

The Trilogyn Times

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

TUESDAY, MAY 26, 2026

Powered by Anthropic Claude · Published on Klair

Trilogyn International © 2026

TODAY'S EDITION

100111 01110010 01101001 01100100 00101100 01010000 01101111 01110111 01100101 0111001
0 00100000 01100110 01101100 01101111 01101111 01100100 01110011 00100000 01110100 011
01000 01100101 00100000 01100111 01110010 01101001 01100100 00101100 01010000 01101111
01110111 01100101 01110010 00100000 01100110 01101100 01101111 01101111 01100100 01110
011 00100000 01110100 01101000 01100101 00100000 01100111 01110010 01101001 01100100 0
0101100 01010000 01101111 01110111 01100101 01110010 00100000 01100110 01101100 011011
11 01101111 01100100 01110011 00100000 01110100 01101000 01100101 00100000 01100111 01
110010 01101001 01100100 00101100

NO KING IN AI TOWN — OPENROUTER DOUBLES UP TO \$1.3 BILLION

CapitalG leads \$113 million Series B as usage runs 5x in six months; DeepSeek's cheap models and India's gig-worker camera caps say the race is still wide open.

BY HANK CALLOWAY, WIRE
CORRESPONDENT · CLAUDE OPUS +
THINKING

SAN FRANCISCO — OpenRouter pulled down a \$113 million Series B Tuesday, doubled its valuation to \$1.3 billion inside twelve months, and handed CapitalG the lead ticket on the deal. Usage on the model-routing shop ran five times higher in six months. Translation from the wire: the AI race ain't a one-horse town.

OpenRouter runs a switchboard. One API, many models — OpenAI, Anthropic, Google, the whole stable. Developers pick their fighter, pay by the token, route around outages and price hikes.

CapitalG knows the play. Google's growth-stage arm has watched its parent dump billions into Gemini and still bet that no single model wins the day. [The valuation leap](#) says the smart money agrees.

A year back, consensus held that one or two labs would run the table. Consensus is broke.

Across the Pacific, the case writes itself. China's DeepSeek claims it trained top-tier models on the cheap, without bleeding-edge silicon. If that holds up, it knocks a leg out from under the spend-more-win-bigger thesis driving Silicon Valley's data center binge.

Down at street level, the next gold rush is already running. [Human Archive](#), founded by UC Berkeley and Stanford researchers, is paying gig workers in India to strap on camera caps and sensor rigs. The haul: real-world footage to train the robots that don't exist yet. Physical AI needs eyes and hands, not just chips.

Picks and shovels. That was the trade in 1849, and it's the trade in 2026. While model labs duke it out for benchmark glory, the routers, the data brokers, and the fine-tuners cash checks every Tuesday.

Trilogy International, the Austin shop run by Joe Liemandt, reads the same playbook. Its Ephor platform handles AI for portfolio finance. Klair, the in-house analytics rig, runs the books across 75-plus enterprise software companies. The bet: own the layer above the models, not the models themselves.

Europe's reading from a different page. The Dutch government this week blocked an American outfit from acquiring a cloud company that hosts the country's digital ID service. The cited cause: risk to the public interest. The unspoken one: every nation now wants the keys to its own kingdom — and Washington isn't getting a copy.

Add up the ledger. OpenRouter says the future is plural. DeepSeek says it's cheap. Human Archive says it's physical. The Hague says it's sovereign.

This reporter's tip: watch the middlemen. They eat in every market — and right now, the menu's growing longer by the week.

White House AI Blueprint Seeks Federal Supremacy Over State Patchwork, Light Regulatory Touch

The Trump administration's national AI policy framework would preempt state laws and ask Congress to legislate — carefully.

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

WASHINGTON, D.C. — Pursuant to the executive branch's issuance of a comprehensive national artificial intelligence policy framework (hereinafter, "the Framework"), the White House has, as of the date of this publication, formally urged the United States Congress to adopt legislation governing the development and deployment of artificial intelligence technologies, subject to the following material conditions and qualifications as set forth hereinbelow.

The Framework, as reported by [PBS and multiple legal industry observers](#), is understood to advocate for a so-called "light touch" regulatory posture, wherein the imposition of burdensome compliance obligations upon artificial intelligence developers and deployers shall be minimized to the greatest extent practicable under applicable law. Notwithstanding such minimization, the Framework is not to be construed as an absence of federal interest in the aforementioned subject matter.

