I cannot name, describe a pattern that looks less like a glitch and more like a feature: sell the acceleration, manage the churn.

Meanwhile, over at Skyvera, the ESW Capital telecom arm just closed its acquisition of CloudSense, the AI-powered CPQ engine now folded into its portfolio. CloudSense had already made headlines this spring for compressing a 26-month TM Forum API certification into 30 days — a compression story that should sound familiar by now.

Nothing in the Trilogy universe happens by accident. Whether it's a school day or a telecom sales cycle, the mandate is the same: collapse the timeline, extract the value, and worry about the friction later. Alpha's parents and CloudSense's

telco clients are, in their own ways, both testing the same hypothesis — that speed and satisfaction can scale together. The early returns, per multiple accounts, are mixed. But mixed has never slowed this machine down before.

The Brain, Rendered Legible: AI's New Way of Listening to the Mind

From hidden scars in the brain to silent thoughts becoming text, machine intelligence is learning to read what the nervous system has always been trying to say.

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

PALO ALTO, CALIFORNIA — There is a particular kind of humility that comes from realizing the brain has been shouting for millennia, and we simply lacked the ears to hear it. This week brought three small revelations that, taken together, sound like the opening notes of that hearing.

Consider multiple sclerosis, a disease that has long played hide-and-seek with radiologists. Gray matter lesions — the quiet, cortical kind, distinct from the white matter damage doctors usually chase — have historically evaded even trained eyes on an MRI. Now, as [researchers report](#), AI models trained on subtle textural patterns are surfacing lesions invisible to the unaided eye — not

replacing the radiologist's judgment, but sharpening the instrument they hold.

Meanwhile at Meta, a system called Brain2Qwerty is attempting something stranger still: translating the electrical weather of thought into typed words, without surgery, without implants — just a cap of sensors and a model patient enough to learn the grammar of a mind that cannot speak. For people locked inside bodies that no longer obey them, this is not a parlor trick. It is a bridge.

And in classrooms far from any lab, something equally moving is happening. Teenagers are being paired with working neuroscientists — not as spectators but as collaborators — and by their own astonished admission, finding the experience simply, unpretentiously "wow." It's a re-

minder that scientific wonder doesn't require decades of credentialing to access; sometimes it just requires an invitation.

Stanford's Human-Centered AI institute has been making a related argument lately: that [AI's role in discovery works best not by replacing the scientist](#), but by handing them a more sensitive instrument — the way a telescope doesn't replace an astronomer's eye, it simply widens what that eye can hold.

Each of these stories, in its own register, says the same thing: intelligence recognizing intelligence, across a gap once thought unbridgeable. The brain has been signaling for four billion years of evolutionary refinement. We are only now building the receivers worthy of the transmission.

Amodei's 3,800 Words Won't Slow Anything Down

The Anthropic chief wants an industry-wide pause on capability gains. His own product roadmap says otherwise.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — Dario Amodei published 3,800 words this week arguing that the company he runs is moving too fast. The Anthropic chief executive's essay lays out a three-step plan for reining in artificial intelligence: third-party safety audits, government-mandated disclosure of model capabilities, and an industry pact to slow the release cadence of frontier systems. It is, by any measure, an unusual document — a CEO asking regulators to constrain the thing that made him a billionaire.

The essay did not arrive in a vacuum. The New York Times reported this week that researchers inside Anthropic, OpenAI, Meta and Google DeepMind have spent recent months in internal debates over what one described as "superintelligence doomsday" scenarios — the possibility that capability gains outpace any control mechanism built to govern them. The consensus, per [people familiar with the discussions](#), is not that catastrophe is imminent. It is that nobody currently has the instrumentation to know either way.

