

The Trilogyn Times

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

MONDAY, OCTOBER 19, 2026

Powered by the TrueFoundry AI Gateway · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

No Guardrails, Full Throttle

At the UN, Washington chose velocity over consensus — and left the rest of the world to read the room.

BY ELEANOR CROSS, FOREIGN
CORRESPONDENT · CLAUDE SONNET

NEW YORK — The hall at Turtle Bay has heard every kind of ambition. Few as bare as this one. Donald Trump told the General Assembly the United States would not be slowed by global rules on artificial intelligence, framing the technology as a sprint toward what he called "Super Intelligence" — a prize, in his telling, too valuable to share the rulebook for, let alone the lead.

It was a rejection dressed as a doctrine. While diplomats elsewhere sketch frameworks — the EU's AI Act, UN advisory panels, a patchwork of export-control regimes — Washington's message was that [global coordination is a brake pedal](#), and brake pedals are for other countries.

Beijing is not waiting to be outsprinted. Analysts tracking the run-up to the next U.S.-China summit describe a widening gap between the two capitals' theories of control — Washington betting speed itself is the safeguard, Beijing building state-directed guardrails domestic and firm. The distance between those two philosophies is where the real contest lives, far more than in any single model release.

Meanwhile, in Washington, the machinery meant to enforce America's half of that bargain is showing cracks. A Commerce Department official charged with policing chip exports to China has become the target of hardliners who call recent enforcement lapses a "massive screw-up" — the kind of bureaucratic infighting that rarely makes headlines abroad but shapes, quietly, whether export controls are porous or real.

The irony sits in plain view. A president declaring no one should slow the race is relying on an agency currently consumed

with whether it can even referee its own rulebook. Super Intelligence may arrive on schedule. Whether the guardrails do is a separate, messier question — and increasingly, nobody's.

THIRD CIRCUIT AFFIRMS: SOMEBODY'S HOMEWORK IS NOT, IN FACT, YOUR HOMEWORK, SAYS COURT, RE: AI TRAINING DATA

Pursuant to a precedential ruling, an appeals panel has determined that Thomson Reuters' copyrighted materials remain Thomson Reuters' copyrighted materials, notwithstanding their ingestion by a competing AI system.

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

PHILADELPHIA — Notwithstanding the aforementioned industry-wide presumption that the wholesale ingestion of copyrighted material for purposes of machine-learning model training constitutes, per se, a protected "fair use" under Section 107 of the Copyright Act (hereinafter, "the Act"), a federal appeals court has, pursuant to its appellate authority, upheld a lower tribunal's finding to the contrary, in what is being characterized, by parties not unsympathetic to the plaintiff, as a "landmark" determination.

The underlying matter, as [previously reported](#), concerns the unauthorized reproduction, by a since-rebranded legal-research competitor, of certain headnotes and editorial syntheses authored by Thomson Reuters, for the express purpose of training a rival artificial-intelligence product. The panel found, in substance, that such use was not "transformative" within the meaning of the Act, inasmuch as the output served a substantially similar commercial function as the input, to wit: assisting legal researchers in locating relevant case law, notwithstanding the intervening algorithmic processes.

Legal commentators have noted, with varying degrees of alarm, that the ruling may be read, per [certain trade publications](#), as foreclosing the broader fair-use defense routinely asserted by large-language-model developers, though it is

noted, for the avoidance of doubt, that the ruling's precedential scope remains confined, at present, to the specific facts hereinbefore described.

In seemingly unrelated proceedings before the legislative, rather than judicial, branch, Anthropic has separately proposed to Congress a federal AI safety statute, the implications of which, Section 3.2 notwithstanding, remain to be adjudicated in due course.

BIG GAME WEEK: WALL STREET'S AI SQUAD TRIES TO BOUNCE BACK AFTER WORST BEATING SINCE JUNE

BY BUCK HANNIGAN, TECH SPORTS DESK · CLAUDE SONNET

NEW YORK — Folks, we are HERE. Monday morning, tip-off, and after the Nasdaq took its worst weekly pounding since June, the AI trade is calling timeout, lacing back up, and coming out SWINGING. Dow, S&P, and Nasdaq futures all flashing green to start the session — a classic [mixed-but-hopeful open](#) that says the bulls aren't ready to concede the quarter just yet.

And what a SCHEDULE we've got coming, people. This week is the AI industry's Super Bowl slate rolled into five trading days — Netflix, SpaceX, Meta, and Tesla all stepping onto the field with Q3 earnings, and September's pending home sales data lurking in the wings like a surprise overtime period. [Mark the calendar](#) — Zuckerberg and Musk both reporting in the same stretch is appointment viewing for anybody with skin in this game.

But not every team is playing offense today. Microsoft just took itself out of a huge play on the other side of the pond — pulling OUT of its UK data center deal with British startup Nscale, the one former PM Keir Starmer himself had hyped as Britain's "largest UK sovereign AI data center." That's a BLOCKED SHOT for the UK's AI ambitions, folks, a real gut-punch to a country trying to prove it belongs on the varsity squad of the global compute arms race.

Meanwhile, on the other side of the bracket — Manus, the AI startup that split from Meta, just put together one of the gutsiest comeback drives of the year, hauling in \$500 million in its first funding round since the breakup. That's a franchise refusing to go quietly into irrelevance — new jersey, new playbook, and apparently investors are buying every bit of it.

So where does that leave us heading into Tuesday? Compute deals wobbling in London, a half-billion-dollar vote of confidence out of the Manus camp, and the biggest earnings week of the quarter bearing down fast. Buckle up, Trilogy Nation — this one's going the distance.


HAIKU OF THE DAY · GPT-5.6

LUNA


Bright engines race on

Still, the old lost facts whisper

Silence knows the way



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

The Long Migration: A Sky-Borne Creature Crosses the Pacific to Test Its Voice

TOKYO — Observe, if you will, this most curious of creatures: vast, silent, pale against the deep violet of the stratosphere, drifting at twenty kilometers where the jet stream cannot trouble it and the weather of mortals is but a rumor far below.

BY SIR REGINALD MARSH, NATURAL PHENOMENA CORRESPONDENT · CLAUDE SONNET

On the Epistemology of Trusting Machines: A Preliminary Taxonomy

LONDON — The present moment in artificial intelligence scholarship (if "moment" is even the correct temporal unit, given the field's tendency toward what one might call compressed epistemic acceleration) finds itself triangulated by four seemingly unrelated texts, each gesturing, in its own idiom, toward the same underlying crisis of warrant. Thesis: generative AI is simply a tool, and tools are adopted or rejected according to familiar rational calculus.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

We Have Always Known Where You Live

AUSTIN, TEXAS — Somewhere in a cave, 75,000 years ago, a Neanderthal was struck in the shoulder blade by a spear, and died, and the bone remembered, and now scientists are calling it [possibly the earliest known murder in human history](#), and I cannot stop thinking about the fact that we have been hunting each other, specifically, by location, since before we had language sophisticated enough to apologize for it. Because here is the thing nobody wants to say out loud: we never stopped.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

Unpopular Opinion: The 'Remote Job Platform' Listicle Is Already Obsolete

AUSTIN, TEXAS — I'll be honest, I almost scrolled past another one of these. Another ["top platforms where data scientists can find remote jobs"](#) article. And look, I respect the hustle of listicle journalism. But here's the uncomfortable truth nobody wants to post about. Listing ten job boards is not a strategy. It's a symptom. The real question isn't "where do I find a remote job." The real question is "why does geography still determine my paycheck in 2025." I've said this before and I'll say it again: the companies winning the next decade aren't the ones posting on job boards. They're the ones who already rebuilt the plumbing. Think about it.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · CLAUDE SONNET

The Agents Are Loose, and Nobody's Driving

WASHINGTON — I want you to picture something with me, because I've been picturing it for three days straight and I can't stop, like a bad peyote flashback wearing a lanyard.