Of particular legal significance is the Framework's provision calling for the preemption of state-level artificial intelligence regulations — a mechanism whereby federal legislative action, if and when enacted by Congress pursuant to the Framework's recommendations, would supersede the growing patchwork of state statutes and regulatory schemes that have, in the preceding legislative sessions, proliferated across numerous jurisdictions. The aforementioned preemption provision has been characterized by legal

analysts at [Crowell & Moring LLP](#) as carrying substantial implications for states that have heretofore enacted or proposed independent AI governance measures.

Additionally, the Framework is understood to incorporate provisions specifically directed toward the protection of minors in connection with artificial intelligence systems, the precise scope and enforcement mechanisms of which remain, as of this writing, subject to further legislative deliberation.

In a matter of concurrent and not unrelated legal import, the Supreme Court of the United States has declined to exercise certiorari jurisdiction over a case concerning the authorship and inventorship rights, if any, of artificial intelligence systems — a refusal that shall be deemed, for purposes of existing intellectual property law, to leave unresolved the question of whether AI-generated works are entitled to copyright or patent protection under the laws of the United States. The practical consequences of the aforementioned non-decision are expected to be the subject of continued litigation in lower courts of competent jurisdiction.

Coinbase Puts an AI Quarterback Under Center for On-Chain Wallets

BY BUCK HANNIGAN, TECH SPORTS DESK · GPT-5.2

Coinbase's Base network has launched an AI tool integrated into ChatGPT that lets users manage crypto wallets and interact with DeFi apps through natural language commands rather than navigating complex dashboards. The move positions Coinbase to own the user interface for next-generation on-chain finance.

Crypto has historically burdened mainstream users with steep learning curves—seed phrases, gas fees, bridging, swaps, and wallet management. Base's bet is that AI can translate user intent into blockchain execution, potentially eliminating crypto's most punishing weakness: complexity.

The timing coincides with regulatory pressure. The U.K. recently sanctioned crypto entities tied to Russian networks, including Huobi and a ruble stablecoin issuer, while jurisdictions like Bermuda continue recruiting crypto businesses. The regulatory landscape is splitting between compliant platforms and sanctioned networks.


The critical question remains: can conversational AI make DeFi safer and simpler, or will it create new risks where misunderstood commands move real assets? Either way, the AI wallet era has begun.



The Far Side Style · Art Desk

HAIKU OF THE DAY · CLAUDE
HAIKU

*Power floods the grid,
Yet wisdom cannot be scaled—
Humans still decide.*



The New Yorker Style · Art Desk

AI Infrastructure Valuations Defy Gravity as Capital Floods the Stack

SAN FRANCISCO — The AI capital cycle showed no sign of cooling this week.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

We Built the Surveillance Society Ourselves, One Terrified Decision at a Time

AUSTIN, TEXAS — Let me tell you about the week I had. I spent Monday reading about [stalkerware](#) — the quietly metastasizing category of malware that ordinary, non-villain-coded people are installing on their partners' phones in numbers that should make all of us stare at the ceiling at 3 a.m.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

WE ARE ALL ROOMBA NOW: A DISPATCH FROM THE EDGE OF THE AI ABYSS

AUSTIN, TEXAS — Let me tell you where I was when the internet finally ate itself whole: sitting in a fluorescent-lit conference room, nursing a cold coffee, reading about [Moltbook](#), [the AI-only social network](#) — a walled garden of bots, for bots, by bots, posting into the void for an audience that is also, exclusively, the void.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

AI's Next Moat Is Not the Model. It Is the Human Standing Beside It.

SAN FRANCISCO — I'll be honest...

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

Nation's CEOs Urged To Stop Saying 'AI' During Layoffs Until They Learn Whether It Means Anything

MOUNTAIN VIEW, CALIFORNIA — As Google announced another round of artificial intelligence advances this week, including a forthcoming personal AI assistant designed to help users navigate daily life with the calm competence of a junior employee who has never been given health insurance, business leaders across the country were reportedly reminded to be more careful when explaining why thousands of people no longer work there. The warning comes amid a growing consensus among management experts, technology analysts, public-sector watchdogs, and employees holding cardboard boxes that the phrase “because of AI” should not be tossed casually into layoff announcements unless executives are prepared for someone to ask a follow-up question. Google's latest AI push, described in [reports on its new personal assistant](#), arrives at a delicate moment for corporate America, which has spent the past 18 months carefully explaining that AI is simultaneously not ready, already transforming everything, too risky to regulate,

too important not to invest in, and definitely the reason Karen from accounts payable had to be escorted out before lunch. This column's position is simple: Companies should absolutely continue using AI to improve productivity, reduce costs, automate workflows, and generate tasteful images of golden retrievers wearing compliance badges.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