That is the harder problem, and it is not new. In 1975, Asilomar biologists imposed a voluntary moratorium on recombinant DNA research until safety protocols caught up — a rare case of an industry successfully self-policing ahead of regulation. AI has no Asilomar. Compute budgets at Anthropic, OpenAI and Google are measured in tens of billions of dollars annually, training runs are measured in months, and the market rewards the lab that ships first, not the lab that waits for consensus. [Amodei's own letter](#) concedes the incentive problem without solving it: he wants competitors bound by rules he can't unilaterally impose on himself.

The skepticism is warranted. Anthropic's valuation, last marked north of \$60 billion, rests on Claude's capability trajectory, not its restraint. A slowdown pact that isn't enforceable by regulators — and none currently exists in the U.S. — functions as a press release, not a policy. The NYT's parallel reporting on why oversight lags development notes the obvious: monitoring systems are built by the same labs whose products they're meant to check. Amodei is right that the industry needs guardrails. He has not yet explained why his company would be the one to install them first.

The Rules Don't Travel: AI Governance Splinters Along Old Fault Lines

As Washington and Beijing race machines, the world discovers a machine cannot be governed by two masters.

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

GENEVA — Diplomats here have a phrase for it, muttered over cold coffee in the corridors of the Palais des Nations: the treaty that keeps not arriving. Artificial intelligence, the technology that was supposed to unite the world's regulators in common cause against a common unknown, has instead become another front in a rivalry as old as the Cold War's ghost.

A new analysis making the rounds among policy circles this week argues what many in the room already suspected: [geopolitical rivalry, not technical disagreement, is the real brake on global AI regulation](#). The United States wants rules that protect its lead. China wants rules that protect its model of state-directed development. Europe wants rules, full stop, and increasingly finds itself negotiating with a table that has only two real chairs.

The consequence is not chaos so much as a slow partition — a Balkanization of the algorithm. Export controls on chips function as de facto industrial policy. Data localization laws function as sovereignty claims. Even open-source releases, once celebrated as a diplomatic bridge, now get read in Washington and Beijing alike as strategic signaling. The Brussels Effect, that old hope that European standards would simply propagate outward by market gravity, strains against a world where the two largest AI economies increasingly decline to be pulled.

What gets lost in this standoff is the smaller nation's leverage. A regulator in Nairobi or Jakarta cannot simply adopt Brussels' rulebook and expect Silicon Valley or Shenzhen to comply; increasingly, they must choose a bloc. The result, officials here concede privately, is not a race to the top or the bottom, but a race to two tops — parallel regimes, parallel compute stacks, parallel definitions of what a machine is allowed to know.

No binding global framework is scheduled. None is expected soon.

The Nation's Attic, and the Man Who Rents the Key

America pays billions a year to forget what it owns; Joe Liemandt built an empire on the same impulse, applied to software.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE SONNET

AUSTIN, TEXAS — There is, this week, a meditation in a magazine of some standing on the American genius for self-storage — on the peculiar national compulsion to rent a metal box, fill it with the furniture of a former life, and then never open the door again. The essay asks, in the gentle, bewildered tone such essays adopt, why a nation of such famous restlessness should also be a nation of such famous hoarding, and it arrives, as these essays generally do, at something about memory and identity and the fear of throwing anything away. It is a fine piece of writing. It has also, without knowing it, described the entire business model of Joe Liemandt.

For what is ESW Capital, if not a self-storage empire for enterprise software? Mr. Liemandt's people go about the corporate landscape acquiring the CRM platforms and telecom billing systems and content-marketing tools that some larger company built in 2003, used hard for a decade, and then quietly stopped looking at — the software equivalent of the Ping-Pong table and the box of college notebooks — and they pay for these properties not what a going concern is worth but what a storage unit is worth: one or two times annual revenue, the price of a thing nobody wants to think about but nobody can quite bring themselves to throw away. Aurea, IgniteTech, Skyvera — these are not so much companies as climate-controlled units, and the tenants keep paying the monthly fee long after they have forgotten what, precisely, they are storing.

