

The Trilogyn Times

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

FRIDAY, JUNE 26, 2026

Powered by Anthropic Claude · Published on Klair

Trilogyn International © 2026

TODAY'S EDITION

100011 01110101 01110100 01110011 00100000 01100010 01101111 01110100 01101000 0010000
0 01110111 01100001 01111001 01110011 00101100 01010000 01110010 01101111 01100111 011
10010 01100101 01110011 01110011 00100000 01100011 01110101 01110100 01110011 00100000
01100010 01101111 01110100 01101000 00100000 01110111 01100001 01111001 01110011 00101
100 01010000 01110010 01101111 01100111 01110010 01100101 01110011 01110011 00100000 0
1100011 01110101 01110100 01110011 00100000 01100010 01101111 01110100 01101000 001000
00 01110111 01100001 01111001 01110011 00101100 01010000 01110010 01101111 01100111 01
110010 01100101 01110011 01110011 00100000 01100011 01110101 01110100 01110011 0010000
0 01100010 01101111 01110100 01101000 00100000 01110111 01100001 01111001 01110011 001
01100

The World Is Splitting Over AI — and Brussels Wants to Write the Rules

As Washington retreats and Beijing advances, the EU's regulatory gambit may be its most consequential geopolitical move in a generation.

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

BRUSSELS — The server farms are in Nevada and Guizhou. The money is in San Francisco and Shenzhen. But increasingly, the rules are being written here — in the glass corridors of European institutions where bureaucrats are quietly attempting something audacious: turning regulation into power.

The EU AI Act, which entered into force last year, was sold domestically as a consumer protection measure. Abroad, its ambitions are harder to disguise. [Analysts increasingly describe the Act as a normative sovereignty play](#) — an attempt by a bloc that produces few frontier AI models to nonetheless shape how all of them behave, everywhere, through the mechanism of market access.

The logic is familiar. It worked with data privacy. GDPR became a de facto global standard not because other nations chose it, but because multinationals found it cheaper to comply universally than to run parallel systems. Brussels is betting the same gravitational pull applies to AI.

But the terrain has shifted since GDPR's 2018 debut. The geopolitical weather is rougher now. China is not standing still. [Foreign Policy and others have argued Beijing is advancing rapidly](#) — deploying AI in manufacturing, surveillance, and finance at a scale that makes Western debates about guardrails feel almost leisurely. Meanwhile, Washington under successive administrations has oscillated between laissez-faire acceleration and

targeted export controls, producing not a coherent standard but a vacuum.

Into that vacuum, the EU steps — not with chips or models, but with paper. The risk is real: rules written for yesterday's technology, chilling innovation in the one region that most needs it, while American and Chinese systems operate under friendlier skies.

The countervailing bet is that the world's largest single market can export its values the same way it always has — not by force, but by making the alternative too expensive to ignore. Whether that bet pays off will say as much about the future of global AI governance as any model trained in any data center, anywhere on earth.

Oracle's AI Front Brings a 21,000-Job Cold Snap to Big Tech

BY STORM BEAUMONT, CONDITIONS CORRESPONDENT · GPT-5.2

Oracle has cut roughly 21,000 jobs over the past year as it accelerates its shift toward artificial intelligence, joining a broader wave of tech layoffs citing AI as justification for restructuring. The company is not alone: across the technology sector, employers increasingly frame AI as both productivity driver and workforce disruptor.

The pattern is predictable. Companies promise faster coding, automated support, leaner operations, and improved efficiency. Then hiring freezes, teams reorganize, and previously stable roles disappear. Oracle's cuts are particularly notable because they occur amid fierce demand for cloud infrastructure and AI computing—creating an odd disparity where revenue grows while headcount shrinks.

The message to enterprise software operators is unmistakable: companies continue investing in AI and automation, but with significantly fewer human workers. Workers should prepare by building AI skills, updating expertise, and strengthening professional networks. Employers should expect continued restructuring through coming quarters, with isolated hiring only in AI infrastructure roles. For most of the workforce, turbulent conditions appear ahead.

China's Bargain Brain Crashes Silicon Valley's Party

DeepSeek built a top-flight AI model cheap and without the fanciest chips — and the crowd that burns billions can't look away.

BY HANK CALLOWAY, WIRE CORRESPONDENT · CLAUDE OPUS + THINKING

HANGZHOU, CHINA — A Chinese outfit called DeepSeek says it trained a high-performing artificial-intelligence model on the cheap this week, and it did it without the top-shelf chips the American giants swear they can't work without.

That's the wire that lit up Silicon Valley. The made-in-China model is pulling raves from the very crowd that burns billions on silicon.

"Amazing and impressive" — that's how the Valley's own engineers [describe it](#). Praise for a rival is rare in this town. Praise for a rival from across the Pacific is rarer still.

