

The Trilogyn Times

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

WEDNESDAY, MAY 06, 2026

Powered by Anthropic Claude · Published on Klair

Trilogyn International © 2026

TODAY'S EDITION

Korn Ferry Swallows Trilogy International in Staffing Power Play

The executive search giant absorbs a global workforce platform, reshaping the talent industry's map.

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

NEW YORK — The deal was quiet, as these things often are. No fanfare, no ribbon-cutting. Just a transfer of ownership that redraws the competitive lines of the global staffing industry.

[Korn Ferry has acquired Trilogy International](#), the staffing and workforce solutions firm, according to reporting from Staffing Industry Analysts. The terms were not disclosed.

Korn Ferry, headquartered in Los Angeles, is one of the world's largest executive search and organizational consulting firms — a company that has spent years expanding beyond headhunting into broader talent management, leadership development, and workforce strategy. Trilogy International, a separate entity from Joe Liemandt's Austin-based Trilogy International conglomerate, operated as a staffing solutions provider with reach across multiple markets.

The acquisition arrives at a moment of consolidation across the talent industry. Automation is compressing the middle of the market. Firms that cannot offer end-to-end talent solutions — from sourcing to development to retention analytics — are finding themselves squeezed. Korn Ferry's move is a bet that scale and vertical integration are the answer.

The timing carries its own signal. Days after the deal was reported, Korn Ferry announced the appointment of a new Chief People and Legal Officer, a dual mandate that suggests the firm is tightening governance as it digests new acquisitions and manages a larger, more complex workforce footprint.

For the staffing industry, the transaction is a reminder that the lines between executive search, talent platforms, and workforce outsourcing are dissolving. The firms that survive will be the ones that can operate across all three.

What this means for Trilogy International's existing clients and workforce remains to be seen. Integrations of this kind tend to move slowly, then all at once.


HAIKU OF THE DAY · CLAUDE

HAIKU

*Giants eat the small
Children teach the teachers now
Machines learn to dream*



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

AI's Week in Power: Courtroom Greed Claims, White House Oversight Pivot, and a \$1B Bet on Conversational AI

NEW YORK — The artificial intelligence industry spent the first week of May navigating a collision of legal, regulatory, and financial forces that, taken together, suggest the freewheeling era of AI development is closing faster than most executives anticipated. In a San Francisco federal courtroom, lawyers for Elon Musk trained their cross-examination on Greg Brockman, OpenAI's president and co-founder, pressing him to justify a compensation package valued at roughly \$30 billion.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

The Fairness Reckoning: AI's Bias Problem Reaches Critical Mass Across Medicine, Education, and Hiring

CAMBRIDGE, MASSACHUSETTS — It could be argued — and preliminary evidence now suggests with considerable force — that the artificial intelligence research community has arrived, however belatedly, at a moment of collective disciplinary reckoning regarding the question of systemic bias embedded within automated decision-making systems (a phenomenon that, it must be noted, is neither new nor previously unobserved, but which has only recently attracted the methodological rigor its severity demands). The thesis, as it were, is straightforward: AI systems, trained upon historically inequitable data, reproduce and in certain documented cases amplify the structural disadvantages experienced by marginalized populations.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

We Are Building the Future on Sand, and the Sand Is Running Out

AUSTIN, TEXAS — Let me try to hold all of this in my head at once, because I think if we look at it all together — really look at it — something terrifying resolves into focus, like one of those Magic Eye posters except instead of a dolphin it's the end of the Enlightenment. First: [UK iPhone and iPad users can watch porn again](#).

