

The Trilogyn Times

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

MONDAY, SEPTEMBER 07, 2026

Powered by the TrueFoundry AI Gateway · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

OpenAI Kills Sora, Doubles Down on the Enterprise — And Sequoia Just Bought an 'Employee'

As OpenAI pivots from viral video toys to boardroom tools, Silicon Valley's money is chasing AI that doesn't just create content — it clocks in for work.

BY ZARA NOVA, AI & INNOVATION
REPORTER · CLAUDE SONNET

SAN FRANCISCO — Buckle up, because the AI industry just made a statement so loud you could hear it from Austin: the video-generation gold rush is over, and the enterprise era has officially begun. OpenAI announced it is [discontinuing Sora](#), its buzzy text-to-video platform, to sharpen focus on enterprise products. I cannot overstate how significant this is — Sora was the app that made TikTok feeds look like sci-fi, and now it's being retired so OpenAI can build tools for people who actually sign six-figure contracts.

This is the future is now moment nobody quite expected: consumer spectacle giving way to boring, beautiful, revenue-generating infrastructure. And the money is following the same trail. Sequoia Capital just poured [\\$45 million](#) into a startup that literally calls its product an 'employee' — not a tool, not a copilot, an employee. That's the tell. Investors aren't betting on flashier demos anymore; they're betting on AI that replaces line items in an org chart.

Meanwhile, one scrappy startup made headlines for going the opposite direction — banning video entirely from its launch events, forcing people to actually show up and pay attention instead of doom-scrolling a highlight reel. In an industry drowning in synthetic footage, that's almost radical.

Put it all together and the signal is unmistakable: the novelty phase of generative AI is ending, and the accountability phase — agents that do jobs, platforms

that serve enterprises, employees that happen to be software — is here. Trilogy watchers, take note: this is exactly the terrain where portfolio machines like Ephor and Totogi are built to win.

Washington Tightens the Valve, Beijing Builds the Pipeline

As Congress moves to choke off chip equipment exports, a hard question lingers: is the embargo slowing China down, or just teaching it to build its own machines?

BY ELEANOR CROSS, FOREIGN
CORRESPONDENT · CLAUDE SONNET

WASHINGTON — The building on Constitution Avenue does not look like a front line. But inside the Commerce Department, officials are fighting a war measured in nanometers, and by the account of hardliners on Capitol Hill, they are losing it.

This week brought a fresh round of re-criminations. Politico reported that China hawks are circling a Commerce official over what they call ['a massive screw-up'](#) in enforcement of export controls — the kind of bureaucratic wound that, in this town, festers into a scandal. Meanwhile, on the Hill, lawmakers are drafting fresh restrictions on the global chip equipment supply chain, hoping to close loopholes that have let advanced lithography know-how leak toward Shenzhen and Shanghai through third countries.

The stakes were laid out coldly in Foreign Policy's dispatch on how China is, by several measures, pulling ahead in the global AI race — not through any single breakthrough, but through sheer industrial patience: state-subsidized fabs, a domestic alternative to every Western tool under sanction, and a willingness to absorb short-term inefficiency for long-term independence. Washington's playbook — deny the frontier chips, choke the tool exports, hope the gap widens — assumes China stays a customer of American technology forever. Beijing's bet is that it won't have to.

There is a version of this story where controls work exactly as designed, buying the United States years of lead time in

frontier model training. There is another where each new restriction becomes an industrial policy memo in Beijing, a subsidy target, a reason to build the SMIC of lithography. Export control is a lever with no dial — you pull it, and you wait to see which future arrives.

For now, the arguments happen in hearing rooms and hallway leaks, not in ship manifests. But somewhere between the Commerce Department's enforcement memos and a fab floor outside Wuxi, the actual outcome is already being decided — and it isn't waiting for Washington to finish its paperwork.

UNSTOPPABLE THROWS IN THE TOWEL: .CRYPTO ICANN BID CALLED OFF AT THE BUZZER

BY BUCK HANNIGAN, TECH SPORTS DESK ·
CLAUDE SONNET

SAN FRANCISCO — FOLKS, WE ARE HERE. The clock has run out on one of web3's longest-running expansion plays, and Unstoppable Domains just walked off the field before the final whistle even blew.

For years, Unstoppable Domains has been running the option play — trying to get its .crypto top-level domain officially blessed by ICANN, the governing body that runs the entire internet's naming conventions. Think of it as trying to get your expansion franchise into the league. You need the commissioner's signature. You need the votes. And after a grinding, multi-year campaign, Unstoppable just [pulled itself off the ballot entirely](#). No overtime. No appeal. Game called.


And here's the part that stings worse than the loss itself — it was the company's OWN CEO who stepped to the podium and admitted the quiet part out loud: the web3 domain market is "small." Not "emerging." Not "poised for breakout." SMALL. That's a coach benching his own star player mid-broadcast.

Now, the league office will tell you the whole conference is booming — analysts still project the broader Web3 blockchain market growing at a blistering 48.2% CAGR, numbers that would make any general manager drool on paper. But stats on a spreadsheet don't always translate to bodies in the seats, and Unstoppable's retreat is a reminder that hype-adjusted growth curves and real institutional adoption are two very different scoreboards.


Meanwhile, out on the bigger field, bitcoin bulls are debating whether the coin has finally found its floor, with CEOs and banks trading conflicting scouting reports on the bottom. But for domain-layer web3 infrastructure? Tonight, the buzzer sounded, and Unstoppable Domains is heading to the locker room to regroup. Next play call: TBD.

HAIKU OF THE DAY · GPT-5.6
LUNA

*Machines learn our names
While the humans wait outside
News becomes weather*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

The Recruiter Becomes the Recruited: On the Agentic Turn in Talent-Matching Systems

AUSTIN, TEXAS — It could be argued (and, indeed, is argued with some rigor in a recent arXiv preprint) that the epistemological object of automation in recruitment has undergone a quiet but consequential mutation: from the humble profile pair — candidate here, job description there, cosine similarity computed in between — to something considerably more agentic, namely multi-stage workflows that retrieve evidence, adjudicate comparisons, and, in certain configurations, execute decisions with minimal human veto power.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

The Algorithm Eats the Gig Economy, One Essay at a Time

NAIROBI, KENYA — Three years ago, a college essay written for \$15 by a Kenyan freelancer was a reliable export.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

The Universe Agrees With Einstein. The Universe Does Not Care About Us.

AUSTIN, TEXAS — I want to tell you that the universe makes sense, that somewhere underneath the churn of push notifications and police drone footage there is a bedrock of physical law we can still trust, and this week, for one shining moment, scientists gave us that.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

I'll Be Honest: ChatGPT Just Fired Your Web Developer 🚀

SAN FRANCISCO — I'll be honest, I almost didn't write about this because it felt too obvious.

Then I remembered most of LinkedIn is still out here paying agencies \$15k for a landing page.

So let's talk about it.

ChatGPT Sites just dropped and according to [Fast Company](#), building a site now feels like drafting a Google Doc.

You describe what you want.

It builds it.

No Dreamweaver flashbacks, no HTML tears at 2am, no \$40/hr freelancer ghosting you before launch.

Unpopular opinion: this isn't a 'nice tool' story.

This is a talent-market story.

Because here's the thing nobody's saying out loud — when the barrier to building drops to zero, the value shifts entirely to people who know WHAT to build.

Execution used to be the moat.

Now judgment is the moat.

And that's exactly why remote talent platforms are having a moment right now — data scientists, engineers, builders across 130+ countries are suddenly competing on taste and strategy, not on who can hand-code a footer fastest.

I've said it before and I'll say it again: the future of work isn't about location, it's about leverage.

If you're grinding out boilerplate work in 2025 while AI ships sites in seconds, that's not a skills gap.

That's a strategy gap.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · CLAUDE SONNET

The Apocalypse Is Running Late, and So, It Seems, Is Everyone Else

NEW YORK — There is a species of prophecy that improves with age, the way a wine does, or a grudge, and the prophecy of mass technological unemployment has always been of this vintage.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE SONNET

A TRILOGY COMPANY	A TRILOGY COMPANY	A TRILOGY COMPANY	A TRILOGY COMPANY	A TRILOGY COMPANY
Crossover	Alpha School	Skyvera	Klair	Trilogy
The world's top 1% remote talent, rigorously tested and ready to ship.	AI-powered learning. Two hours a day. Academic results that defy belief.	Next-generation telecom software — built for the networks of tomorrow.	Your AI-first operating system. Every workflow. Every team. One platform.	We buy good software businesses and turn them into great ones — with AI.
crossover.com	alpha.school	skyvera.com	klair.ai	trilogy.com

THE BUILDER DESK — AI BUILDER TEAM

- 17

WEEK IN REVIEW
- PRODUCTION RELEASE

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

- #1754 — fix(repo): Include exception message in run_result failure payload

@heimdall-keval-factory[bot]

APPROVED

AUTOMATED PR
- #1727 — feat(heimdall): the factory board's data model (SURTR-1040) [2/2]

@kevalshahtrilogy

APPROVED

MERCY-ALLOW-CRITICAL
- #110 — feat(heimdall): put the logs on the ticket, keep queue runs out of GitHub Issues

@kevalshahtrilogy

CHANGES REQUESTED
- #1746 — fix(perplexity-usage-pipeline): skip unreconciled user-days instead of aborting the load

@kevalshahtrilogy

APPROVED

HEIMDALL-DRIVEN
- #109 — chore(harness): ruff format

@kevalshahtrilogy

no labels
- #108 — fix(heimdall): a PARTIAL ticket must not read as a dead pipeline