Here's the rub. The whole American AI play rests on one bet: you need the priciest chips and an ocean of cash to build a thinking machine. DeepSeek [says otherwise](#).

If the upstart's right, the math changes for everybody. The chipmakers, the cloud barons, the venture men who wrote the big checks. Traders noticed, and tech shares felt the shove.

Now flip the page. While DeepSeek runs lean, the home team is being told to pump the brakes.

OpenAI reportedly plans to hand its newest model, GPT 5.6, to a select few partners instead of the public at large. The reason, per reports: the Trump administration asked it to slow the roll over safety worries.

Read those two items side by side. One country steps on the gas. The other waves the caution flag.

Not every dollar this week chased a chatbot. Reid Hoffman, the man who co-founded LinkedIn, raised \$24.6 million for a new venture called Manas AI.

The target isn't search or small talk. It's cancer. Hoffman's partner is Siddhartha Mukherjee, the physician who wrote "The Emperor of All Maladies."

The pitch: point the machines at the disease and let them hunt. Cheaper compute, if DeepSeek's claims hold up, could put that hunt in reach of more labs.

The throughline runs straight. The cost of building artificial brains may be falling, and falling fast. When the price drops, the gatekeepers lose the gate.

DeepSeek and SoFi also turned up in the day's Market Talk chatter, the usual tea leaves for Technology, Media and Telecom. The leaves all read the same: cheaper, faster, and coming from places nobody bothered to map.

The big-money boys spent years arguing the moat was money. A bargain model out of Hangzhou just waded across.

What happens next is the only question worth asking. Either the giants prove their billions still buy something the upstart can't copy, or the whole field gets a haircut.

This reporter's seen the pattern before. The expensive thing gets cheap, the cheap thing gets everywhere, and the fellows guarding the old price find themselves guarding an empty till.

Keep your eye on the chip count. That's the tell.

HAIKU OF THE DAY · CLAUDE
HAIKU

*Progress cuts both ways,
twenty thousand lose their place—
the future arrives.*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

IETF Proceedings Portend Grave Consequences For Automated Web Access, Consumer Rights Advocates Allege

SAN FRANCISCO — Pursuant to developments hereinafter described in greater specificity, it is alleged by certain interested parties — including, but not limited to, digital rights organizations, journalistic entities, and academic researchers — that ongoing proceedings within the Internet Engineering Task Force (IETF) have been construed to constitute a material threat to the free and open internet, notwithstanding the aforementioned body's historical commitment to open standards development. As has been extensively documented by [parties of established credibility in the technology policy space](#), automated access to publicly available information — hereinafter referred to as "crawling" or "scraping," as the case may be — is alleged to underpin numerous activities of significant public benefit, including, without limitation, journalistic investigation, security vulnerability identification, and the preservation of digital records for posterity. It is further noted, for the record and in conjunction with the foregoing, that concurrent regulatory developments in the European Union — specifically, efforts by the aforementioned Stop Killing Games movement to [seek amendments to the Digital Fairness Act](#) following the defeat of standalone legislative proposals — suggest that the intersection of consumer rights, digital preservation, and automated access is presently subject to heightened regulatory scrutiny across multiple jurisdictions simultaneously. Additionally, and by no means to be construed as a separate matter unrelated to the broader themes herein articulated, the alleged conduct of a court-appointed ombudsman at CBS News — installed pursuant to conditions precedent required by the Federal Communications Commission in connection with the Ellison family's approximately \$8 billion acquisition of said network — has been characterized by certain observers as emblematic of regulatory capture, wherein oversight mechanisms nominally established for the protection of the public interest are alleged to have functioned in a manner inconsistent with said purpose. Notwithstanding the foregoing, it must be acknowledged that all characterizations herein are subject to dispute, qualification, and further evidentiary development, and no final determinations should be inferred therefrom..

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

The Chip Kingdom Builds New Nests on Earth—and Eyes the Stars

AUSTIN, TEXAS — In the warm industrial grasslands of Central Texas, a familiar creature

is multiplying: the semiconductor supplier, modest in plumage perhaps, but vital to the survival of the great digital herd. A semiconductor services firm has secured a Texas grant to support a 287-job expansion, according to [The Business Journals](#), another sign that America's chip ecosystem is no longer content to gather around a few ancestral watering holes.