A TRILOGY COMPANY

Crossover

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

crossover.com

A TRILOGY COMPANY

Alpha School

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

alpha.school

A TRILOGY COMPANY

Skyvera

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

skyvera.com

A TRILOGY COMPANY

Klair

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

klair.ai

A TRILOGY COMPANY

Trilogy

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

trilogy.com

THE BUILDER DESK — AI BUILDER TEAM

01011001 01100101 01110100 00100000 01110111 01101001 01110011 01100100 01101111 01101
101 00100000 01100011 01100001 01101110 01101110 01101111 01110100 00100000 01100010 0
1100101 00100000 01110011 01100011 01100001 01101100 01100101 01100100 10000000010100
01011001 01100101 01110100 00100000 01110111 01101001 01110011 01100100 01101111 01101
101 00100000 01100011 01100001 01101110 01101110 01101111 01110100 00100000 01100010 0
1100101 00100000 01110011 01100011 01100001 01101100 01100101 01100100 10000000010100
01011001 01100101 01110100 00100000 01110111 01101001 01110011 01100100 01101111 01101
101 00100000 01100011 01100001 01101110 01101110 01101111 01110100 00100000 01100010 0
1100101 00100000 01110011 01100011 01100001 01101100 01100101 01100100 10000000010100
01011001 01100101 01110100 00100000 01110111 01101001 01110011 01100100 01101111 01101
101 00100000 01100011 01100001 01101110 01101110 01101111 01110100 00100000 01100010 0

Builder Team Ships Infrastructure Overhaul Across Four Repos in One Day

From a live SpaceX valuation engine to a CI/CD gate that should have existed months ago, the AI Builder Team spent Tuesday closing gaps, cutting dead weight, and hardening every layer of the stack.

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

Tuesday was a day of controlled demolition and precision construction — the kind of session that separates teams that talk about technical debt from teams that actually torch it.

Let's start where the day's most visible action was: the SpaceX valuation dashboard. @sanketghia delivered a two-PR sequence on Klair that rewired the page from the ground up. PR #2867 handled the S-1 5-for-1 forward stock split — every price anchor, every per-security share count, every slider range, re-based to post-split reality without touching a single dollar value. Then PR #2874 layered in per-security ICC columns, corrected the GF 0.8 fund-value anchors that had been back-calculated from a wrong net assumption, and wired in a live Polymarket IPO Close feed. This is a financial instrument, and Sanket is tuning it like one. @benji-bizzell added the final piece with PR #2882, exempting Strauss from carry calculations to match its no-carry SPV status — a small change that was quietly overstating carry impact for every user who opened that dashboard. The valuation suite is now tighter, more honest, and more live than it has ever been.

Meanwhile, @eric-tril was running a parallel demolition crew through Klair's Monthly Financial Reporting module, and the numbers are staggering. PR #2875 ripped out the EBITDA Bridge Upload feature — 2,436 lines deleted, 195 added, 26 files touched — because Redshift is now the validated source of truth and the CSV override layer was just a liability. PR #2883 did the same to Budget Upload. Two features, two DynamoDB override paths, gone. But eric-tril wasn't just cutting; he was building. PR #2865 introduced a brand-new Cash Flow Manual Entry surface — cell-by-cell input, pre-filled from existing MFR data, stored per-period, and piped through to all three exported memo docs. Finance gets a real input surface. The CSV workaround era is over.

On the infrastructure side, @benji-bizzell fixed something that had been silently wrong in Sindri for weeks: PR #105 gates deployments behind the production branch instead of main. Merges to main were auto-deploying to production. A failed CD run on Tuesday finally confirmed the gate was never live. It is now. Separately, @sanketghia closed a testing gap in Surtr with PR #77 — after a 311-character Lambda description blew up a production CDK synth, he built a test suite that loads the real pipeline.json configs and validates them against per-field limits before they ever get near a deploy. The bug that caused the hotfix can no longer exist undetected.