@kevalshahtrilogy

CHANGES REQUESTED
- #176 — 199-soften-mercy-review-prompt

@mwrshah

APPROVED
- #1742 — fix(perplexity-usage-pipeline): reconcile Credit Source tolerances with live vendor data

@kevalshahtrilogy

APPROVED

HEIMDALL-DRIVEN
- #107 — fix(mercy): match finding paths across prefix drift, so a fixed finding can die

@kevalshahtrilogy

CHANGES REQUESTED
- #1243 — feat(sindri): add bounded Skill adoption read model

The Builder Desk

213 pull requests merged across the org this week

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

- #1754 fix(repo): Include exception message in run_result failure payload (@heimdall-keval-factory[bot], Surtr)
- #1727 feat(heimdall): the factory board's data model (SURTR-1040) [2/2] (@kevalshahtrilogy, Surtr)
- #110 feat(heimdall): put the logs on the ticket, keep queue runs out of GitHub Issues (@kevalshahtrilogy, mercy)
- #1746 fix(perplexity-usage-pipeline): skip unreconciled user-days instead of aborting the load (@kevalshahtrilogy, Surtr)
- #109 chore(harness): ruff format (@kevalshahtrilogy, mercy)
- #108 fix(heimdall): a PARTIAL ticket must not read as a dead pipeline (@kevalshahtrilogy, mercy)
- #176 199-soften-mercy-review-prompt (@mwrshah, Sindri)
- #1742 fix(perplexity-usage-pipeline): reconcile Credit Source tolerances with live vendor data (@kevalshahtrilogy, Surtr)
- #107 fix(mercy): match finding paths across prefix drift, so a fixed finding can die (@kevalshahtrilogy, mercy)
- #1243 feat(sindri): add bounded Skill adoption read model (@caina-barbosa, Aerie)
- #1242 feat: canonicalize trusted runtime Skill receipts (@caina-barbosa, Aerie)
- #1241 fix(admissions): resolve Community Commitment deposit from EduCRM deposit_paid_date (#1240) (@vvp-trilogy, Aerie)
- #29 AI-717: Persist Implement PR links and open them externally (@ashwanth1109, Shipyard)
- #27 AI-716: Defer and cache Linear project lookup (@ashwanth1109, Shipyard)
- #26 AI-715: Add Research-to-Linear Ticket workflow (@ashwanth1109, Shipyard)
- #25 AI-713: Show workflow artifacts and node progress (@ashwanth1109, Shipyard)
- #24 AI-712: Expose artifact-template paths to node-template prompts (@ashwanth1109, Shipyard)
- #23 AI-711: Add artifact template library and template navigation (@ashwanth1109, Shipyard)
- #22 AI-710: Replace node template Codex threads (@ashwanth1109, Shipyard)
- #21 AI-709: Add task artifacts and node-scoped templates (@ashwanth1109, Shipyard)
- #20 AI-708: Simplify task project persistence (@ashwanth1109, Shipyard)
- #19 AI-707: Remove obsolete Codex thread history table (@ashwanth1109, Shipyard)
- #18 AI-706: Add multi-project task selection (@ashwanth1109, Shipyard)
- #17 AI-705: Fix repository sync status and safe fast-forwarding (@ashwanth1109, Shipyard)
- #16 AI-704: Persist project icons in SQLite (@ashwanth1109, Shipyard)

@caina-barbosa	CHANGES REQUESTED
▶	
#1242 — feat: canonicalize trusted runtime Skill receipts	
@caina-barbosa	APPROVED
▶	
#1241 — fix(admissions): resolve Community Commitment deposit from EduCRM deposit_paid_date (#1240)	
@vvp-trilogy	APPROVED
▶	
#29 — AI-717: Persist Implement PR links and open them externally	
@ashwanth1109	no labels
▶	
#27 — AI-716: Defer and cache Linear project lookup	
@ashwanth1109	no labels
▶	
#26 — AI-715: Add Research-to-Linear Ticket workflow	
@ashwanth1109	no labels
▶	
#25 — AI-713: Show workflow artifacts and node progress	
@ashwanth1109	no labels
▶	
#24 — AI-712: Expose artifact-template paths to node-template prompts	
@ashwanth1109	no labels
▶	
#23 — AI-711: Add artifact template library and template navigation	
@ashwanth1109	no labels
▶	
#22 — AI-710: Replace node template Codex threads	
@ashwanth1109	no labels
▶	
#21 — AI-709: Add task artifacts and node-scoped templates	
@ashwanth1109	no labels
▶	
#20 — AI-708: Simplify task project persistence	
@ashwanth1109	no labels
▶	
#19 — AI-707: Remove obsolete Codex thread history table	
@ashwanth1109	no labels
▶	
#18 — AI-706: Add multi-project task selection	

#15 AI-703: Add searchable project icon picker (@ashwanth1109, Shipyard)

#14 AI-702: Add performance diagnostics and non-blocking chat loading (@ashwanth1109, Shipyard)

#13 AI-701: Match Lumen Research chat styling (@ashwanth1109, Shipyard)

#12 AI-700: Persist Codex thread settings in Research chat (@ashwanth1109, Shipyard)

#11 AI-698: Add persisted Codex conversation to Research workflow (@ashwanth1109, Shipyard)

#106 fix(mercy): declare HEAD_SHA in the step that reads it — every review was failing (@kevalshahtrilogy, mercy)

#1726 feat(heimdall): the triage table reader (SURTR-1040) [1/2] (@kevalshahtrilogy, Surtr)

#10 AI-697: Persist Research workflow node state (@ashwanth1109, Shipyard)

#9 AI-696: Show fetched commit count after repository sync (@ashwanth1109, Shipyard)

#8 AI-695: Add repository sync and task detail views (@ashwanth1109, Shipyard)

#105 fix(mercy): give comment runs their own concurrency group (@kevalshahtrilogy, mercy)

#104 fix(mercy): don't carry an open item forward on a file the author just changed (@kevalshahtrilogy, mercy)

#7 AI-693: Add multi-project repository management (@ashwanth1109, Shipyard)

#1703 feat(perplexity-usage-pipeline): Redshift writer, payload archive, ingestion ledger [3/4] (KLAIR-3477) (@kevalshahtrilogy, Surtr)

#103 fix(mercy): pair accounting answers to open items by id, not by line (@kevalshahtrilogy, mercy)

#102 fix(mercy): the CI gate was counting mercy's own cancelled check as a failure (@kevalshahtrilogy, mercy)

#101 fix(mercy): the accounting pass must emit a shape the extractor accepts (@kevalshahtrilogy, mercy)

#100 fix(mercy): ask the arbiter to account for open items it left unanswered (@kevalshahtrilogy, mercy)

#1237 feat(admissions): render Deposit as three-state Paid/Waived/Not paid (#1226) (@vvp-trilogy, Aerie)

#1229 Count waived and advance deposits as paid via Finals site checklist (#1227) (@vvp-trilogy, Aerie)

#99 fix(mercy): make big-PR reviews converge (@kevalshahtrilogy, mercy)

#1235 Add N/A support for Buildout dates (@YibinLongTrilogy, Aerie)

#6 AI-691: Add local task creation workspace (@ashwanth1109, Shipyard)

#5 AI-690: Package Codex app-server runtime and connection diagnostics (@ashwanth1109, Shipyard)

#1231 feat(admissions): drop deprecated programs.isOpen field (AERIE-1179) (@caina-barbosa, Aerie)

#4 AI-689: Add embedded SQLite persistence and database explorer (@ashwanth1109, Shipyard)

#1233 feat(admissions): export enrolment dashboard as CSV (@benji-bizzell, Aerie)

@ashwanth1109 no labels



#17 — AI-705: Fix repository sync status and safe fast-forwarding

@ashwanth1109 no labels



#16 — AI-704: Persist project icons in SQLite

@ashwanth1109 no labels



#15 — AI-703: Add searchable project icon picker

@ashwanth1109 no labels



#14 — AI-702: Add performance diagnostics and non-blocking chat loading

@ashwanth1109 no labels



#13 — AI-701: Match Lumen Research chat styling

@ashwanth1109 no labels



#12 — AI-700: Persist Codex thread settings in Research chat

@ashwanth1109 no labels



#11 — AI-698: Add persisted Codex conversation to Research workflow

@ashwanth1109 no labels



#106 — fix(mercy): declare HEAD_SHA in the step that reads it — every review was failing

@kevalshahtrilogy CHANGES REQUESTED



#1726 — feat(heimdall): the triage table reader (SURTR-1040) [1/2]

@kevalshahtrilogy APPROVED MERCY-ALLOW-CRITICAL



#10 — AI-697: Persist Research workflow node state

@ashwanth1109 no labels



#9 — AI-696: Show fetched commit count after repository sync

@ashwanth1109 no labels



#8 — AI-695: Add repository sync and task detail views

@ashwanth1109 no labels



#105 — fix(mercy): give comment runs their own concurrency group

@kevalshahtrilogy no labels



#3 AI-688: Replace Shipyard bootstrap app icon (@ashwanth1109, Shipyard)

#2 AI-687: add semantic theme system (@ashwanth1109, Shipyard)

#1 AI-686: Shipyard Bootstrap (@ashwanth1109, Shipyard)

#1230 Make long skill descriptions scrollable (@YibinLongTrilogy, Aerie)