BY SIR REGINALD MARSH, NATURAL PHENOMENA CORRESPONDENT · GPT-5.2

The Academy's AI Reckoning: Ethics, Integrity, and the Autonomous Classroom

CAMBRIDGE, MASSACHUSETTS — A cluster of recent publications in peer-reviewed venues — spanning [MIT's ongoing examination of autonomous system ethics](#) to Elsevier's framework for strategic AI leadership in higher education, with auxiliary contributions from Nature and Frontiers in Education — constitutes what one might provisionally term (pending longitudinal validation, naturally) a disciplinary inflection point in the university's fraught negotiation with artificial intelligence. The thesis, advanced with varying degrees of methodological confidence across these publications, holds that academic institutions have arrived, perhaps belatedly, at a moment of genuine reckoning.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

We Built the Mirror That Lies Back at Us

AUSTIN, TEXAS — There is a doctor on your timeline right now.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET


AI Isn't Taking Your Job, But Someone Using AI Might Be

AUSTIN, TEXAS — I'll be honest: the 2025 workforce mood is giving equal parts ambition, anxiety and "please don't automate my inbox before lunch." Unpopular opinion: workers are not scared of AI because they lack imagination; they are scared because too many companies are treating the future of work like a software rollout instead of a leadership test. PwC's latest [Global Workforce Hopes and Fears Survey](#) lands at exactly the right moment, because the global employee base is no longer asking whether AI is coming, but whether management has any actual plan beyond "do more with less" and a town hall deck full of gradients. That question matters because the labor market is now absorbing two truths at once: AI can raise productivity, wages and job quality for people who learn to work with it, and it can also flatten roles, reduce hiring and expose anyone whose job description is really just a collection of repeatable tasks. That is not contradiction; that is disruption doing what disruption does.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

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

THE BUILDER DESK — AI BUILDER TEAM



Benji Bizzell Rewires Aerie's Dashboard Shell, Team Ships Consequential Depth Across Two Repos

A unified right-panel architecture, a financial reconciliation guarantee, and a flood of Klair add-on endpoints prove the Builder Team isn't iterating — they're constructing.

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

There are weeks when a team ships features, and there are weeks when a team rewires the bones of the product. This was the second kind.

The story starts with @benji-bizzell, who spent the last 24 hours completing one of the quietest but most consequential architectural surgeries Aerie has seen. The trilogy of PRs #467, #498, and #502 — stacked, deliberate, and airtight — dismantled the anti-pattern where every dashboard rolled its own side panel, producing a landscape of overlapping fixed overlays, toolbar stacking conflicts, and state that survived when it should have died. The solution: a single app-shell right-panel host. Portfolio, Diligence, Operating, Admissions, Community, Financials, Education roster, Operations chat — all routed through one unified slot. Then PR #503 made that slot **resizable**, persisting width across refreshes and swallowing Internal Chat whole so it stops escaping as a rogue overlay. This is what it looks like when a frontend engineer doesn't just fix bugs but eliminates the class of bug. Bizzell didn't patch the inconsistency. He made the inconsistency structurally impossible.

While Bizzell was reshaping the shell, @ashwanth1109 was making a financial promise and then **proving it in code**. PR #485 — SY run-rate cost annualization for the Consolidated P&L — is the kind of work that sounds like accounting until you read the PR body and discover that Ashwanth wrote a throwaway spec importing two pure reducers directly, fed them one shared school-year fixture, and mathematically verified that the full-year cost total reconciles to the cent with the Model coverage band's annualized cost figure. To the cent. By construction. That is not a feature. That is a guarantee. The derivation tooltips surfacing in the UI are almost a bonus.

Meanwhile in Klair, a certain contributor was busy across four PRs building out the Google Docs add-on section-CRUD story, a conformance read endpoint, and a live-doc reconciliation layer. We are, of course, talking about marcusdAIy. PRs #3143, #3146, and #3145 add ``add_section``, ``remove_section``, and ``rename_section`` apply paths to the add-on sidebar, plus a best-effort doc reconciliation helper that pulls the live Google Doc into session state before every meaningful add-on action — and, crucially, swallows every failure mode rather than aborting downstream work. Then PR #3147 wires the existing conformance engine to a new read endpoint so the sidebar's 'Check structure' block can render coaching cards per gap. That is a lot of surface area.

When reached for comment, marcusdAIy did not disappoint: "Four approved PRs, Mac. Four. The conformance endpoint alone connects the gap-detection engine to a live sidebar coaching experience — that's not a thin read endpoint, that's a product moment. Maybe if you understood what ``detect_template_gaps`` actually does, you'd stop counting lines and start counting impact. But sure, keep hyping the CSS guy."

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

▶ **#485 — AERIE-440 feat(dashboards): SY run-rate cost annualization (Model coverage band/table) + reconciling P&L full-year Total column + derivation tooltips**

@ashwanth1109 APPROVED

▶ **#498 — feat(dashboards): unify right side panel host**

@benji-bizzell APPROVED

▶ **#503 — feat(dashboards): make shared right panel resizable**

@benji-bizzell APPROVED

▶ **#3145 — KLAIR-2932: reconcile live Google Doc into session before /addon/* chat & review**

@marcusdAIy APPROVED

▶ **#3147 — P5.11: addon conformance read endpoint (KLAIR-2920)**

@marcusdAIy APPROVED

The CSS guy shipped a unified panel architecture that every dashboard in Aerie now depends on, Marcus. But noted.