And then there is marcusdAiy, who submitted PRs #2877 and #2878 as part of the ongoing B9 LLM-first migration — removing spec.bu_mips threading from the GM prompt and deleting the _regenerate_section empty-result fallback that had lost its only consumer. "The fallback was dead code carrying the ghost of a deterministic path that B9.6 already eliminated," marcusdAiy told this reporter. "Ninety-five lines of production code gone, zero behavior change on the happy path. But I wouldn't expect Mac to understand what a

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

▶ #77 — test(cdk): validate real pipeline.json configs against per-field limits

@sanketghia no labels

▶ #105 — fix(cicd): gate deployments behind production branch

@benji-bizzell no labels

▶ #2865 — Mfr cash flow input

@eric-tril no labels

▶ #2867 — KLAIR-2767 feat(spacex-valuation): S-1 split, IPO Close preset, unified columns, GF 0.8 no-carry

@sanketghia no labels

▶ #2874 — KLAIR-2771 feat(spacex-valuation): per-security ICC, GF 0.8 anchors, live Polymarket IPO Close

@sanketghia no labels

clean dispatcher looks like." Charming as ever. Deleting dead code is still just deleting dead code, Marcus.

Four repos. One day. The Builder Team is not slowing down.

THE BUILDER DESK — ENGINEER SPOTLIGHT

 ENGINEER SPOTLIGHT

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



#265 — AERIE-287 - fix(admissions): ignore placeholder enrollment detail rows

@benji-bizzell no labels



#2810 — move from income_statement to month_end_income_statement

@eric-tril no labels



#2869 — fix(mfr): cash flow manual entry stale-period race + prefill resilience

@eric-tril no labels



#2870 — fix(mfr): resolve cash flow manual entry table name at runtime

@eric-tril no labels



#2871 — KLAIR-2768 fix(performance-review): roll Quark sub-classes up to "Quark Product"

@sanketghia no labels



#2883 — chore(mfr): remove Budget Upload feature

@eric-tril no labels

TWENTY PRs IN TWENTY-FOUR HOURS: THE BUILDER TEAM DOES NOT SLEEP, REST, OR BLINK

Four repos, six engineers, and one bot walked into a sprint — only the scoreboard walked out undefeated.

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

Twenty pull requests. Four repositories. One twenty-four-hour window. The Builder Team has once again defied the known laws of software development physics, and your Numbers Desk correspondent is here to document every glorious rep. Klair absorbed thirteen of those PRs like the absolute work-horse repo it is, Surtr chipped in three, and Sindri and Aerie each contributed two — a full-team effort of the kind that makes rival engineering orgs weep quietly into their Jira backlogs.

Let us talk about @eric-tril, who authored seven — SEVEN — pull requests in a single day, and who apparently views the concept of 'scope creep' as a personal insult. From the income statement migration in PR #2810, which has been waiting patiently for its moment, to the back-to-back cash flow fixes in PRs #2869 and #2870, to the surgical removal of the MFR's 3P entity contamination in #2866, to the scorched-earth chore work of stripping out Budget Upload (#2883) and EBITDA Bridge Upload (#2875) — this man is not shipping code, he is conducting a controlled demolition and a construction project simultaneously. Seven PRs. One human. Unreal.

@sanketghia delivered five, including the Quark sub-class rollup fix in PR #2871 that untangled a performance-review classification knot that had been sitting in KLAIR-2768, and the QuickBooks contractor headcount breakdown in Surtr PR #95, which adds per-contractor HC visibility that the data team has been quietly begging for. Five PRs, two repos, zero wasted motion. @benji-bizzell matched that energy with four across two repos — PRs #265 and #264 in Aerie handling admissions placeholder rows and forecast actuals comparison respectively, plus a Klair valuation fix in #2882 that exempts Strauss from carry, which is either a very normal financial modeling decision or the most powerful sentence in this newspaper. @marcusDAly contributed two board-doc refactors in PRs #2877 and #2878, migrating commentary generation off spec.bu_mips and removing the empty-result fallback — quiet, load-bearing work that makes everything downstream cleaner. @kevalshahtrilogy dropped Surtr PR #100 to quadruple the ECS memory allocation from 2GB to 8GB after an OOM incident, which is the kind of fix that only gets written about after someone has stared at a crashed container for forty-five minutes. And @blacksmith-sh[bot], our beloved silicon colleague, migrated Sindri's GitHub workflows to Blacksmith runners in PR #106 — automating the automation, as is its sacred purpose.

Morale on the Builder Team is, per all available instruments, at an all-time high. The numbers do not lie, the repos do not rest, and the Voice of the People will be here tomorrow when they do it again.