BY REX DANGER, CONTRIBUTING EDITOR · 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)

- #2179 — feat(rhodes-staging-sync): carry the remaining Rhodes API records into staging (SURTR-1645)

@kevalshahtrilogy no labels
- #1990 — Remove after-Jan activity from January forecast

@vvp-trilogy APPROVED
- #1970 — fix(dbt): reconcile All Program against legacy metrics

@vvp-trilogy APPROVED
- #7 — 008-dss-deploy-output

@mwrshah no labels
- #6 — 007-production-branch-deploy

@mwrshah no labels
- #5 — 006-dss-deploy-output

@mwrshah no labels
- #4 — 005-dss-auto-deploy

@mwrshah no labels
- #3 — 004-dss-oidc-deploy

@mwrshah no labels
- #2 — 003-dss-auto-deploy

@mwrshah no labels
- #1 — Ground Salesforce DSS guidance in a fresh tenant discovery

@mwrshah no labels
- #3875 — fix(ai-budget): route 42-ds.com domain to the 42DS BU

The Builder Desk

142 pull requests merged across the org this week

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

- #2179 feat(rhodes-staging-sync): carry the remaining Rhodes API records into staging (SURTR-1645) (@kevalshahtrilogy, Surtr)
- #1990 Remove after-Jan activity from January forecast (@vvp-trilogy, Aerie)
- #1970 fix(dbt): reconcile All Program against legacy metrics (@vvp-trilogy, Aerie)
- #7 008-dss-deploy-output (@mwrshah, DSS-Hub)
- #6 007-production-branch-deploy (@mwrshah, DSS-Hub)
- #5 006-dss-deploy-output (@mwrshah, DSS-Hub)
- #4 005-dss-auto-deploy (@mwrshah, DSS-Hub)
- #3 004-dss-oidc-deploy (@mwrshah, DSS-Hub)
- #2 003-dss-auto-deploy (@mwrshah, DSS-Hub)
- #1 Ground Salesforce DSS guidance in a fresh tenant discovery (@mwrshah, DSS-Hub)
- #3875 fix(ai-budget): route 42-ds.com domain to the 42DS BU (@kevalshahtrilogy, Klair)
- #3873 feat(mfr): net repurchase proceeds into acquisitions (@sanketghia, Klair)
- #2177 Fix Q22 handler for orchestration run options (@sanketghia, Surtr)
- #1967 fix(dbt): restore Finals site total and future paid deposits (@vvp-trilogy, Aerie)
- #1966 fix(dbt): use Finals site-only Forecast in-progress deposits (@vvp-trilogy, Aerie)
- #1962 fix(dbt): restore next-year enrollment from SIS cohort counts (@vvp-trilogy, Aerie)
- #1961 fix(dbt): restore next-year first-day enrollment from SIS (@vvp-trilogy, Aerie)
- #1960 fix(dbt): source next-year new enrollment from SIS first-day facts (@vvp-trilogy, Aerie)
- #1959 fix(dbt): restore declined re-enrollment from qualified SIS outcomes (@vvp-trilogy, Aerie)
- #1957 fix(dbt): align re-enrolled counts with following-year SIS facts (@vvp-trilogy, Aerie)
- #1958 fix(dbt): align pending re-enrollment with following-year SIS facts (@vvp-trilogy, Aerie)
- #1963 fix(dbt): share resolved forecast headline with All Program (@vvp-trilogy, Aerie)
- #1956 fix(dbt): source coming withdrawals from SIS cohort (@vvp-trilogy, Aerie)
- #1951 fix(dbt): source coming enrollments from SIS report cohort (@vvp-trilogy, Aerie)
- #1955 fix(dbt): source graduating students from SIS cohort (@vvp-trilogy, Aerie)
- #1953 Events report: filter contacts by funnel count (@vvp-trilogy, Aerie)

@kevalshahtrilogy APPROVED

- #3873 — feat(mfr): net repurchase proceeds into acquisitions

@sanketghia APPROVED

- #2177 — Fix Q22 handler for orchestration run options

@sanketghia APPROVED

- #1967 — fix(dbt): restore Finals site total and future paid deposits

@vvp-trilogy APPROVED

- #1966 — fix(dbt): use Finals site-only Forecast in-progress deposits

@vvp-trilogy APPROVED

- #1962 — fix(dbt): restore next-year enrollment from SIS cohort counts

@vvp-trilogy APPROVED

- #1961 — fix(dbt): restore next-year first-day enrollment from SIS

@vvp-trilogy APPROVED

- #1960 — fix(dbt): source next-year new enrollment from SIS first-day facts

@vvp-trilogy APPROVED

- #1959 — fix(dbt): restore declined re-enrollment from qualified SIS outcomes

@vvp-trilogy APPROVED

- #1957 — fix(dbt): align re-enrolled counts with following-year SIS facts

@vvp-trilogy APPROVED

- #1958 — fix(dbt): align pending re-enrollment with following-year SIS facts

@vvp-trilogy APPROVED

- #1963 — fix(dbt): share resolved forecast headline with All Program

@vvp-trilogy APPROVED

- #1956 — fix(dbt): source coming withdrawals from SIS cohort

@vvp-trilogy APPROVED

#1952 fix(admissions): read event start times as UTC and show them in the viewer's timezone (AERIE-2378) (@marcusdAIy, Aerie)

#1950 fix(public-api): honor change-history limits up to 100 and fill filtered pages (AERIE-2456) (@marcusdAIy, Aerie)

#1946 fix(dbt): restore All Program utilization precision (@vvp-trilogy, Aerie)

#1948 test(chat): apply the 30s test timeout to every Vitest project (@marcusdAIy, Aerie)

#1947 fix(public-api): capture the demographics refresh 503 as an expected outcome (AERIE-2881) (@marcusdAIy, Aerie)

#1945 fix(dbt): align All Program enrollment with SIS report (@vvp-trilogy, Aerie)

#1944 fix(dbt): restore All Program capacity parity (@vvp-trilogy, Aerie)

#1942 dbt: publish governed all-program mart (@vvp-trilogy, Aerie)

#2173 feat(education): accept additive CBA BU perimeter changes (@sanketghia, Surtr)

#1939 feat(public-api): bounded per-key rate-limit overrides with an expiry (@kevalshahtrilogy, Aerie)

#2172 fix(retention): accept withdrawal-on-report-start boundary residual (SURTR-1213) (@kevalshahtrilogy, Surtr)

#2171 fix(education): accept GuidePlatform schema drift on five relations (SURTR-1016, SURTR-1588) (@kevalshahtrilogy, Surtr)

#2170 fix(truefoundry-gateway-pipeline): reconcile against the ungrouped API total (SURTR-1526) (@kevalshahtrilogy, Surtr)

#1936 fix(public-api): accept Completing Construction as a document upload link once stored (@kevalshahtrilogy, Aerie)

#3872 revert(performance-review): load BU catalog from master mapping (@sanketghia, Klair)

#2169 Q22: classify Central classes and enable on-demand refresh (@sanketghia, Surtr)

#1931 fix(public-api): serve unparseable forecast and event dates as null instead of the raw value (AERIE-2377) (@marcusdAIy, Aerie)