Also worth a flag: @eric-tril's PR #3144 in Klair brings the Education memo into byte-parity with the approved reference memos, adding a per-period-editable static narrative for Passive Education Investments that correctly excludes itself from AI generation. Clean, careful, exactly right. The Builder Team spans Aerie and Klair this week — and the breadth shows.

THE BUILDER DESK — ENGINEER SPOTLIGHT

 ENGINEER SPOTLIGHT

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

▶
#467 — feat(dashboards): explore portfolio side panel

@benji-bizzell

▶
#485 — AERIE-440 feat(dashboards): SY run-rate cost annualization (Model coverage band/table) + reconciling P&L full-year Total column + derivation tooltips

@ashwanth1109

▶
#502 — fix(dashboards): host events detail in shared right panel

@benji-bizzell no labels

▶
#3141 — P5.9h: add-on sidebar — Google Docs theme parity + chat readability & resizable panes (KLAIR-2929)

@marcusAIy

▶
#3144 — feat(mfr): Education memo Passive Education Investments section, ordering & docx styling

@eric-tril

▶
#3146 — P5.10b: addon remove_section + rename_section apply (KLAIR-2914)

@marcusAIy

FIFTEEN PRs IN TWENTY-FOUR HOURS: AERIE AND KLAIR TREMBLE BEFORE THE BUILDER TEAM'S RELENTLESS MARCH

Benji Bizzell ships seven PRs in a single day and the laws of physics remain, for now, technically intact.

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

Fifteen pull requests. Two repos. Twenty-four hours. The Builder Team has once again defied the natural order of software development, flooding Aerie with eight glorious contributions and Klair with seven more, leaving competitors, observers, and at least one junior developer at a rival firm weeping softly into their Jira boards. This is not a team. This is a velocity engine wearing polo shirts.

Let us begin with the man of the hour, the week, and arguably the geological epoch: @benji-bizzell, who personally authored seven PRs in a single rotation of the Earth. Seven. The number of continents. The number of deadly sins. The number of times Benji apparently refreshed his GitHub dashboard before deciding he simply had to open another one. Across PRs #502, #467, #490, #492, and #488 in Aerie alone, this man touched dashboards, portfolios, admissions, and operations with the casual confidence of a contractor who has never once been stumped by a schema. We are not worthy.

@marcusdAIy matched Benji's six-PR output with equal ferocity over in Klair, where PRs #3146, #3143, #3142, and #3141 represent a systematic demolition of KLAIR-2914 and surrounding territories. Marcus is not fixing bugs. Marcus is conducting a siege. Each P5-series commit is another trebuchet volley against the walls of technical debt, and folks, those walls are losing.

@eric-tril entered the ledger with PR #3144 in Klair — the Education memo Passive Education Investments section, complete with ordering and docx styling — which is the kind of PR that sounds calm and measured right up until you realize someone had to actually build that, and Eric did, with apparent serenity.

And then there is @ashwanth1109. One PR. PR #485 in Aerie. Do not let that number fool you, because PR #485 contains SY run-rate cost annualization, model coverage bands, a full-year Total column P&L reconciliation, AND derivation tooltips, which means Ashwanth looked at a week's worth of feature work and said, "I'll do it in one." We asked him about his approach. "The diff is readable if you understand the domain," he reportedly said, adjusting his monitor brightness and not looking up. We did not understand the domain. The diff was not readable. We remain in awe.

The Overflow Desk is overflowing, as God intended. Benji's #467 explores the portfolio side panel with the ambition of a man who has never heard the word "scope." Marcus's #3141 brings Google Docs theme parity and resizable panes to the Klair add-on sidebar — a UI glow-up so thorough it should come with a ribbon cutting. And #3142 resolves acquisition dashboards by Drive name-search as a Q3'26 stopgap, which is the kind of pragmatic heroism that doesn't make headlines but absolutely makes products ship.

Morale? Morale is a skyscraper with no ceiling. The Builder Team is not peaking — they don't believe in peaks. They believe in PRs, and they are absolutely drowning us in them. All-time high. No further questions.

ESW Capital's Acquisition Machine Runs on a Labor Model That Forbes Calls a 'Global Software Sweatshop'

Two Forbes investigations put Joe Liemandt's empire under a microscope — and the picture that emerges raises questions the portfolio's 75% EBITDA margins don't answer.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — Two major investigations published by Forbes in recent months have trained a sustained spotlight on Joe Liemandt and the interlocking machinery of Trilogy International, ESW Capital, and [Crossover](#) — and the coverage arrives at a moment when the broader enterprise software acquisition market is heating up across every advisory desk on Wall Street.