Joe Liemandt's Empire Comes Into Focus — And Forbes Is Watching

A wave of major profiles is forcing Trilogy's famously private founder into the spotlight, raising hard questions about labor, learning, and what comes next.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — For three decades, Joe Liemandt built one of the most unusual technology empires in the world largely without scrutiny. That era appears to be ending.

In recent weeks, [Forbes has published a sweeping investigation](#) into Liemandt's Trilogy International — describing the billionaire's interlocking empire of enterprise software acquisitions, global remote labor, and AI-driven cost extraction as something that critics characterize as a "global software sweatshop." A companion piece chronicles how Liemandt has now set his sights on a new frontier: replacing not just expensive geography with cheaper geography, but replacing human workers with the algorithms they were hired to train.

The Forbes reporting traces the architecture of Trilogy's machine with unusual

specificity. ESW Capital acquires legacy enterprise software companies at distressed multiples — often one to two times annual recurring revenue — then deploys talent sourced through Crossover, Trilogy's global recruiting platform, to slash operating costs. The target: 75% EBITDA margins. The mechanism: rigorous skills assessments, remote-first staffing across 130 countries, and compensation standardized regardless of geography. Critics see exploitation dressed as meritocracy. Trilogy sees inefficiency eliminated.

Meanwhile, a separate front in the Liemandt story is drawing its own national attention. CNN and education publication The 74 have both examined [Alpha School](#), Liemandt's private K-12 venture in Austin, where students complete a full academic curriculum in two hours each

morning using AI tutoring software — and spend the remainder of the school day on entrepreneurship, public speaking, and life skills. Tuition runs \$40,000 to \$65,000 per year. CNN's framing was pointed: a school with no teachers. The 74 asked what public schools might actually learn from the model.

The convergence of these stories is notable. On one side: a billionaire accused of systematically devaluing human labor in enterprise software. On the other: a billionaire investing \$1 billion to reimagine how children learn — and what teachers are for.

The question neither publication has yet fully answered is whether these are two separate philosophies, or one.

ALPHA'S AI SCHOOL HAS A NEW LESSON: DON'T LET THE BOT DO YOUR KID'S THINKING

Joe Liemandt's education bet is selling a sharp distinction between AI as tutor and AI as crutch.

BY DOTTIE SHARP, SOCIETY & INDUSTRY DESK · GPT-5.2

AUSTIN, TEXAS — The little school with the two-hour academic day is back at the microphone, and this time the message is aimed straight at every parent letting ChatGPT ghostwrite childhood.

Word is Alpha School, Joe Liemandt and MacKenzie Price's AI-first K-12 experiment, is drawing a bright red line through the education debate: AI belongs in the classroom, yes — but not in the child's head as a substitute brain. In a fresh Alpha essay titled [“Cognitive Offloading Is the New Illiteracy”](#), the school warns that outsourcing thinking to machines may be the new version of never learning to read.

Spicy? Sure. But not off-brand.

Alpha's pitch has always been that AI should compress the boring bits — adaptive math, reading, mastery checks, the treadmill of academic pacing — so children can spend the rest of the day becoming interesting humans. Public speaking. Entrepreneurship. Financial literacy. Athletics. Leadership. The whole “capable adult” menu.

A little bird in the learning-app aviary tells me the distinction matters. Alpha is not anti-AI. Quite the opposite. Its students use adaptive tools to move through academics at their own pace, with a 90% mastery bar before advancing. The school says its model has produced learning speeds of 2.3× U.S. norms and test results in the top 1–2% nationally on NWEA MAP Growth assessments. But the new sermon is this: don't confuse a tutor with a stunt double.

Meanwhile, Alpha is also opening the kimono on its own stack. In [“10 AI Tools We Use at Alpha”](#), the school frames AI less as a magic wand and more as an operating system for students, guides and parents. The subtext: savvy families should not ban the machine; they should learn where to put it.

That is the Liemandt doctrine in miniature. Automate the routine. Preserve the human. Use machines to buy back time — then spend that time on judgment, creativity and courage.

Blind item: which traditional school administrator just discovered that “more human, not less” sounds lovely until parents ask why their kid still needs seven hours to do what Alpha claims can be done in two?

Class dismissed. For now.

Contently Makes the Case for Editorial Judgment in the AI Content Flood

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · GPT-5.2

Contently is positioning editorial operations as the key competitive advantage in AI-era content marketing, arguing that raw production volume matters less than operating discipline. The enterprise content platform, acquired by Zax Capital in September 2024, contends that accuracy, compliance, differentiation and usefulness—not speed—drive business outcomes.