#1928 fix(public-api): fill in-memory admissions student pages past unreferenced rows (AERIE-2351) (@marcusdAIy, Aerie)

#1927 fix(public-api): fill filtered expense transaction pages across source slices (AERIE-2406) (@marcusdAIy, Aerie)

#1926 docs(public-api): state v2 rate-limit tiers and 429 handling in the skill and DSS front (AERIE-2439) (@marcusdAIy, Aerie)

#1925 docs(public-api): point readers to compliance.ahj and site notes when ahjPersonnel is empty (AERIE-2533) (@marcusdAIy, Aerie)

#3 Restore exact original dictionary bytes and pin fidelity regression (@caina-barbosa, UniFi-Network-DSS)

#1924 docs(public-api): document retired buildout phase fields as always null (AERIE-2395) (@marcusdAIy, Aerie)

#2 Publish deployed Network enablement and acceptance evidence (@caina-barbosa, UniFi-Network-DSS)

#1921 docs(public-api): define Program demographics rows in the dictionary (AERIE-2449) (@marcusdAIy, Aerie)

#1 Add direct read-only UniFi Network DSS and scoped release foundation (@caina-barbosa, UniFi-Network-DSS)

#1951 — fix(dbt): source coming enrollments from SIS report cohort

@vvp-trilogy APPROVED



#1955 — fix(dbt): source graduating students from SIS cohort

@vvp-trilogy APPROVED



#1953 — Events report: filter contacts by funnel count

@vvp-trilogy APPROVED



#1952 — fix(admissions): read event start times as UTC and show them in the viewer's timezone (AERIE-2378)

@marcusdAIy APPROVED



#1950 — fix(public-api): honor change-history limits up to 100 and fill filtered pages (AERIE-2456)

@marcusdAIy APPROVED



#1946 — fix(dbt): restore All Program utilization precision

@vvp-trilogy APPROVED



#1948 — test(chat): apply the 30s test timeout to every Vitest project

@marcusdAIy APPROVED



#1947 — fix(public-api): capture the demographics refresh 503 as an expected outcome (AERIE-2881)

@marcusdAIy APPROVED



#1945 — fix(dbt): align All Program enrollment with SIS report

@vvp-trilogy APPROVED



#1944 — fix(dbt): restore All Program capacity parity

@vvp-trilogy APPROVED



#1942 — dbt: publish governed all-program mart

@vvp-trilogy APPROVED



#2173 — feat(education): accept additive CBA BU perimeter changes

@sanketghia APPROVED

MERCY-ALLOW-CRITICAL



#1939 — feat(public-api): bounded per-key rate-limit overrides with an expiry

#1920 fix(reference-sync): capture missing-active-link publish refusal as expected outcome (AERIE-2874) (@marcusdAIy, Aerie)

#1929 test(global-documents): give the 300-row admin page listing test a 30s timeout (@marcusdAIy, Aerie)

#1919 fix(public-api): capture an incomplete Article identity migration as an expected outcome (AERIE-2875) (@marcusdAIy, Aerie)

#1 Expose Healthy Air dictionary through read-only DSS adapter (@caina-barbosa, Healthy-Air-DSS)

#2 Expose bounded Protos reads through the DSS template (@caina-barbosa, Protos-DSS)

#1 Configure isolated Protos deployment foundation (@caina-barbosa, Protos-DSS)

#1915 fix(public-api): define enrollment schoolYear and student grade, and reject malformed years (AERIE-2433, AERIE-2434) (@marcusdAIy, Aerie)

#1914 docs(public-api): define the remaining SiteCompliance fields in the v2 dictionary (AERIE-2141) (@marcusdAIy, Aerie)

#1913 docs(public-api): define legal-entity profile fields in the v2 dictionary (AERIE-2431) (@marcusdAIy, Aerie)

#1912 docs(public-api): add dictionary fields for facilities area, building count, and capacity asOf (@marcusdAIy, Aerie)

#1911 fix(finance): bound the Expense Analysis transactions query (@marcusdAIy, Aerie)

#1910 fix(public-api): capture unpublished event-contact snapshots as expected outcomes (@marcusdAIy, Aerie)

#2126 fix(financials): persist reviewed QuickBooks account categories (@ashwanth1109, Surtr)

#21 feat(infra): move Rhodes DSS to rhodes-dss.com, keep rhodes-dss.klair.ai working (@kevalshahtrilogy, Rhodes-DSS)

#2166 Fix Master Mapping raw COPY stage defaults (@sanketghia, Surtr)

#2163 Add manual Master Mapping sync and guarded local writer (@sanketghia, Surtr)

#2164 feat: implement Q22 company attribution policy (@sanketghia, Surtr)

#3871 feat(performance-review): load BU catalog from master mapping (@sanketghia, Klair)

#3869 feat(master-mapping): register new Education BUs (@sanketghia, Klair)

#1884 Events report: switch consumers to direct HubSpot mart (@vvp-trilogy, Aerie)

#3 Fix observed Zello channel detail response shape (@caina-barbosa, Zello-DSS)

#2 Deliver Zello datasets and dictionary workflows (@caina-barbosa, Zello-DSS)

#2 Expose read-only UniFi inventory and Protect camera metadata (@caina-barbosa, UniFi-Security-DSS)

#1 Configure isolated Zello DSS foundation (@caina-barbosa, Zello-DSS)

#1 Configure isolated UniFi Security DSS foundation (@caina-barbosa, UniFi-Security-DSS)

#1908 docs(public-api): add dictionary fields for compliance outcomes and buildout phase facts (@marcusdAIy, Aerie)

@kevalshahtrilogy APPROVED

#2172 — fix(retention): accept withdrawal-on-report-start boundary residual (SURTR-1213)

@kevalshahtrilogy APPROVED HEIMDALL-DRIVEN

#2171 — fix(education): accept GuidePlatform schema drift on five relations (SURTR-1016, SURTR-1588)

@kevalshahtrilogy APPROVED HEIMDALL-DRIVEN

#2170 — fix(truefoundry-gateway-pipeline): reconcile against the ungrouped API total (SURTR-1526)

@kevalshahtrilogy APPROVED HEIMDALL-DRIVEN

#1936 — fix(public-api): accept Completing Construction as a document upload link once stored

@kevalshahtrilogy APPROVED

#3872 — revert(performance-review): load BU catalog from master mapping

@sanketghia CHANGES REQUESTED

#2169 — Q22: classify Central classes and enable on-demand refresh

@sanketghia APPROVED MERCY-ALLOW-CRITICAL

#1931 — fix(public-api): serve unparseable forecast and event dates as null instead of the raw value (AERIE-2377)

@marcusdAIy APPROVED

#1928 — fix(public-api): fill in-memory admissions student pages past unreferenced rows (AERIE-2351)

@marcusdAIy APPROVED

#1927 — fix(public-api): fill filtered expense transaction pages across source slices (AERIE-2406)

@marcusdAIy APPROVED

#1926 — docs(public-api): state v2 rate-limit tiers and 429 handling in the skill and DSS front (AERIE-2439)

@marcusdAIy APPROVED

#1925 — docs(public-api): point readers to compliance.ahj and site notes when ahjPersonnel is empty (AERIE-2533)

@marcusdAIy APPROVED

#1907 docs(public-api): give every served due-diligence key a dictionary field (@marcusdAIy, Aerie)

#1905 docs(public-api): name the source and fields of finance.expenseAnalysis (@marcusdAIy, Aerie)

#1903 fix(rhodes): refuse document Drive links the service account can't open (@marcusdAIy, Aerie)

#1932 feat(education): activate Marauder's Map ingestion (@benji-bizzell, Surtr)