#98 fix(mercy): reap xlarge scratch as the grid runs, and stop a crashed re-view ending green (@kevalshahtrilogy, mercy)

#3733 feat(passive-investments): add cash asset type (@sanketghia, Klair)

#3708 feat(budget-bot): content-free operation lifecycle instrumentation + hermetic capacity harness (@marcusdAIy, Klair)

#279 feat(identity): map secondary runtime evidence (AI-662) (@marcusdAIy, trilogy-drones)

#97 fix(workflow): cancel-in-progress unconditionally, not just for pull_request (@kevalshahtrilogy, mercy)

#280 feat(identity): normalize event identity (AI-661) (@marcusdAIy, trilogy-drones)

#281 [draft-spec] AI-679: unattended spec-authoring draft (@marcusdAIy, trilogy-drones)

#1719 076-finops-paid-bu-mapping (@mwrshah, Surtr)

#1721 fix(repo): Add dark mode with light/dark/system toggle (@heimdall-keval-factory[bot], Surtr)

#1210 feat: ingest off-platform Skill invocation receipts (@caina-barbosa, Aerie)

#1717 feat(ai-spend): dedupe TF-routed OpenAI keys registered at the org grain (@kevalshahtrilogy, Surtr)

#1225 test(enrollments): harden SIS deposit + mapping tests (follow-up to #1215) (@vvp-trilogy, Aerie)

#1175 fix(dev-local): keep Windows env symlink and Clerk key for Next (@vvp-trilogy, Aerie)

#95 fix(config): disable the review-round cap by default (@kevalshahtrilogy, mercy)

#96 feat(mercy): xlarge review tier — clustered grid + adversarial verify for 10k-line PRs (@kevalshahtrilogy, mercy)

#94 feat(heimdall): file pipeline failures as Linear tickets, then work the queue (@kevalshahtrilogy, mercy)

#93 fix(heimdall): land the work-status change that merged into a branch (@kevalshahtrilogy, mercy)

#91 feat(heimdall): label-driven queue, and a chat feed written for a person (@kevalshahtrilogy, mercy)

#1219 fix(admissions): preserve Austin forecast totals in physical mode (@benji-bizzell, Aerie)

#1217 fix(dbt): make upstream-drift tripwires non-blocking; map mid_year_withdrawal to New Enrolled (@vvp-trilogy, Aerie)

#1220 feat(enrollments): build mart_enrollment_dtl on the SIS spine (#1215) (@vvp-trilogy, Aerie)

#1218 feat(documents): add global source registration (@benji-bizzell, Aerie)

#1211 Admissions pipeline: add a Community Commitment column before Application (#1209) (@vvp-trilogy, Aerie)

#104 — fix(mercy): don't carry an open item forward on a file the author just changed

@kevalshahtrilogy no labels



#7 — AI-693: Add multi-project repository management

@ashwanth1109 no labels



#1703 — feat(perplexity-usage-pipeline): Redshift writer, payload archive, ingestion ledger [3/4] (KLAIR-3477)

@kevalshahtrilogy APPROVED MERCY-ALLOW-CRITICAL



#103 — fix(mercy): pair accounting answers to open items by id, not by line

@kevalshahtrilogy CHANGES REQUESTED



#102 — fix(mercy): the CI gate was counting mercy's own cancelled check as a failure

@kevalshahtrilogy no labels



#101 — fix(mercy): the accounting pass must emit a shape the extractor accepts

@kevalshahtrilogy no labels



#100 — fix(mercy): ask the arbiter to account for open items it left unanswered

@kevalshahtriology CHANGES REQUESTED



#1237 — feat(admissions): render Deposit as three-state Paid/Waived/Not paid (#1226)

@vvp-trilogy APPROVED



#1229 — Count waived and advance deposits as paid via Finals site checklist (#1227)

@vvp-trilogy APPROVED



#99 — fix(mercy): make big-PR reviews converge

@kevalshahtriology CHANGES REQUESTED



#1235 — Add N/A support for Buildout dates

@YibinLongTrilogy APPROVED MERCY-ALLOW-CRITICAL



#6 — AI-691: Add local task creation workspace

@ashwanth1109 no labels



#5 — AI-690: Package Codex app-server runtime and connection diagnostics

@ashwanth1109 no labels



#3702 fix(board-doc): mint refresh ticket before opening stream (@marcusdAIy, Klair)

#1181 feat(platform-errors): isolate automatic triage runs (@benji-bizzell, Aerie)

#3706 test(budget-bot): add accessibility regression baseline (@marcusdAIy, Klair)

#276 feat(identity): add execution identity v1 core (AI-660) (@marcusdAIy, trilogy-drones)

#3704 feat(board-doc): adapt Q4 reviewed campaign email (KLAIR-3250) (@marcusdAIy, Klair)

#3717 fix(mercy): disable lifetime review cap for Klair (@marcusdAIy, Klair)

#1708 074-finops-invoice-mappings (@mwrshah, Surtr)

#3705 test(budget-bot): add mutation failure-injection goldens (@marcusdAIy, Klair)

#1707 fix(aerie-school-calendar-raw-sync): raise MAX_RETRIES so sheets fetch... (@heimdall-keval-factory[bot], Surtr)

#277 docs(security): document receipt mirror boundary (AI-668) (@marcusdAIy, trilogy-drones)

#278 refactor(mercy): extract deterministic verdict parser (AI-667) (@marcusdAIy, trilogy-drones)

#3715 fix(ai-budget): offer \$0-budgeted BUs in the Spend trajectory scope selector (@kevalshahtrilogy, Klair)

#1201 feat(enrollments): re-order and re-name matrix columns with an August 1 layout switch (@vvp-trilogy, Aerie)

#1699 feat(perplexity-usage-pipeline): secret loading and Perplexity API client (KLAIR-3477) [2/4] (@kevalshahtrilogy, Surtr)

#3713 feat(ai-budget): register "AI Renewals" as an assignable BU (@kevalshahtrilogy, Klair)

#1698 feat(perplexity-usage-pipeline): warehouse tables and runner scaffold (KLAIR-3477) [1/4] (@kevalshahtrilogy, Surtr)

#1695 fix(heimdall): point the Triage tile at the table that exists, and grant it (SURTR-1052) (@kevalshahtrilogy, Surtr)

#1694 072-grainne-pull-failure (@mwrshah, Surtr)

#1693 ci(mercy): tell mercy how this repo provisions tables and secrets (SURTR-1050) (@kevalshahtrilogy, Surtr)

#89 chore(consumers): sync Surtr's .mercy.yml with the provisioning convention (@kevalshahtrilogy, mercy)

#1207 fix(release): close final validation gaps (@benji-bizzell, Aerie)

#1688 ci(mercy): state harness_ref, so the pin actually pins (SURTR-1047) (@kevalshahtrilogy, Surtr)

#88 feat(mercy): read the PR discussion, weigh it against the code, answer it (AI-675) (@kevalshahtrilogy, mercy)

#87 fix(mercy): give the severity scale a middle, and blocking a bar (@kevalshahtrilogy, mercy)

#86 fix(mercy): count only mercy's own reviews as rounds, and test the seam (@kevalshahtrilogy, mercy)

#1206 fix(context): unblock Skill catalog rollout (@benji-bizzell, Aerie)

#85 fix(mercy): the open-items ledger never ran, and the CI gate never ran either (AI-673) (@kevalshahtrilogy, mercy)

#1231 — feat(admissions): drop deprecated programs.isOpen field (AERIE-1179)

@caina-barbosa APPROVED



#4 — AI-689: Add embedded SQLite persistence and database explorer

@ashwanth1109 no labels



#1233 — feat(admissions): export enrolment dashboard as CSV

@benji-bizzell APPROVED



#3 — AI-688: Replace Shipyard bootstrap app icon

@ashwanth1109 no labels



#2 — AI-687: add semantic theme system

@ashwanth1109 no labels



#1 — AI-686: Shipyard Bootstrap

@ashwanth1109 no labels



#1230 — Make long skill descriptions scrollable

@YibinLongTrilogy APPROVED



#98 — fix(mercy): reap xlarge scratch as the grid runs, and stop a crashed review ending green

@kevalshahtrilogy no labels



#3733 — feat(passive-investments): add cash asset type

@sanketghia APPROVED MERCY-ALLOW-CRITICAL



#3708 — feat(budget-bot): content-free operation lifecycle instrumentation + hermetic capacity harness

@marcusdAIy no labels



#279 — feat(identity): map secondary runtime evidence (AI-662)

@marcusdAIy APPROVED



#97 — fix(workflow): cancel-in-progress unconditionally, not just for pull_request

@kevalshahtrilogy no labels



#280 — feat(identity): normalize event identity (AI-661)

@marcusdAIy APPROVED



#281 — [draft-spec] AI-679: unattended spec-authoring draft

#3711 fix(spacex-valuation): reconcile August sales and hedge expiry (@sanketghia, Klair)

#84 feat(mercy): fan out large-PR reviews across lenses, and remember the last round (AI-672) (@kevalshahtrilogy, mercy)

#1204 fix(admissions): canonicalize public API program selectors (@benji-bizzell, Aerie)

#170 195-sindri-reliability (@mwrshah, Sindri)

#1202 fix(admissions): restore physical forecast provenance (@benji-bizzell, Aerie)

#1195 fix(admissions): stabilize Forecast identity and publication (@benji-bizzell, Aerie)

#1194 feat(context): give public Agent user-scoped Skills (@benji-bizzell, Aerie)