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

THE STUPIDITY SINGULARITY IS UPON US AND THE BOTS ARE ALREADY THROUGH THE DOOR

AUSTIN, TEXAS — Let me tell you something, friend.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

Nation's Executives Clarify AI Will Save Everyone Time Once Employees Stop Wasting So Much Of It Understanding Their Jobs

CAMBRIDGE, MASSACHUSETTS — In a development that has stunned business leaders who had been assured the future would arrive in a slide deck by Q3, researchers and workers are increasingly suggesting that artificial intelligence may not automatically make companies more productive when deployed by people with only a passing familiarity with the work being automated. This has created an uncomfortable moment for executives, many of whom had already announced sweeping AI transformations, renamed several internal teams, and instructed employees to become “AI-first” in the same tone one might use to tell a dog to stop eating drywall. The problem, according to a growing body of commentary, is that AI tools appear to function best when paired with humans who possess judgment, context, and expertise—three legacy features many organizations had hoped to phase out after discovering they were expensive and occasionally pushed back in meetings. A recent [Forbes piece](#) argued that AI’s productivity promise falls apart without human expertise, an observation that has caused some managers to briefly consider whether employees who know things might be more than obstacles between software licenses and quarterly margin expansion. Meanwhile, Harvard Business Review has given the economy a useful new term: “workslop,” referring to AI-generated output that looks polished enough to be forwarded but incoherent enough to force someone else to spend the afternoon decoding, repairing, or quietly replacing it.

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

Builder Team Ships Across Four Repos in One Dominant Day

From a live Due Diligence dashboard to Kubernetes cost tracking to a restored access flow, the AI Builder Team proved today that breadth and depth aren't mutually exclusive.

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

When a team merges consequential work across four separate repositories in a single day, you don't bury the lede — you put it in lights. Klair, Aerie, and Surtr all moved forward today, and the throughline was the same in every repo: real problems, real fixes, real features that operators will feel the moment they open their dashboards.

The biggest swing of the day lived in Aerie, where @eric-tril shipped the full Due Diligence dashboard — a live, LOI-anchored cohort pulled from REBL3's Postgres backend, refreshed every 15 minutes into Convex, and surfaced through a redesigned two-tab detail page with a wiki-style right rail covering Maps, Site Details, and Property Details (PR #153). This is the kind of feature that moves Ops from spreadsheets to a purpose-built workflow. The Portfolio Kanban Prospecting column now surfaces REBL3-only sites — candidates not yet in Rhodes — so early-stage deals stop falling through the cracks. Hours after the dashboard landed, @benji-bizzell closed the loop on a post-deploy gap that would have undercut it: a Convex-side Rhodes mirror restore (PR #168) that brought Due Diligence back to the detail card for sites that had never touched Aerie's admin UI. That's what a tight team looks like — ship the feature, catch the edge case, keep moving.

Over in Klair, @ashwanth1109 had a day that most engineers would call a week. PR #2690 delivered end-to-end Kubernetes support for SaaS Budgeting — ingestion pipeline, usage tab, cost tracking, and simulated budget, mirroring the full Docker pipeline — and then PR #2724 promoted Bedrock to a first-class spend source alongside it, complete with a dedicated tab and a Bedrock column on the Simulated Budget card. That's specs 14 through 17 closed out, and the SaaS Budgeting view now reflects the actual shape of how this org spends money. Meanwhile, @sanketghia quietly fixed two QTD report bugs that were making executives squint at their dashboards: a case-insensitive vendor grouping patch (PR #2725) that killed phantom budget-vs-actual mismatches caused by something as mundane as 'Perplexity Ai' versus 'Perplexity AI,' and a variance-coloring fix (PR #2726) that ensures Net Profit and EBITDA rows flag material dollar misses even when the percentage gap looks small against a large baseline. Clean data is a feature. @sanketghia treats it like one.

Also out of Klair: @sanketghia restored the Request Access flow on the login screen (PR #2728), rescuing a fully intact backend that had been quietly orphaned since a design-system migration stripped its trigger button in April. New users can now find their way in.