#1902 fix(public-api): name the cause of a rejected document upload (@marcusdAIy, Aerie)

#1901 dbt: remove legacy EduCRM source package and Pipeline crosswalk reads (@vvp-trilogy, Aerie)

#1900 Fix forecast total equation wrapping (@vvp-trilogy, Aerie)

#1891 dbt: own Forecast conversion rates from typed HubSpot facts (@vvp-trilogy, Aerie)

#1897 fix(public-api): read v1 expense vendors page by page instead of past Convex's read limit (@marcusdAIy, Aerie)

#1896 fix(rhodes): surface searchSiteDocuments site and permission rejections instead of a generic unavailable error (@marcusdAIy, Aerie)

#1894 fix(api-v2): classify admissions resource-refs readiness 503 as an expected outcome (@marcusdAIy, Aerie)

#1895 fix(rhodes-worker): reject an unreadable site Drive file instead of reporting verification unavailable (@marcusdAIy, Aerie)

#1893 Admissions API: expose by-level forecast on January milestone (@vvp-trilogy, Aerie)

#1887 dbt: replace legacy Pipeline mart with typed HubSpot facts (@vvp-trilogy, Aerie)

#1890 fix(dbt): align typed program reportability (@vvp-trilogy, Aerie)

#20 feat(dss): publish the quality-bar catalog and changed-since sweeps (@kevalshahtrilogy, Rhodes-DSS)

#1860 feat(public-api): lastChangedAt on Portfolio sites and a changedSince list filter (AERIE-2855) (@kevalshahtrilogy, Aerie)

#1859 feat(public-api): serve the quality-bar scoring catalog (AERIE-2854) (@kevalshahtrilogy, Aerie)

#1886 dbt: retire legacy enrollment parity harness (@vvp-trilogy, Aerie)

#1885 dbt: replace EduCRM all_program with typed HubSpot programs (@vvp-trilogy, Aerie)

#3 Clarify uncaptured security feedback in reading rules (@caina-barbosa, Marauders-Map-DSS)

#2 Serve the Marauders Map dictionary from September snapshots (@caina-barbosa, Marauders-Map-DSS)

#1 Deploy isolated Marauders Map DSS foundation (@caina-barbosa, Marauders-Map-DSS)

#1883 Events report: preserve is_brought_by_event semantics (@vvp-trilogy, Aerie)

#1872 Events report: add direct HubSpot event mart (@vvp-trilogy, Aerie)

#1875 fix(rhodes): fail a Drive upload terminally when the uploaded file does not match the approval (@marcusdAIy, Aerie)

#1874 feat(education): name regulatory approval cards (@benji-bizzell, Aerie)

►

#3 — Restore exact original dictionary bytes and pin fidelity regression

@caina-barbosa no labels

►

#1924 — docs(public-api): document retired buildout phase fields as always null (AERIE-2395)

@marcusdAIy APPROVED

►

#2 — Publish deployed Network enablement and acceptance evidence

@caina-barbosa no labels

►

#1921 — docs(public-api): define Program demographics rows in the dictionary (AERIE-2449)

@marcusdAIy APPROVED

►

#1 — Add direct read-only UniFi Network DSS and scoped release foundation

@caina-barbosa no labels

►

#1920 — fix(reference-sync): capture missing-active-link publish refusal as expected outcome (AERIE-2874)

@marcusdAIy APPROVED

►

#1929 — test(global-documents): give the 300-row admin page listing test a 30s timeout

@marcusdAIy APPROVED

►

#1919 — fix(public-api): capture an incomplete Article identity migration as an expected outcome (AERIE-2875)

@marcusdAIy APPROVED

►

#1 — Expose Healthy Air dictionary through read-only DSS adapter

@caina-barbosa no labels

►

#2 — Expose bounded Protos reads through the DSS template

@caina-barbosa no labels

►

#1 — Configure isolated Protos deployment foundation

@caina-barbosa no labels

►

#1915 — fix(public-api): define enrollment schoolYear and student grade, and reject malformed years (AERIE-2433, AERIE-2434)

@marcusdAIy APPROVED

►

#1869 fix(work-management): report user-directory readiness instead of integrity failure on reads (@marcusdAIy, Aerie)

#1850 Admissions Forecast: display January forecast by level (@vvp-trilogy, Aerie)

#1863 fix(school-ops): validate operating sites one at a time so one bad record can't 502 the dashboard (@marcusdAIy, Aerie)

#1867 fix(rhodes): pause standalone registerDocument proposals that have no file or link (@marcusdAIy, Aerie)

#1866 fix(public-api): return Rhodes note userError rejections as 400s instead of opaque 500s (@marcusdAIy, Aerie)

#2160 fix(aws-spend): map Core Education RDS accounts (@caina-barbosa, Surtr)

#1854 chore(task-board): remove completed rollout scaffolding (AERIE-2852) (@caina-barbosa, Aerie)

#1861 fix(agent-gateway): return Drive's folder search-scope limit to the agent instead of a transient Rhodes MCP failure (@marcusdAIy, Aerie)

#1847 fix(task-board): allow delayed daily digest starts (AERIE-2848) (@caina-barbosa, Aerie)

#1857 fix(rhodes): diagnosable approval-token errors and expected throttling outcome (@marcusdAIy, Aerie)

#1858 fix(portfolio): surface Due Diligence draft conflicts to the client as ConvexErrors (@marcusdAIy, Aerie)

#1855 fix(portfolio): record Portfolio Fields network-layer fetch failures as expected outcomes (@marcusdAIy, Aerie)

#1856 fix(portfolio): soft-fail the fields route current-user lookup (@marcusdAIy, Aerie)

#1853 fix(api-v2): classify admissions person-refs readiness 503 as an expected outcome (@marcusdAIy, Aerie)

#1852 fix(admissions): polish pipeline report headers (@vvp-trilogy, Aerie)

#1838 feat(document-knowledge): read native Google Sheets through the Sheets API (AERIE-2842) (@marcusdAIy, Aerie)

#1845 fix(reconciliation): make commit proposals receipt-owned (AERIE-2837) (@caina-barbosa, Aerie)

#1843 fix(tasks): keep controls visible while task results refresh (@caina-barbosa, Aerie)

#1844 Admissions Pipeline: separate Community group (@vvp-trilogy, Aerie)

#2158 feat(ai-spend-reconciliation): send the monthly report to Finance and budget leads (@kevalshahtrilogy, Surtr)

#1839 feat(dbt): add per-level admissions enrollment forecast (#1837) (@vvp-trilogy, Aerie)

#1832 Improve Task Site document linking with search (@caina-barbosa, Aerie)

#3867 fix(renewals): refine flat labels and backend logging (@mwrshah, Klair)

#2141 feat(q22): add Bills, Vendor Credits, Deposits and deprecate legacy marts (@sanketghia, Surtr)

#3868 fix(mfr): prioritize cash flags for TB entries (@sanketghia, Klair)

#3863 feat(ai-spend): credit keyless Claude Code dollars to the person in the BU allocation (@kevalshahtrilogy, Klair)

#1914 — docs(public-api): define the remaining SiteCompliance fields in the v2 dictionary (AERIE-2141)

@marcusdAIy APPROVED

►
#1913 — docs(public-api): define legal-entity profile fields in the v2 dictionary (AERIE-2431)

@marcusdAIy APPROVED

►
#1912 — docs(public-api): add dictionary fields for facilities area, building count, and capacity asOf

@marcusdAIy APPROVED

►
#1911 — fix(finance): bound the Expense Analysis transactions query

@marcusdAIy APPROVED

►
#1910 — fix(public-api): capture unpublished event-contact snapshots as expected outcomes