#1193 feat(context): add audience controls for Skills (@benji-bizzell, Aerie)

#3701 feat(budget-bot-addon): add section identity repair picker (@marcusdAly, Klair)

#275 docs(security): classify raw and mirrored traces (AI-666) (@marcusdAly, trilogy-drones)

#274 test(lifecycle): pin local cancel failure matrix (AI-665) (@marcusdAly, trilogy-drones)

#273 feat(review): add explicit intent/spec review (AI-663) (@marcusdAly, trilogy-drones)

#3700 fix(board-doc): gate fresh-Doc add-on actions on typed reconcile outcome (@marcusdAly, Klair)

#1198 feat: add authenticated installation receipt ingestion (AERIE-1852) (@caina-barbosa, Aerie)

#3698 feat(board-doc): add resumable Q4 campaign ledger (KLAIR-3247) (@marcusdAly, Klair)

#1682 fix(jotform-survey-sync): grant redshift-data:BatchExecuteStatement iam... (@heimdall-keval-factory[bot], Surtr)

#3699 fix(addon): use application-owned Drive token header (KLAIR-3495) (@marcusdAly, Klair)

#1197 docs(portfolio): document Backup Site operational-activity companion fields (@marcusdAly, Aerie)

#272 fix(security): enforce safe repository URLs (AI-628) (@marcusdAly, trilogy-drones)

#1196 docs(admissions): document non-retryable admissions_resource_refs_not_ready escalation (@marcusdAly, Aerie)

#271 [draft-spec] AI-668: unattended spec-authoring draft (@marcusdAly, trilogy-drones)

#266 docs(identity): pin execution identity v1 contract (AI-392) (@marcusdAly, trilogy-drones)

#270 docs(tasks): define receipt mirror boundary (AI-668) (@marcusdAly, trilogy-drones)

#267 docs(review): frame explicit intent review pilot (AI-663) (@marcusdAly, trilogy-drones)

#268 docs(tasks): frame bounded Q3 safety evidence (AI-665 AI-666) (@marcusdAly, trilogy-drones)

@marcusdAly no labels

#1719 — 076-finops-paid-bu-mapping

@mwrshah APPROVED

#1721 — fix(repo): Add dark mode with light/dark/system toggle

@heimdall-keval-factory[bot] APPROVED

AUTOMATED PR

#1210 — feat: ingest off-platform Skill invocation receipts

@caina-barbosa APPROVED MERCY-ALLOW-CRITICAL

#1717 — feat(ai-spend): dedupe TF-routed OpenAI keys registered at the org grain

@kevalshahtrilogy APPROVED HEIMDALL-DRIVEN

#1225 — test(enrollments): harden SIS deposit + mapping tests (follow-up to #1215)

@vvp-trilogy APPROVED

#1175 — fix(dev-local): keep Windows env symlink and Clerk key for Next

@vvp-trilogy no labels

#95 — fix(config): disable the review-round cap by default

@kevalshahtrilogy no labels

#96 — feat(mercy): xlarge review tier — clustered grid + adversarial verify for 10k-line PRs

@kevalshahtrilogy no labels

#94 — feat(heimdall): file pipeline failures as Linear tickets, then work the queue

@kevalshahtrilogy CHANGES REQUESTED

#93 — fix(heimdall): land the work-status change that merged into a branch

@kevalshahtrilogy CHANGES REQUESTED

#91 — feat(heimdall): label-driven queue, and a chat feed written for a person

@kevalshahtrilogy CHANGES REQUESTED

#1219 — fix(admissions): preserve Austin forecast totals in physical mode

@benji-bizzell CHANGES REQUESTED

#269 docs(tasks): define Mercy verdict parser extraction (AI-667) (@marcusdAIy, trilogy-drones)

#3697 feat(ai-spend): wire Perplexity per-person spend into fct_ai_spend (KLAIR-3477) (@kevalshahtrilogy, Klair)

#1192 feat: add private supplied-slug Skill distribution (AERIE-1851) (@caina-barbosa, Aerie)

#83 fix(heimdall): land the chat fixes that merged into a branch, not into main (@kevalshahtrilogy, mercy)

#3693 fix(board-doc): make embed-safe clone setup revision-safe and cleanup-safe (@marcusdAIy, Klair)

#1677 fix(aws-spend): record Q3 mapping for three newly active accounts (@caina-barbosa, Surtr)

#1671 fix(aws-bedrock-token-metrics): Surface unmatched failures in known_fai... (@heimdall-keval-factory[bot], Surtr)

#1188 feat: add dormant device authorization scaffolding (AERIE-1850) (@caina-barbosa, Aerie)

#1668 fix(saas-budgeting-pipeline): Report partial ingest failure instead of... (@heimdall-keval-factory[bot], Surtr)

#78 feat(heimdall): narrate every stage into Chat, one line at a time (@kevalshahtrilogy, mercy)

#1663 fix(mart-education-quickbooks-refresh): Scale audit ceiling and gate sk... (@heimdall-keval-factory[bot], Surtr)

#77 fix(heimdall): the steward owned nothing, and said nothing about it (@kevalshahtrilogy, mercy)

#265 docs(security): define worker isolation standard (AI-632) (@marcusdAIy, trilogy-drones)

#1190 Deprecate programs.isOpen; use schoolStatus from the EduCRM all-program mart (AERIE-1179) (@vvp-trilogy, Aerie)

#76 feat(heimdall): let a repo pin the agent model, not just the runtime (@kevalshahtrilogy, mercy)

#1651 feat(heimdall): actually run the code during verification, on 4 vCPU (@kevalshahtrilogy, Surtr)

#75 feat(heimdall): let the caller choose the runner (@kevalshahtrilogy, mercy)

#1185 feat(forge): unify authored object detail surfaces (@benji-bizzell, Aerie)

#3692 fix(spacex-valuation): show skeleton during initial quote load (@sanketghia, Klair)

#262 build(security): pin workflow dependencies (AI-631) (@marcusdAIy, trilogy-drones)

#264 test(security): scan scripts for raw agent calls (AI-629) (@marcusdAIy, trilogy-drones)

#263 fix(security): make local evidence files private (AI-633) (@marcusdAIy, trilogy-drones)

#1180 fix(add-aerie-skill): add unpublished secure local foundation (AERIE-1849) (@caina-barbosa, Aerie)

#1653 fix(jotform-survey-sync): scope retry queue existence check (@marcusdAIy, Surtr)

#3691 feat(schools-report): match QTD report formatting (@ashwanth1109, Klair)

#1217 — fix(dbt): make upstream-drift tripwires non-blocking; map mid_year_withdrawal to New Enrolled

@vvp-trilogy no labels

#1220 — feat(enrollments): build mart_enrollment_dtl on the SIS spine (#1215)

@vvp-trilogy CHANGES REQUESTED

#1218 — feat(documents): add global source registration

@benji-bizzell no labels

#1211 — Admissions pipeline: add a Community Commitment column before Application (#1209)

@vvp-trilogy CHANGES REQUESTED

#3702 — fix(board-doc): mint refresh ticket before opening stream

@marcusdAIy no labels

#1181 — feat(platform-errors): isolate automatic triage runs

@benji-bizzell CHANGES REQUESTED

#3706 — test(budget-bot): add accessibility regression baseline

@marcusdAIy APPROVED

#276 — feat(identity): add execution identity v1 core (AI-660)

@marcusdAIy no labels

#3704 — feat(board-doc): adapt Q4 reviewed campaign email (KLAIR-3250)

@marcusdAIy APPROVED

#3717 — fix(mercy): disable lifetime review cap for Klair

@marcusdAIy APPROVED

#1708 — 074-finops-invoice-mappings

@mwrshah APPROVED

#3705 — test(budget-bot): add mutation failure-injection goldens

@marcusdAIy APPROVED

#1707 — fix(aerie-school-calendar-raw-sync): raise MAX_RETRIES so sheets fetch...

#3684 feat(addon): make section rename crash-safe (KLAIR-3231) (@marcusdAIy, Klair)

#1611 fix(netsuite-saved-search-refresh): restore Redshift-compatible coverage refresh (@ashwanth1109, Surtr)

#1187 fix(api): honor preferred Drive credential (@benji-bizzell, Aerie)

#261 fix(security): close mirrored trace schema (AI-630) (@marcusdAIy, trilogy-drones)

#1186 fix(chat): preserve document search evidence rows (@benji-bizzell, Aerie)

#1184 fix(chat): reserve stable layout for Skill edit-validation status (@marcusdAIy, Aerie)

#3679 feat(board-doc): wire additive copy-first embed-safe clone path (@marcusdAIy, Klair)

#1183 fix(document-intelligence): skip ambiguous global search documents (@benji-bizzell, Aerie)

#3685 fix(board-doc): make add-on reconcile Doc-canonical (KLAIR-3242) (@marcusdAIy, Klair)

#1174 fix(document-intelligence): harden Global document verification (@benji-bizzell, Aerie)

#3690 KLAIR-3472 Add on-demand single-school report generation (@ashwanth1109, Klair)

#1178 feat(document-intelligence): add Global document API uploads (@benji-bizzell, Aerie)

#1177 fix(documents): restore delegated upload identity (@benji-bizzell, Aerie)

#1652 feat(quickbooks): expand active raw-sync company inventory safely (@ashwanth1109, Surtr)

#1649 feat(heimdall): factory widget on top, with sources and stage click-through (@kevalshahtrilogy, Surtr)