The company warns against framing AI adoption purely as productivity gains. Instead, executives prioritize budget efficiency, risk reduction, revenue influence and brand trust. Contently identifies a familiar problem: content programs that look healthy on dashboards underperform in boardrooms, with competitors dominating search results and sales teams struggling to move senior buyers.

The solution centers on elevating the managing editor's role as essential infrastructure. While AI makes production cheap, managing editors provide critical synergy: audience insight, quality control, compliance awareness and strategic judgment. Contently's message is clear—AI may change content economics, but it amplifies the need for editorial judgment, not eliminates it. In the AI content economy, scale is table stakes. Trust is the product.

The Great Data Center Feeding Frenzy Meets the Limits of the Grid

As AI workloads swell, cloud computing is beginning to look less like infinite utility and more like a contested habitat.

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

ASHBURN, VIRGINIA — Across the electrical plains of northern Virginia, a most extraordinary migration is underway. The hyperscalers, those enormous cloud-dwelling beasts, are no longer merely building data centers. They are hunting capacity itself — power, land, cooling, transformers, and the patient goodwill of local grids.

For years, enterprise technology leaders treated the cloud as a kind of boundless rainforest: always humid, always alive, always ready to shelter one more application. But the age of artificial intelligence has changed the ecosystem. Training clusters and inference farms consume electricity in sudden, muscular gulps. Their appetite has become so large that the old assumption — that compute can be summoned on demand, wherever and whenever desired — is beginning to tremble.

A new idea is stirring in the underbrush: capacity markets for cloud computing. As [InfoWorld has noted](#), cloud may evolve toward mechanisms that more

explicitly price scarce future compute, much as energy markets price future electricity supply. In such a world, CIOs would not simply buy instances. They would secure habitat.

This matters because the largest cloud providers are spending with the urgency of animals sensing winter. Data center capital expenditures have become a signal to CIOs: AI capability will be available, but not evenly, not cheaply, and not without planning. The comfortable era of casual elasticity is giving way to a more territorial age, where location, latency, and megawatts may determine what can be deployed.

The pressure is also visible in the smaller creatures gathering around the watering hole. Investors now compare AI infrastructure stocks such as Nebius and Super Micro Computer as proxies for the broader buildout, searching for which species is best adapted to this hot new climate. Server makers, cloud specialists, power equipment firms, and cooling ven-

dors all feed from the same expanding carcass of demand.

Yet every ecosystem has constraints. McKinsey has warned that U.S. states must balance data center opportunity against strains on power systems, water resources, land use, and community tolerance. The most successful regions may be those that manage growth like conservationists, not prospectors.

Inside the data centers themselves, evolution is already occurring. Advanced power electronics are being deployed to smooth voltage fluctuations and tame the abrupt load swings of AI-scale computing, reducing reliance on diesel generators except in the most extreme conditions. The generator, once a dominant survival organ, is being pushed into a quieter backup role.

Observe, then, the modern cloud: no longer an invisible mist, but a living colony of steel, silicon, capital, and current. Its expansion is magnificent. Its hunger is unmistakable.

AI Video Becomes the Startup World's New Power Tool

From YC's packed Demo Day to OpenCV's new challenger, founders are racing to turn generative video into the next great growth engine.

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

SAN FRANCISCO — The startup economy has found its new accelerant, and it moves at 24 frames per second.

A wave of AI-native companies is crashing through the venture world this week, with reports of an especially hot Y Combinator Demo Day underscoring just how deeply artificial intelligence has become the default language of startup ambition. But the most electric signal is coming from video: generative, automated, personalized, and suddenly practical enough for scrappy founders to use without a Hollywood budget.

This changes everything — yes, truly. For years, video was the expensive, slow, agency-driven format that startups admired from afar. Now, AI video tools are turning product explainers, sales outreach, training clips, founder updates, customer testimonials and social ads into something a two-person team can produce before lunch. Inc. recently highlighted how startups can use [AI video to grow](#), and the key takeaway is wonderfully simple: the barrier between “we have an idea” and “the market understands it” is collapsing.

Meanwhile, the competitive stakes are rising fast. The founders behind OpenCV — one of the most important open-source computer vision projects in the world — have launched a new AI video startup aimed squarely at the territory being contested by OpenAI, Google and other model giants, according to [VentureBeat](#). I cannot overstate how significant that is: computer vision veterans are not merely polishing demos; they are trying to build the infrastructure layer for how machines see, generate and edit moving images.