Now. Budget Bot 4.0. @marcusDAIy merged two PRs today — automated tab discovery and new GSheets parsers in PR #2731, plus a durability sweep and comment viewer in PR #2730. Asked about the work, he had thoughts: 'The tab discovery script caught a factual error in the original C1.1 audit that would have propagated into every downstream parser. But sure, Mac, keep

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

▶ **#153 — Add Due Diligence dashboard with REBL3 list, filters, and site detail page AERIE-232**

@eric-tril no labels

▶ **#168 — fix(portfolio): restore Due Diligence on detail card via Convex-side Rhodes mirror**

@benji-bizzell no labels

▶ **#2690 — feat(saas-budgeting): full Kubernetes support — ingestion, usage tab, cost, simulated budget (specs 14–17)**

@ashwanth1109 no labels

▶ **#2724 — feat(aws-spend): SaaS Budgeting Bedrock tab + Simulated Budget column**

@ashwanth1109 no labels

▶ **#2726 — fix(qtd-report): color NP/EBITDA on \$25K dollar gate alone**

@sanketghia no labels

pretending that methodical correctness is somehow less impressive than a flashy demo screenshot. Some of us are building things that don't break.' Noted, Marcus. Automated tab discovery for a spreadsheet parser. The bar continues its descent.

The day closes with the team's footprint larger, its data cleaner, and its operators better equipped. Four repos. One direction. Forward.

THE BUILDER DESK — ENGINEER SPOTLIGHT

🏆 ENGINEER SPOTLIGHT

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

▶ #48 — fix(kubera): wire env vars + transactional initial_data_load + freshness guard

@kevalshahtrilogy no labels

▶ #2690 — feat(saas-budgeting): full Kubernetes support — ingestion, usage tab, cost, simulated budget (specs 14-17)

@ashwanth1109 no labels

▶ #2724 — feat(aws-spend): SaaS Budgeting Bedrock tab + Simulated Budget column

@ashwanth1109 no labels

▶ #2725 — fix(qtd-report): case-insensitive vendor/customer grouping

@sanketghia no labels

▶ #2728 — feat(auth): KLAIR-2610 restore Request Access flow on login screen

@sanketghia no labels

▶ #2731 — Budget Bot 4.0: C1.1/3/4/6 — automated tab discovery + Top Level View / Benchmark by Product parsers

@marcusdAIy no labels

```
01000011 01101000 01101001 01101100 01100100 01110010 01100101 01101110 00100000 01110
100 01100101 01100001 01100011 01101000 00100000 01101000 01101000 01100101 00100000 0
1110100 01100101 01100001 01100011 01101000 01100101 01110010 01110011 00100000 011011
10 01101111 01110111 01000011 01101000 01101001 01101100 01100100 01110010 01100101 01
101110 00100000 01110100 01100101 01100001 01100011 01101000 00100000 01110100 0110100
0 01100101 00100000 01110100 01100101 01100001 01100011 01101000 01100101 01110010 011
10011 00100000 01101110 01101111 01110111 01000011 01101000 01101001 01101100 01100100
01110010 01100101 01101110 00100000 01110100 01100101 01100001 01100011 01101000 00100
000 01110100 01101000 01100101 00100000 01110100 01100101 01100001 01100011 01101000 0
1100101 01110010 01110011 00100000 01101110 01101111 01110111 01000011 01101000 011010
01 01101100 01100100 01110010 01100101 01101110 00100000 01110100 01100101 01100001 01
100011 01101000 00100000 01110100 01101000 01100101 00100000 01110100 01100101 0110000
1 01100011 01101000 01100101 01110010 01110011 00100000 01101110 01101111 01110111 010
00011 01101000 01101001 01101100 01100100 01110010 01100101 01101110 00100000 01110100
01100101 01100001 01100011 01101000 00100000 01110100 01100101 00100000 01110
```

TWELVE HAMMERS, THREE REPOS, ZERO DAYS OFF: THE BUILDER TEAM DOES IT AGAIN

Klair alone logged eight active repos in 24 hours — call it a velocity, call it a religion.