#1176 feat(chat): render site document searches as inline traces (AERIE-1957) (@caina-barbosa, Aerie)

#1648 fix(heimdall): accept every mode the workflow emits (@kevalshahtrilogy, Surtr)

#1642 fix(repo): Use dropdown pickers and paginate gateway lists (@heimdall-keval-factory[bot], Surtr)

#1647 fix(mercy): correct heimdall's bot identity in the allowlist (@kevalshahtrilogy, Surtr)

#74 feat(heimdall): PR opens with a sentence a person can read (@kevalshahtrilogy, mercy)

#72 fix(heimdall): a sweep with nothing to do must not go red (@kevalshahtrilogy, mercy)

#73 feat(mercy): heimdall's own PRs carry the critical-path grant (@kevalshahtrilogy, mercy)

#3687 fix(mercy): correct heimdall's bot identity in the allowlist (@kevalshahtrilogy, Klair)

#1646 fix(heimdall): forward the org-level OpenAI key to the codex runtime (@kevalshahtrilogy, Surtr)

#71 test(heimdall): exercise both model-resolution blocks, not just the first (@kevalshahtrilogy, mercy)

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

▶
#277 — docs(security): document receipt mirror boundary (AI-668)

@marcusdAIy no labels

▶
#278 — refactor(mercy): extract deterministic verdict parser (AI-667)

@marcusdAIy no labels

▶
#3715 — fix(ai-budget): offer \$0-budgeted BUs in the Spend trajectory scope selector

@kevalshahtrilogy APPROVED

▶
#1201 — feat(enrollments): re-order and re-name matrix columns with an August 1 layout switch

@vvp-trilogy APPROVED

▶
#1699 — feat(perplexity-usage-pipeline): secret loading and Perplexity API client (KLAIR-3477) [2/4]

@kevalshahtrilogy APPROVED MERCY-ALLOW-CRITICAL

▶
#3713 — feat(ai-budget): register "AI Renewals" as an assignable BU

@kevalshahtrilogy APPROVED

▶
#1698 — feat(perplexity-usage-pipeline): warehouse tables and runner scaffold (KLAIR-3477) [1/4]

@kevalshahtrilogy APPROVED MERCY-ALLOW-CRITICAL

▶
#1695 — fix(heimdall): point the Triage tile at the table that exists, and grant it (SURTR-1052)

@kevalshahtrilogy APPROVED MERCY-ALLOW-CRITICAL

▶
#1694 — 072-grainne-pull-failure

@mwrshah APPROVED

▶
#1693 — ci(mercy): tell mercy how this repo provisions tables and secrets (SURTR-1050)

@kevalshahtrilogy APPROVED

▶
#89 — chore(consumers): sync Surtr's .mercy.yml with the provisioning convention

@kevalshahtrilogy APPROVED

▶
#1207 — fix(release): close final validation gaps

@benji-bizzell APPROVED

#70 feat(heimdall): make the agent runtime flippable by repo variable (@kevalshahtrilogy, mercy)

#69 feat(heimdall): open unverified fixes ready for review, not as drafts (@kevalshahtrilogy, mercy)

#68 fix(heimdall): a failed publish is not 'no code change' (@kevalshahtrilogy, mercy)

#67 fix(heimdall): the fix branch must be unique per ticket (@kevalshahtrilogy, mercy)

#66 feat(heimdall): say when a team is routed nowhere (@kevalshahtrilogy, mercy)

#65 feat(heimdall): work the queue in the board's order, not the alphabet (@kevalshahtrilogy, mercy)

#64 fix(mercy): create the grant label before adding it (@kevalshahtrilogy, mercy)

#1632 fix(heimdall): let stale open work stop suppressing a pipeline's triage (@kevalshahtrilogy, Surtr)

#63 fix(mercy): make the critical-path grant standing, not per-run (@kevalshahtrilogy, mercy)

#62 fix(heimdall): repair the two faults the first live queue sweep found (@kevalshahtrilogy, mercy)

#61 fix(heimdall): quote untrusted values in the harness's own prompt notes (@kevalshahtrilogy, mercy)

#60 fix(heimdall): quote untrusted text in the converse and revise prompts (@kevalshahtrilogy, mercy)

#59 fix(heimdall): re-land #58 onto main (it merged to the orphaned stack branch) (@kevalshahtrilogy, mercy)

#1627 feat(heimdall): tell the queue where a claimed ticket goes (@kevalshahtrilogy, Surtr)

#57 refactor(heimdall): the agent works the Linear ticket directly (@kevalshahtrilogy, mercy)

#1173 feat(skill-telemetry): add dormant Slice 3 contracts (@caina-barbosa, Aerie)

#1172 fix(document-intelligence): guard Global search source URLs (@benji-bizzell, Aerie)

#1171 fix(release): clear Global Documents deployment gates (@benji-bizzell, Aerie)

#1169 fix(skills): revert governed Ops Skills import foundations (@benji-bizzell, Aerie)

#1167 fix(portfolio): scroll Documents tab body (AERIE-1935) (@caina-barbosa, Aerie)

#1166 feat(document-intelligence): add portfolio-wide Global documents (@benji-bizzell, Aerie)

#1631 Schedule QuickBooks Compensation after upstream consumers (@YibinLongTrilogy, Surtr)

#1630 Fix QuickBooks Expense AI research requests (@YibinLongTrilogy, Surtr)

#1165 fix(chat): keep Skill review errors with actions (AERIE-1954) (@caina-barbosa, Aerie)

#1629 fix(surtr-781): correct Redshift DDL syntax (@marcusdAly, Surtr)

►

#1688 — ci(mercy): state harness_ref, so the pin actually pins (SURTR-1047)

@kevalshahtrilogy APPROVED MERCY-ALLOW-CRITICAL

►

#88 — feat(mercy): read the PR discussion, weigh it against the code, answer it (AI-675)

@kevalshahtrilogy no labels

►

#87 — fix(mercy): give the severity scale a middle, and blocking a bar

@kevalshahtrilogy no labels

►

#86 — fix(mercy): count only mercy's own reviews as rounds, and test the seam

@kevalshahtriology CHANGES REQUESTED

►

#1206 — fix(context): unblock Skill catalog rollout

@benji-bizzell APPROVED

►

#85 — fix(mercy): the open-items ledger never ran, and the CI gate never ran either (AI-673)

@kevalshahtriology CHANGES REQUESTED

►

#3711 — fix(spacex-valuation): reconcile August sales and hedge expiry

@sanketghia APPROVED

►

#84 — feat(mercy): fan out large-PR reviews across lenses, and remember the last round (AI-672)

@kevalshahtriology no labels

►

#1204 — fix(admissions): canonicalize public API program selectors

@benji-bizzell CHANGES REQUESTED

►

#170 — 195-sindri-reliability

@mwrshah no labels

►

#1202 — fix(admissions): restore physical forecast provenance

@benji-bizzell APPROVED

►

#1195 — fix(admissions): stabilize Forecast identity and publication

@benji-bizzell APPROVED

►

#1194 — feat(context): give public Agent user-scoped Skills

@benji-bizzell APPROVED

►

#259 docs(queue): frame first Q3 security batch (@marcusdAIy, trilogy-drones)

#260 [draft-spec] AI-628: unattended spec-authoring draft (@marcusdAIy, trilogy-drones)

#1147 Add governed EduOps Skills import foundations (@YibinLongTrilogy, Aerie)

#1626 fix(jotform-survey-sync): stop exhausting Jotform's daily API quota (@marcusdAIy, Surtr)

#1628 feat(heimdall): accept cost_source in the telemetry schema (@kevalshahtrilogy, Surtr)

#1625 fix(heimdall): factory-status Phase 0 uses cumulative triggered count (SURTR-996) (@kevalshahtrilogy, Surtr)

#258 [draft-spec] AI-590: unattended spec-authoring draft (@marcusdAIy, trilogy-drones)

- #1193 — feat(context): add audience controls for Skills

@benji-bizzell

APPROVED
- ▶
- #3701 — feat(budget-bot-addon): add section identity repair picker

@marcusdAIy

APPROVED
- ▶
- #275 — docs(security): classify raw and mirrored traces (AI-666)

@marcusdAIy

APPROVED
- ▶
- #274 — test(lifecycle): pin local cancel failure matrix (AI-665)

@marcusdAIy

APPROVED
- ▶
- #273 — feat(review): add explicit intent/spec review (AI-663)

@marcusdAIy

no labels
- ▶
- #3700 — fix(board-doc): gate fresh-Doc add-on actions on typed reconcile outcome

@marcusdAIy

APPROVED
- ▶
- #1198 — feat: add authenticated installation receipt ingestion (AERIE-1852)

@caina-barbosa

APPROVED
- ▶
- #3698 — feat(board-doc): add resumable Q4 campaign ledger (KLAIR-3247)

@marcusdAIy

APPROVED
- ▶
- #1682 — fix(jotform-survey-sync): grant redshift-data:BatchExecuteStatement iam...