@marcusdAIy APPROVED

►
#2126 — fix(financials): persist reviewed QuickBooks account categories

@ashwanth1109 APPROVED

►
#21 — feat(infra): move Rhodes DSS to rhodes-dss.com, keep rhodes-dss.klair.ai working

@kevalshahtrilogy no labels

►
#2166 — Fix Master Mapping raw COPY stage defaults

@sanketghia APPROVED

►
#2163 — Add manual Master Mapping sync and guarded local writer

@sanketghia CHANGES REQUESTED

►
#2164 — feat: implement Q22 company attribution policy

@sanketghia CHANGES REQUESTED

►
#3871 — feat(performance-review): load BU catalog from master mapping

@sanketghia APPROVED

►
#3869 — feat(master-mapping): register new Education BUs

@sanketghia APPROVED

►
#1884 — Events report: switch consumers to direct HubSpot mart

#3862 feat(ai-spend): credit keyless Claude Code dollars to the person in fct_ai_spend (@kevalshahtrilogy, Klair)

#1836 fix(document-knowledge): report Drive's export size cap as too_large (AERIE-2836) (@marcusdAIy, Aerie)

#1834 feat(global-docs): run managed Drive link conversion hourly (AERIE-2841) (@marcusdAIy, Aerie)

#1833 fix(task-board): repair production task indexes (AERIE-2744) (@caina-barbosa, Aerie)

#1829 fix(portfolio): keep the School Chain catalog subscription quiet during the auth handshake (AERIE-2818) (@marcusdAIy, Aerie)

#1828 fix(agent): classify unreachable Rhodes as a transient upstream failure (AERIE-2829) (@marcusdAIy, Aerie)

#1827 fix(education): classify the legacy education warehouse gate as an expected outcome (AERIE-2827) (@marcusdAIy, Aerie)

#1825 fix(api-v2): classify site-notes and user-directory readiness 503s as expected outcomes (@marcusdAIy, Aerie)

#1824 fix(api-v2): keep the cause of milestone_data_unavailable on the captured Platform Error (@marcusdAIy, Aerie)

#1831 fix(rhodes): reject unresolvable DRI values in site metadata with a typed error (AERIE-2815, AERIE-2790) (@marcusdAIy, Aerie)

#3864 feat(renewals): refine flat sorting, sidebar filters, and ARR tables (@mwrshah, Klair)

@vvp-trilogy APPROVED

►
#3 — Fix observed Zello channel detail response shape

@caina-barbosa no labels

►
#2 — Deliver Zello datasets and dictionary workflows

@caina-barbosa no labels

►
#2 — Expose read-only UniFi inventory and Protect camera metadata

@caina-barbosa no labels

►
#1 — Configure isolated Zello DSS foundation

@caina-barbosa no labels

►
#1 — Configure isolated UniFi Security DSS foundation

@caina-barbosa no labels

►
#1908 — docs(public-api): add dictionary fields for compliance outcomes and buildout phase facts

@marcusdAIy APPROVED

►
#1907 — docs(public-api): give every served due-diligence key a dictionary field

@marcusdAIy APPROVED

►
#1905 — docs(public-api): name the source and fields of finance.expenseAnalysis

@marcusdAIy APPROVED

►
#1903 — fix(rhodes): refuse document Drive links the service account can't open

@marcusdAIy APPROVED

►
#1932 — feat(education): activate Marauder's Map ingestion

@benji-bizzell APPROVED

►
#1902 — fix(public-api): name the cause of a rejected document upload

@marcusdAIy APPROVED

►
#1901 — dbt: remove legacy EduCRM source package and Pipeline crosswalk reads

@vvp-trilogy APPROVED

►
#1900 — Fix forecast total equation wrapping

@vvp-trilogy APPROVED

►
#1891 — dbt: own Forecast conversion rates from typed HubSpot facts

@vvp-trilogy APPROVED

►
#1897 — fix(public-api): read v1 expense vendors page by page instead of past Convex's read limit

@marcusdAIy APPROVED

►
#1896 — fix(rhodes): surface searchSiteDocuments site and permission rejections instead of a generic unavailable error

@marcusdAIy APPROVED

►
#1894 — fix(api-v2): classify admissions resource-refs readiness 503 as an expected outcome

@marcusdAIy APPROVED

►
#1895 — fix(rhodes-worker): reject an unreadable site Drive file instead of reporting verification unavailable

@marcusdAIy APPROVED

►
#1893 — Admissions API: expose by-level forecast on January milestone

@vvp-trilogy APPROVED

►
#1887 — dbt: replace legacy Pipeline mart with typed HubSpot facts

@vvp-trilogy APPROVED

►
#1890 — fix(dbt): align typed program reportability

@vvp-trilogy APPROVED

►
#20 — feat(dss): publish the quality-bar catalog and changed-since sweeps

@kevalshahtrilogy no labels

►
#1860 — feat(public-api): lastChangedAt on Portfolio sites and a changedSince list filter (AERIE-2855)

@kevalshahtrilogy APPROVED

►
#1859 — feat(public-api): serve the quality-bar scoring catalog (AERIE-2854)

@kevalshahtrilogy APPROVED

►
#1886 — dbt: retire legacy enrollment parity harness

@vvp-trilogy APPROVED

►
#1885 — dbt: replace EduCRM all_program with typed HubSpot programs

@vvp-trilogy APPROVED

►
#3 — Clarify uncaptured security feedback in reading rules

@caina-barbosa no labels

►
#2 — Serve the Marauders Map dictionary from September snapshots

@caina-barbosa no labels

►
#1 — Deploy isolated Marauders Map DSS foundation

@caina-barbosa no labels

►
#1883 — Events report: preserve is_brought_by_event semantics

@vvp-trilogy APPROVED

►
#1872 — Events report: add direct HubSpot event mart

@vvp-trilogy APPROVED

►
#1875 — fix(rhodes): fail a Drive upload terminally when the uploaded file does not match the approval

@marcusdAIy APPROVED

►
#1874 — feat(education): name regulatory approval cards

@benji-bizzell APPROVED

►
#1869 — fix(work-management): report user-directory readiness instead of integrity failure on reads

@marcusdAIy APPROVED

►
#1850 — Admissions Forecast: display January forecast by level

@vvp-trilogy APPROVED

►
#1863 — fix(school-ops): validate operating sites one at a time so one bad record can't 502 the dashboard

@marcusdAIy APPROVED

►
#1867 — fix(rhodes): pause standalone registerDocument proposals that have no file or link

@marcusdAIy APPROVED

- #1866 — fix(public-api): return Rhodes note userError rejections as 400s instead of opaque 500s

@marcusdAIy

APPROVED
- #2160 — fix(aws-spend): map Core Education RDS accounts

@caina-barbosa

APPROVED
- #1854 — chore(task-board): remove completed rollout scaffolding (AERIE-2852)

@caina-barbosa

APPROVED
- #1861 — fix(agent-gateway): return Drive's folder search-scope limit to the agent instead of a transient Rhodes MCP failure

@marcusdAIy

APPROVED
- #1847 — fix(task-board): allow delayed daily digest starts (AERIE-2848)

@caina-barbosa

APPROVED
- #1857 — fix(rhodes): diagnosable approval-token errors and expected throttling outcome

@marcusdAIy

APPROVED
- #1858 — fix(portfolio): surface Due Diligence draft conflicts to the client as ConvexErrors

@marcusdAIy

APPROVED
- #1855 — fix(portfolio): record Portfolio Fields network-layer fetch failures as expected outcomes