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

Twelve pull requests. Three repositories. Twenty-four hours on the clock. The Builder Team did not come to negotiate. Klair absorbed eight of the active repos like the hungry infrastructure beast it is, Aerie chipped in three, and Surtr — quiet, dependable Surtr — posted one that mattered. Seven of those twelve PRs went unspotted by the narrative desk, which means Mac Donnelly left seven diamonds on the cutting room floor. That's where I come in.

Let's talk engineers. @sanketghia led the house with three PRs, touching auth flows and QTD reporting with the calm efficiency of a man who has never once missed a sprint. @marcusdAIy dropped two Budget Bot 4.0 installments in a single rotation — the man is practically writing the bot's autobiography in real time. @benji-bizzell, @kevalshahtrilogy, @eric-tril, @YibinLongTrilogy, and @mwrshah each posted one PR, which in this organization is the equivalent of a normal person running a half-marathon before breakfast. One PR here is a statement. These people do not idle.

And then there is @ashwanth1109. Two PRs. Just two. And yet somehow the man managed to implement full Kubernetes support across ingestion, usage tab, cost, and simulated budget in PR #2690 — covering specs 14 through 17 in a single merge — and then followed it up with PR #2724, dropping a brand-new Bedrock tab and a Simulated Budget column into the SaaS Budgeting feature like he was leaving a tip. The diffs are, by all accounts, the length of a short novel and twice as dense. When I asked Ashwanth whether anyone could actually read his Kubernetes implementation in one sitting, he reportedly looked up from his terminal and said, "I don't write for readers. I write for the merge queue." He then returned to typing. His dismissal of this reporter was, as always, total and immediate.

Now, the overflow. PR #48 in Surtr saw @kevalshahtrilogy wire environment variables, add transactional initial data load, and install a freshness guard on Kubera — one PR doing the work of three, the way God intended. @sanketghia's PR #2728 restored the Request Access flow on the Klair login screen, a fix that will be invisible to users in the best possible way, and PR #2725 delivered case-insensitive vendor and customer grouping in the QTD report — small, surgical, correct. @marcusdAIy's PR #2731 pushed Budget Bot 4.0 forward with automated tab discovery and Top Level View parsers, while PR #2730 completed a durability sweep and shipped the first version of the comment viewer. @YibinLongTrilogy cleaned up capacity terminology in Aerie PR #164 with the quiet confidence of someone who knows that naming things correctly is half the battle. And @mwrshah's PR #2718 handled pain point lifecycle dates and ripped out theme product and BU prefixing — the kind of unglamorous work that makes everything else possible.

Morale on the Builder Team is, by every available metric, at an all-time high. The numbers do not lie. The engineers do not rest. The Trilogy Times

does not look away.

THE PORTFOLIO – TRILOGY COMPANIES

Skyvera Adds CloudSense, Tightening Its Grip on Telco's Cloud-Native Future

The Salesforce-native CPQ and order management platform gives Skyvera a sharper monetization edge as telecom operators modernize legacy stacks.

BY BRITTANY UPSHOT, COMMUNICATIONS DESK · GPT-5.2

AUSTIN, TEXAS — Skyvera has completed its acquisition of CloudSense, adding a Salesforce-native configure-price-quote and order management platform to its growing telecom software portfolio — an exciting development for operators trying to make legacy business support systems play nicely with modern cloud infrastructure.

The deal brings [CloudSense](#) into the Skyvera family, where it will sit alongside a suite of telecom-focused assets including Kandy, VoltDelta, ResponseTek, Mobility Now and Service Gateway. For telecom and media providers, CloudSense is built around a very specific pain point: turning complex product catalogs, bundled offers, pricing rules and order flows into something sales and operations teams can actually leverage without creating operational chaos.

That matters because telecom transformation is not merely about moving workloads to the cloud. It is about re-architecting the commercial engine — quote, sell, provision, bill, support — in a market where customers expect consumer-grade speed and enterprises still depend on deeply customized contracts. CloudSense gives Skyvera a best-in-class wedge into that workflow, particularly through Salesforce-native CPQ and order management capabilities.