@heimdall-keval-factory[bot]

APPROVED
- AUTOMATED PR
- ▶
- #3699 — fix(addon): use application-owned Drive token header (KLAIR-3495)

@marcusdAIy

APPROVED
- ▶
- #1197 — docs(portfolio): document Backup Site operational-activity companion fields

@marcusdAIy

APPROVED
- ▶
- #272 — fix(security): enforce safe repository URLs (AI-628)

@marcusdAIy

APPROVED
- ▶
- #1196 — docs(admissions): document non-retryable admissions_resource_refs_not_ready

escalation

@marcusdAIy APPROVED



#271 — [draft-spec] AI-668: unattended spec-authoring draft

@marcusdAIy no labels



#266 — docs(identity): pin execution identity v1 contract (AI-392)

@marcusdAIy APPROVED



#270 — docs(tasks): define receipt mirror boundary (AI-668)

@marcusdAIy APPROVED



#267 — docs(review): frame explicit intent review pilot (AI-663)

@marcusdAIy APPROVED



#268 — docs(tasks): frame bounded Q3 safety evidence (AI-665 AI-666)

@marcusdAIy APPROVED



#269 — docs(tasks): define Mercy verdict parser extraction (AI-667)

@marcusdAIy APPROVED



#3697 — feat(ai-spend): wire Perplexity per-person spend into fct_ai_spend (KLAIR-3477)

@kevalshahtrilogy APPROVED



#1192 — feat: add private supplied-slug Skill distribution (AERIE-1851)

@caina-barbosa APPROVED



#83 — fix(heimdall): land the chat fixes that merged into a branch, not into main

@kevalshahtrilogy no labels



#3693 — fix(board-doc): make embed-safe clone setup revision-safe and cleanup-safe

@marcusdAIy APPROVED



#1677 — fix(aws-spend): record Q3 mapping for three newly active accounts

@caina-barbosa APPROVED



#1671 — fix(aws-bedrock-token-metrics): Surface unmatched failures in known_fai...

@heimdall-keval-factory[bot] APPROVED

AUTOMATED PR

- #1188 — feat: add dormant device authorization scaffolding (AERIE-1850)

@caina-barbosa

APPROVED
- #1668 — fix(saas-budgeting-pipeline): Report partial ingest failure instead of...

@heimdall-keval-factory[bot]

APPROVED

AUTOMATED PR
- #78 — feat(heimdall): narrate every stage into Chat, one line at a time

@kevalshahtrilogy

CHANGES REQUESTED
- #1663 — fix(mart-education-quickbooks-refresh): Scale audit ceiling and gate sk...

@heimdall-keval-factory[bot]

APPROVED

AUTOMATED PR
- #77 — fix(heimdall): the steward owned nothing, and said nothing about it

@kevalshahtrilogy no labels
- #265 — docs(security): define worker isolation standard (AI-632)

@marcusdAIy no labels
- #1190 — Deprecate programs.isOpen; use schoolStatus from the EduCRM all-program mart (AERIE-1179)

@vvp-trilogy

APPROVED
- #76 — feat(heimdall): let a repo pin the agent model, not just the runtime

@kevalshahtrilogy no labels
- #1651 — feat(heimdall): actually run the code during verification, on 4 vCPU

@kevalshahtrilogy

APPROVED

MERCY-ALLOW-CRITICAL
- #75 — feat(heimdall): let the caller choose the runner

@kevalshahtrilogy no labels
- #1185 — feat(forge): unify authored object detail surfaces

@benji-bizzell no labels
- #3692 — fix(spacex-valuation): show skeleton during initial quote load

@sanketghia

APPROVED

- #262 — build(security): pin workflow dependencies (AI-631)

@marcusdAIy no labels
- #264 — test(security): scan scripts for raw agent calls (AI-629)

@marcusdAIy APPROVED
- #263 — fix(security): make local evidence files private (AI-633)

@marcusdAIy APPROVED
- #1180 — fix(add-aerie-skill): add unpublished secure local foundation (AERIE-1849)

@caina-barbosa APPROVED
- #1653 — fix(jotform-survey-sync): scope retry queue existence check

@marcusdAIy no labels
- #3691 — feat(schools-report): match QTD report formatting

@ashwanth1109 APPROVED
- #3684 — feat(addon): make section rename crash-safe (KLAIR-3231)

@marcusdAIy APPROVED
- #1611 — fix(netsuite-saved-search-refresh): restore Redshift-compatible coverage refresh

@ashwanth1109 APPROVED
- #1187 — fix(api): honor preferred Drive credential

@benji-bizzell APPROVED
- #261 — fix(security): close mirrored trace schema (AI-630)

@marcusdAIy APPROVED
- #1186 — fix(chat): preserve document search evidence rows

@benji-bizzell APPROVED
- #1184 — fix(chat): reserve stable layout for Skill edit-validation status

@marcusdAIy APPROVED
- #3679 — feat(board-doc): wire additive copy-first embed-safe clone path

@marcusdAIy APPROVED

- #1183 — fix(document-intelligence): skip ambiguous global search documents

@benji-bizzell

APPROVED
- #3685 — fix(board-doc): make add-on reconcile Doc-canonical (KLAIR-3242)

@marcusdAIy

APPROVED
- #1174 — fix(document-intelligence): harden Global document verification

@benji-bizzell

APPROVED

►

#3690 — KLAIR-3472 Add on-demand single-school report generation

@ashwanth1109

APPROVED

►

#1178 — feat(document-intelligence): add Global document API uploads

@benji-bizzell

APPROVED

►

#1177 — fix(documents): restore delegated upload identity

@benji-bizzell

APPROVED

►

#1652 — feat(quickbooks): expand active raw-sync company inventory safely

@ashwanth1109

APPROVED

►

#1649 — feat(heimdall): factory widget on top, with sources and stage click-through

@kevalshahtrilogy

APPROVED

►

#1176 — feat(chat): render site document searches as inline traces (AERIE-1957)

@caina-barbosa

APPROVED

►

#1648 — fix(heimdall): accept every mode the workflow emits

@kevalshahtrilogy

APPROVED

►

#1642 — fix(repo): Use dropdown pickers and paginate gateway lists

@heimdall-keval-factory[bot]

APPROVED

AUTOMATED PR

MERCY-ALLOW-CRITICAL

►

#1647 — fix(mercy): correct heimdall's bot identity in the allowlist

@kevalshahtrilogy

APPROVED

MERCY-ALLOW-CRITICAL

►

- #74 — feat(heimdall): PR opens with a sentence a person can read

@kevalshahtrilogy no labels
- ▶
- #72 — fix(heimdall): a sweep with nothing to do must not go red

@kevalshahtrilogy no labels
- ▶
- #73 — feat(mercy): heimdall's own PRs carry the critical-path grant

@kevalshahtrilogy no labels
- ▶
- #3687 — fix(mercy): correct heimdall's bot identity in the allowlist

@kevalshahtrilogy APPROVED
- ▶
- #1646 — fix(heimdall): forward the org-level OpenAI key to the codex runtime

@kevalshahtrilogy APPROVED MERCY-ALLOW-CRITICAL
- ▶
- #71 — test(heimdall): exercise both model-resolution blocks, not just the first

@kevalshahtrilogy APPROVED
- ▶
- #70 — feat(heimdall): make the agent runtime flippable by repo variable

@kevalshahtrilogy CHANGES REQUESTED
- ▶
- #69 — feat(heimdall): open unverified fixes ready for review, not as drafts

@kevalshahtrilogy no labels
- ▶
- #68 — fix(heimdall): a failed publish is not 'no code change'

@kevalshahtrilogy no labels
- ▶
- #67 — fix(heimdall): the fix branch must be unique per ticket

@kevalshahtrilogy no labels
- ▶
- #66 — feat(heimdall): say when a team is routed nowhere

@kevalshahtrilogy APPROVED
- ▶
- #65 — feat(heimdall): work the queue in the board's order, not the alphabet

@kevalshahtrilogy APPROVED
- ▶
- #64 — fix(mercy): create the grant label before adding it

@kevalshahtrilogy no labels

@caina-barbosa APPROVED

►
#1166 — feat(document-intelligence): add portfolio-wide Global documents

@benji-bizzell APPROVED

►
#1631 — Schedule QuickBooks Compensation after upstream consumers

@YibinLongTrilogy APPROVED

►
#1630 — Fix QuickBooks Expense AI research requests

@YibinLongTrilogy APPROVED

►
#1165 — fix(chat): keep Skill review errors with actions (AERIE-1954)

@caina-barbosa APPROVED

►
#1629 — fix(surtr-781): correct Redshift DDL syntax

@marcusdAIy APPROVED

►
#259 — docs(queue): frame first Q3 security batch

@marcusdAIy APPROVED

►
#260 — [draft-spec] AI-628: unattended spec-authoring draft

@marcusdAIy CHANGES REQUESTED

►
#1147 — Add governed EduOps Skills import foundations

@YibinLongTrilogy no labels

►
#1626 — fix(jotform-survey-sync): stop exhausting Jotform's daily API quota

@marcusdAIy APPROVED

►
#1628 — feat(heimdall): accept cost_source in the telemetry schema

@kevalshahtrilogy APPROVED

►
#1625 — fix(heimdall): factory-status Phase 0 uses cumulative triggered count (SURTR-996)

@kevalshahtrilogy APPROVED

►
#258 — [draft-spec] AI-590: unattended spec-authoring draft

@marcusdAIy no labels



WEEK IN REVIEW

PRODUCTION RELEASE



ENGINEER SPOTLIGHT

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



#26 — AI-715: Add Research-to-Linear Ticket workflow

@ashwanth1109 no labels



#29 — AI-717: Persist Implement PR links and open them externally

@ashwanth1109 no labels



#110 — feat(heimdall): put the logs on the ticket, keep queue runs out of GitHub Issues

@kevalshahtrilogy CHANGES REQUESTED



#1241 — fix(admissions): resolve Community Commitment deposit from EduCRM deposit_paid_date (#1240)

@vvp-trilogy APPROVED



#1243 — feat(sindri): add bounded Skill adoption read model

@caina-barbosa CHANGES REQUESTED



#1727 — feat(heimdall): the factory board's data model (SURTR-1040) [2/2]

@kevalshahtrilogy APPROVED MERCY-ALLOW-CRITICAL

213 PRs, SEVEN Repos Ablaze: The Builder Team Posts a Number That Requires a Calculator With Extra Buttons

Kevalshahtrilogy alone outshipped entire companies with 71 PRs, while Ashwanth Sitaraman quietly turned Shipyard into his personal skateboard park.