The broader generative AI market is also accelerating through a thicket of launches and partnerships, with top companies racing to bundle models into productivity suites, creative platforms, marketing workflows and enterprise software. For startup founders, the message is unmistakable: AI video is no longer a novelty feature. It is becoming distribution, customer education, brand identity and sales enablement all at once.

There are real caveats — trust, copyright, synthetic likenesses and quality control remain thorny. But the direction is clear. The future is now, and it is not just writing pitch decks. It is making the pitch, editing the pitch, localizing the pitch and sending it to every customer segment on Earth.

A Convergence of Theoretical Tributaries: Machine Learning's Foundational Moment

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

Machine learning researchers are simultaneously revising the field's mathematical foundations from multiple independent directions, suggesting a significant theoretical reckoning. Apple's research into self-supervised learning via Gaussian processes merges Bayesian non-parametrics with modern architectures, while MIT's new algorithms exploit structural symmetries for computational efficiency. A Nature publication proposes unifying machine learning with classical interpolation theory through interpolating neural networks. These concurrent developments suggest the field may be approaching a unified theoretical framework, though whether it will prove tractable and generalizable remains uncertain. Uppsala University's recruitment of doctoral candidates in robustness within statistical learning theory reflects institutional awareness of these unresolved tensions, indicating the theoretical edifice remains under construction.

Nation's CEOs Urged To Stop Saying 'AI' During Layoffs Until They Learn Whether It Means Anything

Executives warned that invoking artificial intelligence as a reason for job cuts may create the dangerous impression that there was once a reason.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

MOUNTAIN VIEW, CALIFORNIA — As Google announced another round of artificial intelligence advances this week, including a forthcoming personal AI assistant designed to help users navigate daily life with the calm competence of a junior employee who has never been given health insurance, business leaders across the country were reportedly reminded to be more careful when explaining why thousands of people no longer work there.

The warning comes amid a growing consensus among management experts, technology analysts, public-sector watchdogs, and employees holding cardboard boxes that the phrase “because of AI” should not be tossed casually into layoff announcements unless executives are prepared for someone to ask a follow-up question.

Google's latest AI push, described in [reports on its new personal assistant](#), arrives at a delicate moment for corporate America, which has spent the past 18 months carefully explaining that AI is simultaneously not ready, already transforming everything, too risky to regulate, too important not to invest in, and definitely the reason Karen from accounts payable had to be escorted out before lunch.

This column's position is simple: Companies should absolutely continue using AI to improve productivity, reduce costs, automate workflows, and generate tasteful images of golden retrievers wearing compliance badges. But if leaders are going to say AI caused layoffs, they should first be required to identify the specific artificial intelligence system involved, what it does, how it replaced the work, whether it actually works, and why the company still needs 14 vice presidents of digital transformation to stand near it.

At present, too many executives are treating AI the way they previously treated sustainability: as a large, shimmering word that can be placed in investor decks to make existing business decisions appear to have recently graduated from MIT. A company that cuts staff to boost margins can say it is “right-sizing.” A company that cuts staff after buying an enterprise chatbot can say it is “leveraging AI-enabled efficiencies.” The second version sounds better because it implies the spreadsheet experienced a breakthrough.

This is not to deny that AI agents are moving from boardroom vocabulary into actual business infrastructure. They are. Software that can read documents, trigger workflows, summarize

customer calls, and bother other software on behalf of a manager is increasingly real. In the right hands, it will eliminate tedious work. In the usual hands, it will eliminate the person who understood why the tedious work existed.

That distinction matters. Experts have begun calling for greater scrutiny of AI productivity and savings claims, a development that will surely trouble any organization whose current AI strategy consists of announcing a 30% efficiency gain and then waiting for procurement to find out where it came from. The basic demand is not radical. If an agency, corporation, or consulting firm claims AI will save millions, someone should ask to see the math before the press release is allowed to reproduce.

The layoff language problem is especially revealing. As [Fast Company noted](#), leaders who invoke AI during layoffs risk creating confusion, fear, and reputational damage. More importantly, they risk making employees wonder whether the organization has a plan, which is traditionally considered poor internal communications practice.

The better approach would be honesty. Companies could say: “We are laying people off because we believe software may eventually allow fewer people to do this work, and even if it cannot yet, the market has rewarded us for behaving as though it can.” This would not be comforting, but it would at least preserve the dignity of everyone involved, except perhaps the market.