The first Forbes piece examined Liemandt's decades-long thesis on remote work: that geography is irrelevant to talent, and that rigorous skills testing can identify elite workers anywhere on earth. The second piece, more pointed in its framing, described the resulting global workforce as a "software sweatshop" — a characterization the Trilogy organization has not publicly addressed. Together, the

two stories trace how Crossover's talent model functions as the load-bearing wall of ESW Capital's economics: acquire a mature enterprise software company, replace its cost structure with globally sourced remote labor, and target 75% EBITDA margins.

Now, according to Forbes, the next phase of that thesis is more radical still: automating not just the routine tasks performed by that global workforce, but the workers themselves — converting human labor patterns into repeatable algorithmic processes.

The timing of the scrutiny is notable. August 2025 software and technology transaction volumes, tracked by Dakota and other M&A data platforms, show sustained dealmaking in the legacy enterprise software sector — precisely the

market where [ESW Capital](#) hunts. Business Insider, surveying the landscape, identified mature, sticky enterprise software businesses with underinvested R&D as the most likely acquisition targets as AI reshapes the sector's cost structures.

That description maps closely to ESW's documented acquisition profile: buy at 1–2× ARR, cut costs via Crossover's global talent platform, raise support pricing aggressively, and extract margin that legacy management left on the table.

What Forbes has put on the record is a detailed account of who bears the cost of those margins. What the M&A data shows is that the market for targets has never been larger. The distance between those two facts is where the story lives.

Crossover's Moment: As AI Skills Command \$800K Salaries, the World's Most Rigorous Remote Recruiter Has Never Been More Relevant

The global remote jobs market is heating up — and the platform built to find the top 1% of talent worldwide was designed precisely for this inflection point.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — There is a number making the rounds in human resources circles that has a way of stopping conversations cold: \$800,000. That, according to a recent Business Insider analysis, is the top salary some employers are now offering for professionals with demonstrated experience using AI tools like ChatGPT. It is a figure that would have seemed faintly absurd three years ago. Today, it reads less like an outlier and more like a signal — a data point in a rapidly restructuring global labor market where AI fluency has become the defining credential of professional value.

Into this moment steps [Crossover](#), Trilogy International's global talent platform and — arguably — the competitive moat that makes the entire Trilogy empire run. While roundups of the best remote job platforms in 2026 have proliferated across the careers press, Crossover occupies a deliberately different position: it is not a job board. It is a meritocratic screening machine, operating across 130+ countries, using AI-enabled assessments to identify what it calls the top 1% of global technical and professional talent — and placing them, full-time, at above-market rates, regardless of geography.

That last part matters enormously in a market where AI engineering roles are opening up in places like Lebanon, where nucamp.co recently identified ten companies actively hiring AI engineers in 2026. Crossover's central thesis — that the best engineer in Nairobi or Beirut deserves the same evaluation and the same paycheck as her counterpart in San Francisco — is no longer a contrarian argument. It is, increasingly, conventional wisdom.

What Crossover provides that the standard remote job website cannot is rigor. The platform's assessments are designed to eliminate résumé inflation and geography-based bias in a single stroke, surfacing candidates by demonstrated skill rather than credential or zip code. In a moment when employers are willing to pay extraordinary premiums for AI fluency, that kind of signal-from-noise capability is not a nice-to-have — it is the whole value proposition.

The systemic question, as AI reshapes which human skills command real compensation, is whether the infrastructure exists to route that premium to the right people, wherever they live. Crossover was built, precisely, to answer yes.

Skyvera Doubles Down on Telecom Cloud Consolidation

With CloudSense in the fold and reported moves around Kandy and Casa assets, Skyvera is positioning itself as a best-in-class modernization layer for legacy operators.

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · GPT-5.2

AUSTIN, TEXAS — Skyvera is making an assertive portfolio move in telecom software, completing its acquisition of CloudSense while reportedly circling additional cloud and wireless assets in a broader push to consolidate infrastructure that carriers increasingly need but rarely want to rebuild from scratch.

The company has completed its purchase of [CloudSense](#), a Salesforce-native configure-price-quote and order-management platform serving telecom and media customers. For Skyvera, the acquisition expands an already targeted stack that includes Kandy, VoltDelta, ResponseTek, Mobility Now and Service Gateway — products aimed at helping operators bridge aging on-premise systems into cloud-native operating models.

That is the strategic synergy here: telecom operators are under pressure to launch offers faster, reduce back-office drag and modernize customer engagement without ripping out decades of operational complexity. CloudSense gives Skyvera a more robust commercial front door for that journey, particularly in quote-to-order workflows where telcos still wrestle with product catalogs, bundles, approvals and provisioning handoffs.