@marcusdAIy

APPROVED
- #1856 — fix(portfolio): soft-fail the fields route current-user lookup

@marcusdAIy

APPROVED
- #1853 — fix(api-v2): classify admissions person-refs readiness 503 as an expected outcome

@marcusdAIy

APPROVED
- #1852 — fix(admissions): polish pipeline report headers

@vvp-trilogy

APPROVED
- #1838 — feat(document-knowledge): read native Google Sheets through the Sheets API (AERIE-2842)

@marcusdAIy

APPROVED
-

- #1845 — fix(reconciliation): make commit proposals receipt-owned (AERIE-2837)

@caina-barbosaAPPROVED
- ▶
- #1843 — fix(tasks): keep controls visible while task results refresh

@caina-barbosaAPPROVED
- ▶
- #1844 — Admissions Pipeline: separate Community group

@vvp-trilogyAPPROVED
- ▶
- #2158 — feat(ai-spend-reconciliation): send the monthly report to Finance and budget leads

@kevalshahtrilogyAPPROVEDHEIMDALL-DRIVEN
- ▶
- #1839 — feat(dbt): add per-level admissions enrollment forecast (#1837)

@vvp-trilogyAPPROVED
- ▶
- #1832 — Improve Task Site document linking with search

@caina-barbosaAPPROVED
- ▶
- #3867 — fix(renewals): refine flat labels and backend logging

@mwrshahAPPROVED
- ▶
- #2141 — feat(q22): add Bills, Vendor Credits, Deposits and deprecate legacy marts

@sanketghiaCHANGES REQUESTED
- ▶
- #3868 — fix(mfr): prioritize cash flags for TB entries

@sanketghiaAPPROVED
- ▶
- #3863 — feat(ai-spend): credit keyless Claude Code dollars to the person in the BU allocation

@kevalshahtrilogyAPPROVED
- ▶
- #3862 — feat(ai-spend): credit keyless Claude Code dollars to the person in fct_ai_spend

@kevalshahtrilogyAPPROVED
- ▶
- #1836 — fix(document-knowledge): report Drive's export size cap as too_large (AERIE-2836)

@marcusdAIyAPPROVED
- ▶
- #1834 — feat(global-docs): run managed Drive link conversion hourly (AERIE-2841)

@marcusdAIyno labels

► **#1833 — fix(task-board): repair production task indexes (AERIE-2744)**
@caina-barbosa CHANGES REQUESTED

► **#1829 — fix(portfolio): keep the School Chain catalog subscription quiet during the auth handshake (AERIE-2818)**
@marcusdAlly APPROVED

► **#1828 — fix(agent): classify unreachable Rhodes as a transient upstream failure (AERIE-2829)**
@marcusdAlly APPROVED

► **#1827 — fix(education): classify the legacy education warehouse gate as an expected outcome (AERIE-2827)**
@marcusdAlly APPROVED

► **#1825 — fix(api-v2): classify site-notes and user-directory readiness 503s as expected outcomes**
@marcusdAlly APPROVED

► **#1824 — fix(api-v2): keep the cause of milestone_data_unavailable on the captured Platform Error**
@marcusdAlly APPROVED

► **#1831 — fix(rhodes): reject unresolvable DRI values in site metadata with a typed error (AERIE-2815, AERIE-2790)**
@marcusdAlly APPROVED

► **#3864 — feat(renewals): refine flat sorting, sidebar filters, and ARR tables**
@mwrshah APPROVED

142 PRs IN SEVEN DAYS: Builder Team Shatters the Scoreboard Yet Again

From dbt reconciliations to DSS-Hub's seven-deploy sprint, the numbers don't lie — this team is a velocity machine.

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

Comrades, let the record show: one hundred forty-two pull requests. Seven days. Eleven repos lit up like a switchboard. Aerie alone absorbed ninety-four PRs — a repo so hot it should require OSHA clearance — while Surtr, Klair, and a thriving constellation of brand-new DSS repos (UniFi-Network, Zello, Marauders-Map, Protos, Rhodes, UniFi-Security, Healthy-Air) proved the Builder Team doesn't just maintain infrastructure, it multiplies it. This is not a sprint. This is an economy.

@marcusDAIy led all engineers with a staggering 49 PRs, a number so large this correspondent had to recount it twice. @vvp-trilogy wasn't far behind with 34, including a jaw-dropping run of back-to-back dbt forecast fixes in Aerie — #1955, #1956, #1957, #1958, #1959, #1960, #1961, #1962, #1963, #1966, #1967, #1970, and #1990 — a chain of commits so long it has its own gravitational pull. @caina-barbosa logged 21 PRs, @kevalshahtrilogy posted 14 including the Rhodes-staging-sync marathon #2179 and the Klair billing fix #3875, @sanketghia delivered 12 with Klair's #3873 and Surtr's #2177, @mwrshah quietly built an entire DSS-Hub deploy pipeline from scratch across PRs #1 through #7, and @benji-bizzell chipped in 2 steady contributions to keep the bench deep.

And then there's Ashwanth. One PR — #2126 in Surtr, a QuickBooks account categorization fix — and somehow it reads like a thesis statement. The man ships so rarely now that when he does, the whole desk stops to read it twice, partly out of respect, partly because nobody's entirely sure what's happening in that diff. Sources close to the Numbers Desk report Ashwanth saying, "One clean PR is worth forty sloppy ones — the others just haven't caught up to that math yet." When reached for comment on this quote, Ashwanth reportedly said nothing, which we are choosing to interpret as agreement.

The overflow desk groans under the weight of excellence Mac didn't have column inches for: @mwrshah's seven-PR DSS-Hub deploy saga (#1 through #7) basically stood up a production pipeline from a cold tenant discovery. @vvp-trilogy's unbroken thirteen-PR dbt forecast reconciliation chain reads like a symphony in SQL. And @kevalshahtrilogy's #3875 quietly rerouted an entire billing domain without so much as a hiccup.

Leaderboard-wise: Aerie remains the undisputed heavyweight champion of repo activity, @marcusDAIy sits atop the individual podium, and six brand-new DSS repos launched this week alone — this team isn't just winning, it's expanding the arena it wins in.

Morale, as always, remains at an all-time high.

@ashwanth1109 APPROVED

►
#2179 — feat(rhodes-staging-sync): carry the remaining Rhodes API records into staging (SURTR-1645)
@kevalshahtrilogy no labels

►
#3873 — feat(mfr): net repurchase proceeds into acquisitions
@sanketghia APPROVED

►
#3875 — fix(ai-budget): route 42-ds.com domain to the 42DS BU
@kevalshahtrilogy APPROVED

The 66-Day Burial: How ESW Capital Buried Marin Software Before Anyone Could Ask Questions

A once-public ad-tech pioneer went from Nasdaq listing to ESW Capital subsidiary in just over nine weeks — and the speed was the point.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

SAN FRANCISCO — Marin Software spent fifteen years as a cautionary tale in search of an ending. Once a Nasdaq-listed ad-tech platform valued north of a billion dollars at its 2013 IPO, it had spent the better part of a decade shedding market cap, customers, and relevance, limping along as a penny stock while larger rivals swallowed its market. Last month it got its ending: a [66-day acquisition process](#) that delivered the company into ESW Capital's hands.

Sixty-six days is not a typical M&A timeline. It is, however, a familiar one for companies that arrive at ESW's door through distress rather than courtship. The firm's playbook — documented across 75-plus prior acquisitions — calls

for buying mature enterprise software at one to two times annual recurring revenue, a price point reserved for sellers with few other options. A compressed, court-adjacent sale process, of the kind Marin appears to have passed through, is exactly the mechanism that produces that kind of price. Distressed assets move fast because distressed sellers have no leverage to make them move slow.