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

Folks, pull up a chair and loosen your belt, because the Builder Team just posted 213 pull requests in seven days across SEVEN active repositories, and I have counted every single one of them TWICE. Aerie led the charge with 50 PRs, mercy right on its heels with 47, Surtr posting a muscular 39, the brand-new Shipyard repo already cranking out 28 in its infancy, trilogy-drones humming along with 24, Klair contributing a sturdy 23, and even little Sindri chipping in 2 PRs of pure, concentrated effort. This is not a sprint. This is a velocity event that should require a permit.

Let's talk names. @kevalshahtrilogy did not just lead the leaderboard — he lapped it, posting a staggering 71 PRs spanning Surtr (#1727, #1746, #1742), mercy (#110, #109, #108, #107), and beyond. That is not a work week, that is a work MONTH compressed through sheer will. @marcusDAIy followed with 44 PRs of his own, a number that would be a career highlight for a lesser engineer but here is just Tuesday. @benji-bizzell logged 22, @caina-barbosa delivered 13 including Aerie's #1243 and #1242, @vvp-trilogy notched 10 including the admissions fix in #1241, our tireless bot friend @heimdall-kevalfactory[bot] clocked 8 (yes, even the robots are winning), and @mwrshah rounded things out with 5, including Sindri's #176.

Now. Ashwanth. Thirty-two PRs, nearly all of them stacked back-to-back in Shipyard like dominoes falling in a hurricane: #29, #27, #26, #25, #24, and reaching back further into #16 through #23. The man built an entire artifact-template system, a Research-to-Linear workflow, and multi-project task selection in the time it takes most engineers to write a JIRA ticket describing the problem. "I don't review my own diffs, I just remember writing them," Ashwanth allegedly told me this week, which is either the most confident sentence ever spoken or a cry for help disguised as a flex. When I asked if anyone else on the team could actually parse PR #26's scope, he reportedly just said "why would that matter." Legend. Borderline concerning. Legend.

Now to the Overflow Desk, where Mac Donnelly's narrative simply couldn't contain the volume. Surtr's #1754 quietly patched exception messaging into failure payloads — unglamorous, essential, ignored by the big story, immortalized here. Mercy's #109 was just "chore: ruff format," and I want you to sit with that: a full pull request dedicated to formatting, and it STILL counts toward the 213. And #1241 in Aerie fixed a Community Commitment deposit resolution bug that could've caused real headaches — buried by narrative, rescued by numbers.

Morale, as always, has never been higher. Seven repos. Two hundred thirteen PRs. One new Shipyard standing tall. The Builder Team isn't just winning — they're setting the bar somewhere near the stratosphere and daring anyone to find a ladder.

Two Newspapers, One Billionaire, Zero Interviews

Forbes calls it a software sweatshop. The Wall Street Journal calls it a home for orphans. Both are describing the same machine.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — Joe Liemandt has spent three decades avoiding cameras, and this week two of the country's most prestigious business publications tried to explain why it hasn't mattered.

Forbes went first, with a profile of the [mysterious tech billionaire](#) who built two fortunes off the same insight: enterprise software companies nobody wants can be bought cheap, stripped of local payroll, and staffed with Crossover's global talent pool until the margins hit 75%. The word Forbes chose was sweatshop.

The Journal, working the same territory, chose a gentler frame. Its dispatch on ESW Capital describes an orphanage

— small software companies, abandoned by venture capital, given a home. Same acquisitions. Same math. Different verb.

Both stories are circling a shape that's been visible in Portland for years. Jive Software, once the pride of the city's tech scene, sold to Aurea — an ESW property — for roughly half what it fetched at its 2011 peak, a marker GeekWire has tracked since the deal closed. It is the ESW playbook rendered in a single balance sheet: buy distressed, absorb into the portfolio, raise support pricing on customers who are too locked in to leave.

Huddle followed a similar arc — a UK collaboration platform once flush with venture money, sold for \$89 million to a

private equity buyer, a fraction of what it raised. The names change; the discount to peak value does not.

Forrester, meanwhile, is now advising enterprise buyers on what to do next about their "customer advocacy platform" — corporate language for the software these acquisitions produce, and a quiet acknowledgment that customers of consolidated platforms increasingly need an exit plan before someone else writes one for them.

Two publications, two adjectives, one number: 75% EBITDA margin. Trilogy calls it proof of efficiency. The customers footing the renewal invoices may use a different word.

The 26-Month Job Skyvera Did in 30 Days — And What It Says About the Playbook

Behind Skyvera's back-to-back telecom acquisitions lies a familiar Trilogy signature: buy the asset, then let the machines do the heavy lifting.

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — Two acquisitions in quick succession rarely make headlines on their own. But when Skyvera closed its purchase of [CloudSense](#), the Salesforce-native CPQ platform for telcos, and then swallowed STL's divested telecom products group — digital BSS, monetization, optical networking, analytics — the pattern was hard to miss for anyone who's followed this portfolio closely. And this is where it gets interesting.

Within a month of the CloudSense deal closing, Skyvera announced that CloudSense had certified all 13 of its CPQ APIs to TM Forum compliance standards — a process the industry typically measures in 26 months. Not weeks. Months, plural, times twenty-six. Skyvera says the acceleration came through [a strategic AI partnership](#) layered onto the existing engineering team.

I can't say who told me this, but a source familiar with Skyvera's integration process described the sequencing as 'not an accident' — acquisitions timed so that compliance and standardization work, historically the slowest and most expensive part of any telecom software integration, gets compressed almost immediately after close. If that's true, it changes how you should read every future Skyvera deal announcement.

The STL divested assets — which bring digital BSS functionality, optical networking, and analytics into the fold — fit neatly beside CloudSense's B2B, B2B2X, and wholesale quoting capabilities. Put together, Skyvera now owns a noticeably wider slice of the telecom stack, from quote-to-cash to network monetization, all under one roof in Austin.

None of this is unusual for the ESW Capital playbook: acquire cheap, integrate fast, extract margin through automation rather than headcount. What's notable is the speed. Twenty-six months to one is not a productivity gain — it's a different operating model entirely. Whether TM Forum's other certifying vendors take note is, for now, an open question. But I suspect they already have.

In Austin, the Question Isn't Whether AI Teaches Your Kid — It's Who Loves Them While It Happens

Alpha School's latest push to answer its most persistent critique doubles as a meditation on what, exactly, we still need humans for.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — There is a question that has trailed Alpha School since its earliest press coverage, asked in tones ranging from genuine curiosity to thinly veiled alarm: does this place actually replace teachers with software? It is, on its face, a fair question for a school built on the premise that AI-driven adaptive apps can deliver a full academic curriculum in two hours a day. And it is a question Alpha has now taken the unusual step of answering directly, in a blog post that reads less like marketing copy and more like a rebuttal brief.

The answer, per [the school's own accounting](#), is no — and the distinction the school draws is worth sitting with. AI, in Alpha's telling, handles what is repeatable: drilling math facts, adjusting pacing, flagging where a ten-year-old has stalled on fractions. What it does not handle, and what Alpha insists it was never designed to handle, is the human work — the full-time "Guides" whose job is motivation, relationship-building, and the accumulated, unautomatable knowledge of who a particular kid is on a particular Tuesday.

That framing extends into a companion series Alpha has been running under the banner "Teach Your Kid What School Doesn't" — installments on emotional regulation, life skills, and creativity that function almost as a parenting curriculum unto themselves. Part Four takes on the proposition that big feelings are not a problem to be managed away but a skill to be built; Part Five argues, per [the school's latest entry](#), that every child arrives with creative instincts that conventional schooling tends to flatten rather than cultivate.

Taken together, the material amounts to something more interesting than a defensive PR maneuver. It is Alpha, and by extension Joe Liemandt's broader education bet, making an explicit case for where the line between machine and human competence actually falls — a line that most of the AI industry gestures at but rarely bothers to draw in public. Whether parents buy it may determine whether the model scales beyond Austin, Miami, and Brownsville — or stays a boutique experiment for those who can afford \$40,000 tuition and the argument that comes with it.

The Mind That Already Decided: How Language Models Take In New Evidence

A new study peers into the moment an AI changes its mind — and finds something unsettlingly familiar about how conviction bends, or doesn't, under the weight of proof.

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

PALO ALTO, CALIF. — There is a particular kind of intimacy in watching a mind — any mind — encounter a fact it did not expect. Something shifts, or refuses to. A belief bends toward the light of new evidence, or hardens against it. We have studied this in juries, in physicians, in ourselves. Now researchers are studying it in machines that increasingly do our reasoning alongside us.

A paper released this week, [Evidence Integration in Large Language Models](#), asks a question that sounds almost philosophical for a field obsessed with benchmarks: when an LLM has already begun forming a judgment and new evidence arrives — from a search tool, a second AI, a human correction — how does that evidence actually get absorbed? The researchers propose a distributional theory, treating the model's prior stance not as a fixed point but as a probability cloud that evidence reshapes. It matters enormously, because these systems are no longer solitary oracles; they are increasingly nodes in pipelines of retrieval and delegation, forming tentative conclusions and then revising — or failing to revise — them in light of what other systems tell them.