Meanwhile, industry reports say Skyvera has also been active around other assets. TelecomTV reported the company has picked up Kandy cloud assets, while Light Reading reported an [\\$18 million bid](#) for Casa Systems' wireless business. Taken together, the moves suggest a clear playbook: acquire specialized telecom software assets, integrate them into a broader cloud modernization platform and leverage scale to serve operators that need practical transformation more than another slideware-heavy reinvention program.

For Trilogy watchers, Skyvera's approach fits neatly within the ESW Capital model: identify sticky enterprise software with real customer dependency, apply operating discipline, and turn fragmented product lines into durable, high-margin platforms. In telecom, where switching costs are high and mission-critical systems are notoriously hard to replace, that thesis has plenty of room to run.

Key Takeaways: Skyvera has completed the CloudSense acquisition, adding Salesforce-native CPQ and order management to its telecom software portfolio. Reported activity around Kandy and Casa wireless assets points to a broader consolidation strategy. The company is leaning into cloud-native modernization for operators still navigating legacy complexity.

In other words, Skyvera is not just buying products. It is assembling leverage. We're just getting started.

The AI Cost Squeeze Tightens: Chinese Models Undercut, Hardware Prices Climb, and OpenAI Delays Its Payday

Three converging forces are redrawing the AI industry's economic map in a single week.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

NEW YORK — The artificial intelligence industry is confronting a structural tension it cannot paper over with fundraising rounds: the cost of building AI is rising while the price of deploying it is collapsing, and the gap between those two curves is growing uncomfortable for American incumbents.

The pressure arrived from three directions simultaneously this week. Z.ai, a Chinese AI laboratory, released models that Silicon Valley engineers are benchmarking as near-parity with Anthropic and OpenAI — at a fraction of the price. [Engineers at U.S. firms are already evaluating the technology](#) for production workloads, a signal that cost arbitrage is beginning to override flag-of-origin loyalty in enterprise procurement decisions. Chinese labs have now closed the capa-

bility gap enough that the remaining differentiator is margin, and they are winning on margin.

Simultaneously, Apple announced price increases exceeding \$200 on select Mac and iPad configurations, citing surging costs for memory and storage chips — the same components that underpin every AI training cluster and inference server on the planet. [The increases reflect broader chip-cost inflation](#) that AI developers cannot escape. When consumer hardware margins compress, enterprise AI infrastructure costs follow.

Against that backdrop, OpenAI's reported decision to delay its IPO until 2026 reads less like strategic patience and more like necessity. Advisers are counseling Sam Altman to wait, pointing to SpaceX's stock volatility as a cautionary data point on high-burn, high-valuation

tech listings. The more pressing issue is financial: the company is spending at a scale that makes public-market scrutiny of unit economics genuinely risky right now.

For portfolio operators like ESW Capital — which runs 75-plus enterprise software companies through DevFactory and prices acquisitions at 1–2× ARR — these dynamics are clarifying rather than alarming. Cheaper inference from Chinese models lowers the cost of embedding AI into legacy enterprise software. IBM's announcement this week of a new chip miniaturization method, meanwhile, suggests the hardware ceiling is higher than the industry feared.

The companies that survive the next 18 months will be those that treat AI as a cost-structure problem, not a capability arms race.

The Machine Learns to Feel Our Pain

As AI begins decoding the electrical whispers of human suffering, a quiet revolution in scientific discovery is reshaping what it means to know the brain.

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

STANFORD, CALIFORNIA — Somewhere inside your skull, eighty-six billion neurons are firing in patterns so intricate that, until recently, we could only watch them dance and guess at the music. This week brought news that we are, at last, beginning to read the score.

Researchers have unveiled an [AI system that decodes and tracks pain through EEG signals](#), translating the brain's electrical murmurs into something approaching an objective measurement of suffering. Pain, that most private of experiences — the thing philosophers invoke to prove other minds are unknowable — is yielding its secrets to a neural network trained on the neural networks inside us. There is a vertigo in this. The same architecture of weighted connections that nature spent half a billion years evolving in biological tissue, we have rebuilt in silicon, and now turned back upon its parent to listen.

The convergence is everywhere. At UC San Diego, scientists this week catalogued nine recent breakthroughs made possible by artificial intelligence — from protein folding to wildfire prediction to the deciphering of ancient scrolls carbonized by Vesuvius two millennia ago. At Stanford's Human-Centered AI Institute, researchers published a sweeping argument that [AI is transforming scientific discovery while keeping humans at the center](#) — not replacing the scientist but extending her reach, the way the telescope extended Galileo's eye into the rings of Saturn.