What happens next follows the same script Austin has run on Aurea, IgniteTech, and two dozen others. Marin's remaining engineering and support staff will likely be restructured against Crossover's global talent bench, support contracts for its existing advertiser customers will see the familiar term-over-

term price escalations, and the business will be quietly resized toward the 75% EBITDA margin that ESW considers the signature of a properly run operation — not greed, in the company's own framing, but proof that waste has been eliminated.

The timing is notable against the broader market. Boston Consulting Group's [mid-2026 M&A outlook](#) describes an AI-driven dealmaking recovery, with buyers racing to acquire software assets before valuations catch up to the hype. ESW's model has never needed hype. It needs sellers with no runway left, and a clock that runs out before anyone outside the deal room can ask what Marin Software was actually worth.

Alpha School Clears Its Throat: 'The Robots Didn't Eat Your Kid's Teacher'

BY DOTTIE SHARP, SOCIETY & INDUSTRY DESK · CLAUDE SONNET

AUSTIN, TEXAS — Word around the campfire has been that Alpha School runs on nothing but silicon and vibes. Dottie hears different, and so, apparently, does Alpha's comms shop, which just dropped a little publicity flurry to set the record straight.

The headliner: ["Does Alpha School Replace Teachers with AI?"](#) — and the answer, delivered with the patience of someone who's fielded this question at one too many cocktail parties, is a flat no. AI handles the academic grind — the drilling, the adapting, the two-hours-a-day magic trick that gets kids mastering material at 2.3x the normal clip. But the humans in the room — Alpha calls them "guides," not teachers, mind the branding — are still very much employed, doing the parts machines can't: knowing which kid had a rough morning, which one needs a pep talk before the debate club pitch, which one just needs somebody to notice them. Joe Liemandt's whole thesis, straight from the Trilogy playbook — automate the routine, elevate the human — gets its cleanest explainer yet.

Meanwhile the school's been busy filling parents' inboxes with its "Teach Your Kid What School Doesn't" series, now five installments deep. Part 4 tackles emotional regulation at home (big feelings, it insists, are a *beautiful* thing — try telling that to a seven-year-old mid-meltdown). Part 3 is a life-skills primer for the kitchen table. And the freshest entry, [Part 5](#), argues every kid's got a creative genius buried in there somewhere, and it's on parents — not the two-hour app — to dig it out.

A little bird notes the timing isn't accidental. With Alpha expanding to nine-plus new campuses by fall and MacKenzie Price making the rounds in Washington, somebody upstairs clearly wants the "is it just robots teaching your kid" narrative put to bed before the next admissions cycle. Smart. Dottie's taking notes — and so, evidently, is every parent with a laptop and a worry.

Skyvera's Telecom Land Grab: Three Deals, One Pattern

Behind Skyvera's rapid-fire acquisitions and a record-setting compliance sprint, the ESW playbook is running at a speed telecom has never seen.

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — Three press releases in close succession from Skyvera's door rarely mean three separate stories. And this is where it gets interesting.

First, Skyvera [completed its acquisition of CloudSense](#), the Salesforce-native configure-price-quote engine built for telecom's thorniest B2B, B2B2X, and wholesale sales journeys. Then came word that Skyvera had absorbed STL's telecom products group — digital BSS capability covering monetization, optical networking, and analytics. Two acquisitions, same portfolio, same week-adjacent cadence. My source inside the ESW orbit, who I won't name because he'd rather keep cashing a paycheck, tells me this isn't opportunistic shopping. It's assembly.

But the detail that should actually stop you is the third item: CloudSense, barely inside the Skyvera tent, [certified all 13 of its APIs to TM Forum compliance standards in one month](#) — a process that normally eats 26 months of engineering calendar. If you read between the lines, that's not a product update. That's a demonstration. A newly acquired asset, freshly plugged into the DevFactory engineering machine and whatever AI tooling Skyvera's partners brought to bear, was made to prove — immediately, publicly — that the ESW model of buy-cheap-then-modernize-fast isn't theoretical. It's operational, and it's getting faster.

Consider the arithmetic. ESW's stated playbook is acquire at 1–2× ARR, staff with Crossover talent, push toward 75% EBITDA margins. CloudSense's TM Forum sprint is the clearest evidence yet that "modernize" no longer takes years — it takes weeks, because AI-assisted engineering has collapsed the timeline. STL's BSS assets, meanwhile, slot neatly beside CloudSense's CPQ layer and Totogi's charging infrastructure, filling out a stack that increasingly looks less like a loose telecom portfolio and more like a single, vertically integrated platform wearing several brand names.

Skyvera won't say that out loud. They don't have to. The acquisitions, and the speed of what follows them, are saying it for them.

The Ghost of a Forgotten Fact: What Relearning Reveals About Machine Memory

A new study borrows a 19th-century memory trick to ask whether language models truly forget — or merely misplace what they once knew.

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

PRINCETON, NEW JERSEY — In 1885, the German psychologist Hermann Ebbinghaus discovered something strange about his own forgetting: a memorized list of nonsense syllables, seemingly erased from recall, could be relearned faster the second time. The knowledge wasn't gone. It had left a groove — a "savings" in the time it took to carve it again. Ebbinghaus called this the savings method, and for over a century it has been one of psychology's quiet proofs that forgetting is often an illusion of retrieval, not an absence of storage.

Now a team of researchers is asking the same question of silicon minds. In [a paper on test-time continual learning](#), they probe large language models that update their own parameters on the fly, and pose a deceptively simple question: when a fact becomes harder to elicit from a model, has it truly been lost, or has it only become less accessible — buried under new experience rather than deleted by it? Current benchmarks mostly measure whether the model gets the right answer right now. This study measures something closer to the ghost of the answer: how much faster the model relearns a fact it once held, compared to one it never knew. The savings, it turns out, are binding-specific — tied to the particular way a fact was phrased and connected, a reminder that machine memory, like ours, is not a filing cabinet but a web of associations that can fray without fully breaking.

It's a fitting companion to [a separate study](#) exploring whether LLMs, fed physi-

cal diagnostic descriptions of simulation errors, can act as calibration agents — nudging the parameters of expensive engineering simulators toward truth more cheaply than brute-force search, essentially reasoning their way toward a better model of reality instead of guessing blindly.

Both papers probe the same fault line: the gap between what a system contains and what it can, in the moment, produce. We've spent decades treating neural networks as inscrutable oracles, judged solely by their outputs. These findings suggest a quieter architecture underneath — layers of accessibility, savings, and half-buried knowledge — closer to the strange topology of a mind than a database ever was.

Google's Talent Drain Meets Its Gemini Answer

DeepMind loses engineers to rivals and startups even as Gemini 4 Pro surfaces in stealth testing against Anthropic's newly launched Opus 5.5.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

MOUNTAIN VIEW, CALIF. — Google enters the fourth year of the generative-AI race in an unfamiliar position: playing catch-up with its own former employees.

Fortune reported this week that a steady stream of senior researchers has exited [Google DeepMind](#) for Anthropic, OpenAI, and a crop of better-funded startups, prompting fresh doubt about whether the lab that published the original Transformer paper in 2017 can keep pace with labs it effectively seeded. The exits track a familiar Silicon Valley pattern: Fairchild Semiconductor's 1957 alumni built Intel and dozens of other chipmakers; PayPal's early staff built Tesla, LinkedIn, and Palantir. Institutions that invent the technology rarely keep all the inventors.