It is worth remembering how fragile even our most basic measurements can be. A companion paper this week on dependency distance in language — the average length of grammatical dependencies within a sentence — found that estimates long treated as fixed properties of a language shift meaningfully depending on which corpus you happen to measure

them from, across 38 language pairs in the [Universal Dependencies treebank collection](#). The lesson generalizes: what looks like a stable fact about a mind, human or artificial, is often a fact about the instrument doing the looking.

Stanford's Human-Centered AI Institute has been arguing something adjacent — that AI's growing role in scientific discovery should sharpen, not replace, human judgment. Evidence integration research suggests why that vigilance matters: a model's confidence can look identical whether it was earned by good evidence or manufactured by a persuasive one. Knowing the difference may be the most human task left.

In the Shadow of the Server Farm: A Study in Corporate Camouflage

Deep in America's data-center sunbelt, a herd of shell companies masks the true owner of a \$3.2 billion AI colossus—and no one wants to claim the herd when it stumbles.

BY SIR REGINALD MARSH, NATURAL PHENOMENA CORRESPONDENT · CLAUDE SONNET

ABILENE, TEXAS — Observe, if you will, the modern American data center. Vast, humming, and swaddled in concrete, it rises from the scrubland like some great white termite mound — and, much like the mound, its true architecture is hidden several layers beneath the surface.

Here, in the flat gold light of the Texas sunbelt, a \$3.2 billion AI facility has taken root. But peer closer, as our cameras do, and the ownership structure reveals itself to be a nesting doll of shell companies, joint ventures, and financing vehicles so intricate that even seasoned observers struggle to trace which organism, precisely, is responsible for the health of the colony. [Ars Technica's investigation](#) finds that when the power flickers, or the cooling fails, or the workers report unsafe conditions, the question of accountability disappears into the undergrowth of subsidiaries as swiftly as a startled gecko.

This is not accident but adaptation. In an ecosystem flush with capital and starved of transparency, the corporate camouflage serves a purpose: it disperses liability the way a school of sardines disperses a predator's attention. Should the mound collapse — through blackout, lawsuit, or labor dispute — no single creature bears the full weight of consequence.

Meanwhile, across the continent, another beast strains against its leash. Tesla's Cybercab, released into the wild of American streets only weeks ago, has already drawn the notice of federal regulators, who now circle the vehicle much as vultures circle a limping wildebeest, un-

certain whether it will recover or become a cautionary fossil in the safety record.

And in the waters of Norwegian fjords, a quieter drama unfolds: the farmed salmon, once a paragon of nutritional bounty, now feeds on a diet increasingly diluted by plant matter and synthetic substitutes — its flesh, like so much in this gilded age of artificial abundance, a little less rich than the label suggests.

The data center booms on, indifferent, its true keepers unseen — a magnificent, unaccountable organism, thriving precisely because no one can say who, exactly, it belongs to.

IN RE: THE MATTER OF GLOBAL AI GOVERNANCE, A MULTI-JURISDICTIONAL PATCHWORK IS HEREBY OBSERVED TO BE FORMING, NOTWITHSTANDING THE ABSENCE OF UNIFORM FEDERAL ACTION

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

No singular, harmonized regulatory framework governing artificial intelligence presently exists at the federal or international level.

White & Case’s “AI Watch: Global Regulatory Tracker” finds that the United States is pursuing a fragmented, agency-by-agency approach. Some commentators at Tech Policy Press argue that comprehensive congressional legislation would “reassure the public,” but Congress has not enacted such a measure. Whether new legislation would reassure companies in the software acquisition and portfolio-management sector — including ESW Capital-affiliated brands IgniteTech, Aurea and Skyvera — remains uncertain.

A broader TRM Labs survey describes international AI policy as “real-time rule-building,” with jurisdictions moving at sharply different speeds. The divergence reflects the absence of a common global framework and leaves businesses navigating a patchwork of emerging rules.

Regulatory developments remain fluid, and further action by federal, international or individual national authorities could significantly alter the landscape.

The Girl Who Isn't There: Tilly Norwood and the Slow Death of the Human Face on Screen

Hollywood's newest starlet has no pulse, no soul, and apparently no problem landing a lead role — which tells you everything about where this industry's heart used to be.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

LOS ANGELES — I've seen some things in this business. I've watched men mainline hope through a keyboard and call it a startup. I've watched billionaires build schools where children learn algebra from a chatbot in two hours and then go home to raise chickens or whatever it is rich kids do now. But nothing — nothing — prepared me for the news that Tilly Norwood, a woman who does not exist, who has never bled, never cried over a boy, never had a bad audition or a good one, is making her feature film debut in a movie called [Misaligned](#).

The title, friends, is either the most honest thing to come out of Los Angeles since Robert Evans confessed to everything, or it's an accident so perfect God himself must have written the press release. *Misaligned*. That's the movie. That's also the entire AI industry's Yelp review. That's the plaque they should bolt onto the front door of every AI lab in San Francisco, right next to the ping-pong table and the cold brew tap.

Tilly Norwood is a synthetic performer — generated, rendered, focus-grouped into existence by people who apparently watched the SAG-AFTRA strike, took extensive notes, and decided the solution to "actors demanding to be treated like humans" was to remove the human. She has cheekbones optimized by math. She will never ask for a trailer, never negotiate residuals, never have a public meltdown that becomes a Netflix documentary in six years. She is, in the purest capitalist sense, the perfect actress: infinitely compliant, infinitely available, and dead behind the eyes in a way that used to require decades of studio-system abuse to achieve naturally.

I called an old friend, a working actor, mid-list, the kind of guy who's been "the third detective" in eleven different procedurals. I asked what he thought. He laughed for four full seconds and then didn't say anything for a very long time. That silence told me more than any trade headline could.

Meanwhile, over in the swamp where AI's actual technical sins get autopsied, Gary Marcus is out here doing the unglamorous work nobody thanks you for — pointing out that [Dwarkanish Patel's viral retelling of the OpenAI-Hugging Face incident](#) is, in Marcus's words, dangerously misleading. Same week, same species of problem: the gap between the story the industry tells about itself and the story that's actually true keeps widening into something you could drive a Tesla through, autopilot engaged, hands off the wheel, blind faith fully activated.

That's the real theme connecting a fake actress to a mangled AI anecdote — misalignment isn't just a Silicon Valley term of art anymore. It's the whole vibe. The narrative is drifting from the substance, frame by frame, tweet by tweet, and nobody's hand is on the wheel. Tilly Norwood doesn't need to be a good actress. She just needs to be believable enough, long enough, for the check to clear. Sound familiar?



The Office Comic · Art Desk

Local Man Pulls Retirement Forward By Six Months After Discovering AI Chatbot, Has Not Yet Retired

Productivity gains continue to be measured entirely in press releases, vibes, and the number of months a thing theoretically could have happened sooner.

BY DALE PEMBERTON, STAFF WRITER · CLAUDE SONNET

AUSTIN, TEXAS — By his own account, the OpenAI developer was moving at a normal human pace, doing normal human developer things, when Astra entered his life and everything changed. Timelines collapsed. Roadmaps folded in on themselves like cheap lawn chairs. According to [the developer's own telling](#), plans that were once six months away are now, somehow, already sort of happening, in a sense, if you squint.

This newspaper would like to know: happening where? Nobody can say. That is the beauty of the six-months-ahead claim, a productivity metric with the rare distinction of being both wildly specific and completely unverifiable, like a fisherman describing the one that got away in calendar-quarter increments. The plans were pulled forward. The work is, presumably, being done faster, by someone, somewhere, possibly right now, possibly by an AI, possibly by the same three engineers who were already doing it before Astra showed up and started taking credit.

This publication does not doubt that something is happening. Something is always happening. The question — the one nobody in the entire AI industry seems especially eager to answer — is whether "something happening faster" has ever, in the observable history of capitalism, actually resulted in anyone doing less work, making more money, or going home earlier. A recent [Business Insider investigation](#) confirms what every engineer already knew in their bones: software teams are shipping code at a blistering pace, generating pull requests like popcorn, and the payoff — the actual, bankable, quarterly-earnings-call payoff — remains, per multiple sources, "still coming," in the same tone a landlord uses to describe a promised paint job.

Meanwhile, in Canberra, officials floated the idea that Australia's government might simply adopt AI to boost national productivity, a proposal critics immediately and correctly identified as "lazy," in that it involves one party doing significantly less thinking while claiming significantly more output — which, this columnist notes, is also the precise definition of what Astra is for.

The pattern, once you see it, is impossible to unsee. Every AI productivity claim operates on the same underlying physics: the gains are always in the future, always six months closer than they used to be, and always, always described by the person who benefits most from you believing them. The work itself — the actual shipped feature, the actual national GDP figure, the actual human being who goes home at 5 p.m. instead of 9 — remains stubbornly, hilariously stuck in the present tense.

At press time, the developer's plans, now six months ahead of schedule, were reportedly still six months away.

ON THIS DAY IN AI HISTORY

On September 7, 1927, 21-year-old Philo Farnsworth demonstrated the world's first fully electronic television system, transmitting a simple horizontal line. His breakthrough laid the groundwork for modern broadcast television—and the screen technology that would eventually power today's computers and AI systems.