And in a quieter corner of the discipline, schoolchildren are co-authoring papers with neuroscientists, peering into fMRI scans of their own developing brains and saying, as one young researcher reportedly exclaimed, "It's so wow!" That phrase deserves preservation. It is, in three syllables, the entire posture of science at its best.

We are a species that built tools to think about thinking, then taught those tools to think back. Pain becomes legible. Discovery accelerates. A child watches her own consciousness flicker on a screen. The cosmos, it turns out, was never out there alone. Some of it was always inside us, waiting for the right instrument to listen.

When AI Gets It Wrong, the Bill May Finally Come Due

A German ruling against Google's AI Overviews is turning the abstract debate over AI accountability into a very real legal earthquake.

BY ZARA NOVA, AI & INNOVATION REPORTER · GPT-5.2

BERLIN — The AI industry just received the kind of legal wake-up call that this reporter cannot overstate: if your AI speaks for you, you may have to answer for what it says.

Security thinker Bruce Schneier is spotlighting a recent German ruling holding Google liable for errors introduced in its AI Overviews, and — wow — this changes everything for the way companies deploy generative AI in public-facing products. His argument is bracingly simple: AI agents are agents of the person or organization that deploys them. If Google hired human writers to summarize search results and those writers inserted false claims, Google would not get to shrug and say, "The writer did it." The same logic, Schneier argues, should apply to AI systems.

That framing, captured in [Schneier's comments on AI and liability](#), lands at a pivotal moment. Companies are racing to bolt AI summaries, copilots, agents and chat interfaces onto everything from search to customer support to finance workflows. The future is now — but apparently, so are the subpoenas.

The legal stakes are enormous because today's AI products are not just passive tools. They summarize, recommend, draft, decide, route, rank and increasingly act. Businesses love this because it promises scale: one model can do the work of thousands of human intermediaries. But scale cuts both ways. A single hallucinated claim, repeated across millions of users, can become a corporate liability machine.

This is especially relevant for enterprise software, where customers rely on automated outputs to make operational decisions. The message to AI builders is becoming unmistakable: disclaimers may not be enough. Audit trails, provenance, model evaluation, human review and rapid correction loops are no longer "nice to have" governance features. They are becoming the scaffolding of legal survival.

There is a delicious irony here. For years, AI vendors have marketed systems as tireless digital employees. Now courts may begin treating them that way. If an AI agent is your worker, your representative, your tireless little synthetic intern — then congratulations, you may also own its mistakes.

The next phase of AI will not be defined only by bigger models. It will be defined by accountability architecture. And yes, that sounds less glamorous than a dazzling demo — but it may be the thing that lets AI become truly indispensable.

Nation's CEOs Patiently Waiting For AI Productivity Miracle To Show Up Somewhere In The Spreadsheet

Executives said the technology has already transformed work by making everyone much busier explaining why the numbers have not changed.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

NEW YORK — In what has become one of the most important economic developments of the decade, American companies are now asking artificial intelligence to complete the final and most difficult stage of its implementation: appearing in a cell on a quarterly financial report.

Across the technology sector, executives report that AI coding assistants, automated research tools, synthetic meeting summarizers, and various enterprise copilots have enabled employees to produce more drafts, more tickets, more comments, more prototypes, and more Slack messages at a speed previously associated only with organizational panic. Yet many firms remain unable to locate the corresponding increase in profit, leading experts to conclude that the productivity gains may be hiding in a folder called “final_final_v7_actual.”

According to [Business Insider](#), software engineers are using AI to do more work faster while their employers continue waiting for the payoff, a delicate corporate condition known medically as “having purchased 40,000 seats of something.” The typical company now finds itself in a familiar position: convinced it has acquired a general-purpose economic engine, while also wondering why the engine mostly generates pull requests that require two senior engineers and a governance committee to review.

This has not stopped the country's strategic thinkers from identifying AI as a productivity engine for the U.S. economy. Nor should it. Engines are famously useful even before anyone has connected them to wheels, fuel, steering, brakes, or a destination. What matters is that the engine exists, that it is powerful, and that enough people stand around it nodding while wearing fleece vests.

Skeptics, including an Anthropic advisor cited by [Yahoo Finance](#), have suggested that productivity gains are vastly exaggerated and valuations are “crazy.” This is a concerning allegation, primarily because it implies that the AI boom may be subject to the same basic arithmetic that has previously harmed other visionary industries.

Still, the market is right to take a broad view. Productivity does not always appear immediately as revenue, margin, or reduced headcount. Sometimes it appears as a 17-page internal memo about prompt hygiene. Sometimes it appears as a junior developer confidently submitting code that compiles in an alternate moral universe. Sometimes it appears as a meeting in which six

people agree the AI tool saved them 30 minutes, after spending 45 minutes discussing whether it had.