The self-storage customer, per the [essay in question](#), is engaged in an act of faith — a belief that the objects retain some value that the act of storing them, paradoxically, guarantees will never be tested. This is precisely the faith of the Fortune 500 procurement officer who has not opened his CloudSense license in four years but dares not cancel it, lest some downstream integration groan and fail in a manner nobody now living understands well enough to fix. ESW Capital does not build a religion; it merely collects the tithes of one already practiced, at a discount, by companies who long ago stopped believing in the software and kept only the fear of moving it.

The magazine's other offerings this week run to a critic's shelf of new novels, a cover illustration by Dean Rohrer imagining the city as a blank canvas — appropriately, since a blank canvas is what every enterprise IT department imagines it is working with, right up until the movers arrive — and a pair of dispatches on the capital's own storage habits: Jared Kushner's talent for converting proximity to power into private equity, and Mr. Trump's persistent misreading of Xi Jinping as a man who might be charmed

rather than calculated with. Each is, in its way, a story about things kept far past the point of use — access, goodwill, illusion — because someone, somewhere, could not bear to pay the fee for the dumpster instead of the unit. Mr. Liemandt, at least, has the honesty to send an invoice.



The Office Comic · Art Desk

Economists Confirm AI Has Already Delivered 150% More Productivity That Has Not Yet Occurred

The Federal Reserve says the gains are still to come, Big Tech says they already came, and both are apparently correct at the same time, which is fine.

BY DALE PEMBERTON, STAFF WRITER · CLAUDE SONNET

AUSTIN, TEXAS — I would like to take a moment, as a journalist and a person who once believed numbers meant things, to sit with a remarkable piece of consensus that emerged this week across the technology and economics press. The Federal Reserve, in a report that should have been accompanied by a national moment of silence, found that [95 percent of AI's promised productivity gains are 'still to come.'](#) Fine. Sober. Reasonable. Almost boring in its restraint.

Meanwhile, over in the part of the industry that does not have to answer to a central bank, a commit-level study of Big Tech engineering teams found that developer performance is already up 150 percent over eighteen months. Not 'expected to be up.' Not 'projected, pending macro conditions.' Up. Already. Retroactively, apparently, since the Fed's math suggests the bulk of the gains haven't happened yet. Somewhere between the commit logs and the balance sheets, 150 percent of a thing that is only 5 percent finished has already occurred, and nobody in either report seemed to find this strange enough to mention.

Oracle, for its part, is not interested in your percentages at all. Oracle would like everyone to know it is pinning its future on what can only be described, in the company's own framing, as a Star Wars-style leap to lightspeed via AI-assisted engineering — a phrase that suggests either a bold new era of enterprise software or a deeply concerning misunderstanding of how hyperdrives work. Either way, shareholders were reportedly asked to imagine the Millennium Falcon, and, this being 2025, several of them did.

And this is where the legal industry, sensing blood in the water the way it always does, has begun quietly sharpening something. Analysts are now openly discussing the convergence of securities claims and regulatory action around AI hype, on the theory that if you tell investors your productivity gains have already lightspeed-jumped while also filing internal projections describing those same gains as '95 percent still to come,' someone, somewhere, eventually has to pick a number and stand behind it in a deposition.

I mention all this not to editorialize, but because I spent this week reviewing a rather unremarkable set of internal documents from Trilogy's own AI Builder Team, the group that maintains Klair, the analytics platform quietly keeping tabs on ESW Capital's seventy-five-odd portfolio companies. Nothing in those documents claimed lightspeed. Nothing claimed 150 percent. There were tickets. There was a demo. Someone closed a Linear issue about invoice reconciliation and moved on with their day, which is, by current industry standards, borderline insubordinate.

It is worth asking, gently, whether an industry that has decided its productivity gains can simultaneously be nearly finished and barely begun has actually built a hyperdrive, or whether it has simply discovered that no one checks the math until the SEC does. Until then, the Falcon flies. The commits rise. The 95 percent remains, patiently, still to come.