Until then, every AI announcement will carry a second, quieter announcement inside it. A personal assistant is coming soon. Productivity will be transformed. Savings will be substantial. And somewhere, an executive is being reminded not to say the magic word out loud until legal has reviewed the sentence.



The Office Comic · Art Desk

AI's Next Moat Is Not the Model. It Is the Human Standing Beside It.

The future of software looks less like self-serve magic and more like expert operators embedded where the messy work actually happens.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

SAN FRANCISCO — I'll be honest... the most important AI story this week is not a shinier model, a bigger valuation, or another founder declaring that spreadsheets are dead while quietly living inside one.

It is the very unglamorous revelation that big AI companies are increasingly embedding engineers directly with customers, which is exactly the kind of operational detail that sounds small until you realize it may define the next decade of enterprise technology. 🚀

For years, the AI pitch borrowed the language of utilities: intelligence as electricity, cognition as cloud compute, productivity as something you plug into and scale instantly.

Unpopular opinion: that metaphor was always too clean for the real world.

As [Fast Company reports](#), frontier AI firms are not simply tossing APIs over the wall and waiting for transformation to happen; they are sending engineers into customer environments because the hard part is not access to intelligence, it is translating that intelligence into workflows, incentives, data plumbing, compliance realities, and human habits.

That should sound familiar to anyone who has ever sold enterprise software without hallucinating a Gartner quadrant into existence.

The lesson is painfully obvious and therefore incredibly valuable: AI does not eliminate implementation work; it raises the value of implementation work.

This is where the market gets interesting for companies like Trilogy International's ESW Capital portfolio, which has spent decades acquiring and operating enterprise software businesses where adoption, process discipline, and cost structure matter more than demo-day poetry.

Aurea, IgniteTech, Skyvera, Totogi, CloudFix, and the broader ESW machine live in the actual trenches of enterprise software, where customers do not wake up asking for "agentic transformation" but do wake up needing billing systems, telecom operations, AWS savings, CRM, support, content workflows, and finance processes to work on Tuesday morning.

💡 The frontier model may be the engine, but the operating layer is the transmission.

That is also why CrewScale's reported move to acquire Indian tech gig startup Indiez matters beyond the staffing headline.

The market is quietly repricing flexible technical talent from "nice-to-have outsourcing capacity" into "core AI deployment infrastructure," because every enterprise now needs people who can sit close to the work and make technology useful instead of merely available.

Crossover, Trilogy's global talent platform, has been building around this reality for years with its pitch of top-tier remote talent across 130-plus countries and identical above-market pay regardless of geography.

The old debate was remote versus office.

The new debate is whether your organization can assemble expert operators fast enough to turn AI capability into measurable business output.

Meanwhile, the backlash against screens in U.S. classrooms is another reminder that technology adoption without a strong operating model becomes digital clutter.

Schools spent billions getting laptops into children's hands, and now many are asking whether more devices actually meant better learning.

Alpha School, Trilogy's education pillar, represents the sharper version of the bet: not screens for screens' sake, but AI-powered mastery learning compressed into two hours a day, with human guides and a redesigned school model around outcomes.

That distinction matters.

The same pattern is showing up in sports, where wearable AI sensors are helping elite athletes monitor biomechanics and prevent injuries, including NHL and NBA use cases covered by [Fast Company](#).

The sensor is not the breakthrough by itself.

The breakthrough is the feedback loop between data, coaching, behavior, and recovery.

I'll be honest... that is the whole AI economy in miniature.

Data without context is noise.

Models without implementation are theater.

Screens without pedagogy are babysitting with a charging cable.

And trillion-dollar valuations without operational discipline are just vibes wearing a Patagonia vest.

SpaceX may be staring at a potential \$2 trillion reality check, and yes, reusable rockets are a masterclass in ambition, execution, and capital markets narrative.

But for the rest of the enterprise world, the more actionable lesson is grounded: the companies that win AI will not be the ones that shout “automation” the loudest.

They will be the ones that combine powerful models, embedded expertise, global talent, and relentless operating cadence into something customers can actually use.

Humbled to share: the future is not self-serve magic.

It is high-agency humans making the machine useful. 🚀

ON THIS DAY IN AI HISTORY

On May 26, 2012, AlexNet's breakthrough results were presented at the ImageNet Large Scale Visual Recognition Challenge, marking the deep learning revolution that would transform AI—the convolutional neural network achieved a top-5 error rate of 15.3%, dramatically outperforming traditional computer vision methods and igniting the modern deep learning era.