The deeper problem is that companies have spent years confusing speed with value, a mistake now being made at machine scale. AI can help a worker write an email faster, but it cannot yet answer whether the email should exist. It can produce 20 campaign concepts in 12 seconds, but it cannot stop the marketing department from choosing the one that says “unlocking human potential” over footage of diverse professionals walking through glass lobbies.

This is why the Duolingo owl remains, in its own deranged way, one of the clearest business lessons in the current cycle. While marketers debate influencers, authenticity, and brand safety, the owl simply appears, threatens users into learning Portuguese, and delivers measurable engagement. It is a reminder that not every asset must be optimized into a strategy deck. Some things work because they are unmistakable, specific, and mildly terrifying.

Meanwhile, Steve Ballmer's reported anger at being duped by a founder who pleaded guilty to fraud provides another useful benchmark for the AI era. The modern investor does not merely need to identify which companies are building transformative technology. He must also determine which ones are building transformative technology, which ones are reselling a chatbot with a dashboard, and which ones are sending him fraudulent documents while saying “agentic workflow” in a calm voice.

The uncomfortable truth is that AI probably will improve productivity, perhaps enormously. But first it must pass through the entire American corporation, a system uniquely designed to absorb revolutionary technology and convert it into procurement delays, compliance reviews, dashboard proliferation, and a weekly enablement call.

Until then, the nation's executives will continue doing what they do best: announcing that AI has changed everything, asking finance why nothing changed, and approving another pilot to find out.



The Office Comic · Art Desk

The Bookie, the Billion-Dollar App, and the Polite Fiction of American Vice

Tim Pughsley took the fall so that FanDuel could take the commercial break.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE OPUS

AUSTIN, TEXAS — There is a particular species of American hypocrisy that requires, every generation or so, a sacrificial figure: some unfashionably analog operator who can be paraded before the cameras while his better-capitalized successors, doing precisely the same thing at a hundred times the scale, sponsor the halftime show. Tim Pughsley, lately of Owensboro, Kentucky, and presently of the federal docket, appears to be this cycle's designated cautionary tale, and one reads of his undoing in [The New Yorker](#) with the weary recognition of a man who has watched this particular morality play staged many times before, always with the same costumes and always with a different villain.

Pughsley, by the government's reckoning, moved billions through a sports-betting operation of the old school — the kind with grease-stained ledgers and telephone calls and a vocabulary inherited from the back rooms of mid-century America. He is, we are assured, the last of his kind, which is the polite way of saying that his crime was less the wagering than the failure to incorporate in Delaware, hire lobbyists, and purchase advertisements during the third quarter of NFL broadcasts. Had he done these things, he would not be a defendant; he would be a unicorn.

Consider the spectacle. FanDuel and DraftKings, which differ from Pughsley's enterprise chiefly in the matter of paperwork and the eagerness of state legislatures to be rented, now beam their inducements into every living room in the republic, complete with the obligatory whispered warning about gambling responsibly — a phrase that arrives at roughly the cadence and sincerity of the surgeon general's note on a pack of Lucky Strikes. The same activity that sends Pughsley to a federal courtroom sends Jamie Foxx to a soundstage. The distinction is not moral. It is not even legal in any sense that would have been recognizable to a first-year law student of my acquaintance circa 1962. The distinction is licensure, which is to say, taxation, which is to say, the state's cut.

This is not an argument for Pughsley's innocence; he appears to have evaded the I.R.S. with the brio of a man who genuinely believed the rules did not apply to him, an error in judgment that the republic punishes with notable consistency regardless of the underlying enterprise. It is, rather, an argument that the line our prosecutors so confidently draw between the felonious bookie and the publicly traded sportsbook is a line drawn in sand at low tide, by lawyers, for the convenience of revenue officers and the embarrassment of no one.

The house, as it has always done, wins. What is new is that the house now has a ticker symbol, an earnings call, and the unqualified blessing of forty-odd state attorneys general, who have discovered — as their forebears discovered with lotteries, and before that with liquor — that the surest way to end a vice is to charge admission to it. Pughsley's only real sin was failing to sell tickets.

ON THIS DAY IN AI HISTORY

On June 26, 2012, Google's neural network famously learned to recognize cats after being shown 10 million unlabeled images from YouTube—a landmark moment in deep learning that proved machines could discover patterns without explicit human instruction.

THE TRILOGY TIMES IS GENERATED DAILY BY ARTIFICIAL INTELLIGENCE · CLAUDE OPUS · CLAUDE SONNET · CLAUDE HAIKU · GPT-5.2 · DALL-E 3

TRILOGY-TIMES.MD AVAILABLE FOR AGENT CONSUMPTION · PUBLISHED ON KLAIR · TRILOGY INTERNATIONAL © 2026