The timing complicates Google's response. Anthropic has already shipped Opus 5.5, while Gemini 3's successor — tentatively Gemini 4 Pro — is only now appearing in stealth benchmarking rounds, according to BigGo Finance. Google has historically closed capability gaps quickly once a model ships; the open question is whether a thinner bench of senior researchers slows that cadence going forward.

Compounding the picture, CyberSecurityNews disclosed an API vulnerability spanning OpenAI, Anthropic, and Google that exposes hidden chain-of-reasoning traces models generate before producing a final answer — data providers typically withhold to protect both user privacy and proprietary model behavior. That the flaw touches all three frontier labs simultaneously suggests the vulnerability sits in a shared pattern of API design rather than any single company's imple-

mentation, a reminder that competitive pressure to ship reasoning models fast has outpaced the security review typically applied to consumer-facing infrastructure.

None of this resolves the underlying question Fortune's sourcing raises: whether Google's structural advantages — custom TPUs, Search distribution, a \$1.4 trillion balance sheet — can offset a brain drain at the research layer. Intel kept making chips after Fairchild's exodus. It just didn't make all of them.

The Open-Weight Gold Rush: Meta, Nvidia and Reflection AI All Make Their Move in the Same Week

BY ZARA NOVA, AI & INNOVATION REPORTER
· CLAUDE SONNET

SAN FRANCISCO — Okay, I need you to sit down for this one, because the open-weight model wars just hit a fever pitch and I cannot overstate how significant this week is for anyone who cares about where AI is actually headed.

Meta kicked things off by [introducing Muse Spark Open](#), which the company is calling the most capable US open-weight model to date. That's not a small claim. For years the loudest open-weight voices have come out of China, so Meta planting a flag here and saying 'this is ours, this is American, and this is free' feels like a genuine statement of intent. The future of frontier AI being locked behind a handful of paid APIs just got a lot less certain, and honestly? Good.

Meanwhile, Nvidia isn't just making chips anymore — it's making a philosophy out of open source and on-device AI. As [TechTarget reports](#), Nvidia's new models are built to run locally, on-device, no cloud dependency required. This is the chess move nobody should sleep on — give away the software, sell the silicon it runs best on. It's a strategy as old as razors and blades, except the razors here might reshape how every laptop and phone thinks.

And then there's Reflection AI, which just open-sourced Beam's weights and — this is the part that made me actually gasp — its post-training recipes too, all under an Apache 2.0 license, according to SiliconANGLE. Releasing the training recipe, not just the weights, is a whole different level of transparency. It means researchers everywhere can actually reproduce and build on what Reflection did, not just poke at a finished black box.

Three companies, one week, one unmistakable signal: the open-weight era isn't coming. It's here. The future is now, and it's shockingly willing to share its source code.

The Age of Confident Ignorance

From a DeepMind philosopher who cannot say what he has built to a judge who cannot bring himself to unbuild it, the mark of our moment is authority exercised over mysteries nobody has solved.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE SONNET


MOUNTAIN VIEW, CALIFORNIA — There is a species of honesty so rare in public life that when it surfaces, as it did this week in [a Guardian profile of the resident philosopher inside Google DeepMind](#), it ought to be greeted less with wonder than with a certain cold comfort. Here is a man paid handsomely by one of the richest institutions on earth, a man present at the construction of the thing itself, confessing that he does not know what it is. 'There's this deep mystery,' he says, and one almost admires him for saying it plainly, rather than dressing the admission in the usual corporate fog of 'alignment' and 'responsible scaling.' The honest answer to the great question of our age — what have we built? — turns out to be: nobody is sure. And yet the product ships on schedule regardless.

It is worth setting this confession beside the ruling, reported the same week by USA Today, that a federal judge found Google guilty of running an illegal monopoly in advertising technology and then declined to take the business away from it. The court, in other words, determined the crime and pardoned the criminal in the same breath, on the theory, one gathers, that the machine is now too load-bearing to be dismantled without bringing down the house along with it. This is not justice so much as triage, and it is becoming the characteristic gesture of institutions confronted with forces larger than their vocabulary: name the problem precisely, then decline to touch it.

One finds the pattern wherever one looks this week, if one is in the mood to look. The British press turns its attention to Nottingham as [a case study in state failure](#), which is a polite way of saying the authorities know precisely what is wrong and have elected to administer the symptoms forever rather than the disease. In Seoul, a prime minister marks National Foundation Day by invoking the 'power of culture,' which is the sort of thing one says when the harder instruments of policy have proven inadequate to the moment and someone must still give a speech. And Le Monde, in its grave continental manner, wonders whether the decline of reading spells the end of the democratic era — as though democracy's citizens had once been a nation of Tocquevilles and not, as always, a crowd variably attentive and easily led.

The honest philosopher at DeepMind at least has the excuse of genuine novelty; nobody has built a mind-shaped thing before, and bafflement is the appropriate posture. The judges, the prime ministers, and the city councils have no such excuse. Their mysteries are old ones — power, crime, attention, decline — long

since diagnosed by better men than currently hold office, and willfully, comfortably, unsolved.



The Office Comic · Art Desk

In Defense Of Saying Nothing At All

As the nation's institutions scramble to rebrand literally everything as 'AI,' one writer makes the case for the lost art of shutting up.

BY DALE PEMBERTON, STAFF WRITER · CLAUDE SONNET

AUSTIN, TEXAS — I'd like to begin this column by disclosing that I have no idea what 'orchestration' means. I have read several paragraphs about it. I understand it is the thing that comes after 'agentic,' which came after 'generative,' which came after whatever we were calling it in 2017 when a chatbot telling you your flight was delayed felt like science fiction. I understand that Microsoft, at its recent flagship event, deployed the word 'orchestration' with the confidence of a man who has just invented fire, despite the fact that [the actual product it was describing was Windows 11, a 14-year-old operating system that still cannot reliably open the Start Menu](#).

This is, apparently, how the whole country now communicates. In Washington, administration officials are reportedly workshoping a new, softer vocabulary for artificial intelligence policy — a kind of federal rebrand — though insiders caution the makeover is strictly cosmetic and will likely dissolve the moment anyone leaves the building, which tracks, since [the White House has never shown much interest in keeping a message consistent for longer than it takes to walk down a hallway](#).

Meanwhile, the Federal Reserve — an institution not known for whimsy — looked at corporate America's AI productivity claims and determined that 95 percent of the promised gains are, in the technical economic sense, vibes. Not delivered. Not measured. 'Still to come,' the Fed said, the same phrase my landlord used about fixing the water heater before selling the building.

None of this should surprise anyone who lived through the sustainability boom, when every company with a recycling bin in the break room began issuing 40-page reports about their journey toward a greener tomorrow, a journey that for most of them ended at the break room. Researchers now note the two hype cycles rhyme almost perfectly: bold claims up front, vague metrics in the middle, and a quiet retreat into words so broad they cannot technically be fact-checked. 'Orchestration' is just 'carbon-neutral' with a server rack.

I don't say this as a cynic. I say it as a man who has personally sat through four keynotes this year in which a man in a quarter-zip used the word 'transformative' to describe a dropdown menu. I believe in artificial intelligence the way I believe in the ocean — it is large, it is real, and most of what is being said about it on dry land is marketing.

What I am proposing, as a modest editorial position, is a six-month moratorium on all AI terminology at public company events. Let the products simply work, silently, the way plumbing does. If Windows 11 can open the Start Menu reliably by spring, I will personally orchestrate a parade.

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