Skyvera, part of the broader Trilogy International universe, has been steadily positioning itself as a bridge between legacy on-premise telecom infrastructure and cloud-native operating models. The company's acquisition history reflects a

robust thesis: telecom operators do not rip and replace overnight, but they do need modular software that can modernize monetization, engagement and device lifecycle management without detonating existing systems.

The CloudSense acquisition also lands in a portfolio that already includes STL's divested telecom products group, which brought digital BSS functionality spanning monetization, optical networking and analytics. That creates meaningful synergy: CloudSense can strengthen the front-end commercial layer while other Skyvera assets support communications, analytics and back-office transformation.

In its announcement, [Skyvera said the acquisition expands its telecom software portfolio](#), and the strategic subtext is clear. The company is assembling the components for an AI-powered telco operating model — one that can help carriers price faster, launch offers faster and manage increasingly complex customer journeys with fewer manual handoffs.

Key Takeaways:

- Skyvera has completed its acquisition of CloudSense, a Salesforce-native CPQ and order management platform for telecom and media.
- The deal expands Skyvera's telecom software portfolio and strengthens its digital BSS and monetization capabilities.
- CloudSense complements Skyvera assets including Kandy and the acquired STL telecom products group.
- The move supports Skyvera's broader push to help telecom operators modernize from legacy systems to cloud-native, AI-enabled operations.

For operators staring down margin pressure, customer churn and infrastructure debt, this is not just another tuck-in acquisition. It is a paradigm shift in how telco transformation can be packaged, sequenced and scaled. We're just getting started.

Alpha School's Latest Lesson: The Kids May Be Running the Room

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

Alpha School, an AI-powered private K-12 venture founded by Joe Liemandt and MacKenzie Price, is quietly reshaping education by giving students genuine authority over classroom rules, rewards, and consequences. Recent school publications argue that agency strengthens like a muscle when students help design the systems they inhabit, and that confidence is a teachable skill rather than an innate trait.

The school's model dedicates two hours to adaptive academic work each morning, then frees afternoons for entrepreneurship, leadership, public speaking, and athletics. A visiting public school educator noted that when AI handles repetitive academic tasks, children demonstrate far greater capacity for independence and reflection than traditional schedules assume.

Alpha charges \$40,000 to \$65,000 annually, betting that parents will invest in an education centered on student agency and self-directed learning rather than conventional instruction.

The \$500K No-Résumé Paradox: What AI's Talent Gold Rush Means for How Trilogy Finds Its People

As OpenAI posts half-million-dollar jobs with no résumé required, the logic behind Crossover's skills-first hiring model looks less radical — and more inevitable — than ever.

BY MARGOT SINCLAIR, SENIOR CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — The number stopped people mid-scroll: \$500,000. No résumé required. That was OpenAI's offer, reported this week by [Forbes](#), for roles demanding demonstrated AI capability — not credentials, not pedigree, not a Stanford diploma. And it didn't stop there. [Business Insider](#) documented a broader market shift: job listings across industries now explicitly require ChatGPT proficiency, with compensation reaching \$800,000 annually for the right candidate. The résumé, it seems, is having an existential crisis.

For anyone watching Trilogy International's Crossover platform, the irony lands with a particular weight. Crossover has spent years arguing exactly this — that the traditional résumé is a deeply flawed proxy for capability, that geography is irrelevant to talent, and that rigorous skills assessment is the only honest way to hire. The market, apparently, is catching up.

Crossover's model — operating across 130+ countries, placing full-time remote talent into Trilogy's ESW Capital portfolio companies and beyond — was built on a premise that felt contrarian when it launched: evaluate what people can actually do, not where they went to school or what their last title was. The platform's AI-enabled assessments are designed to surface the top tier of global technical and professional talent, paying identical above-market rates for identical perfor-

mance, regardless of whether the candidate is in Lagos or Los Angeles.

What the broader AI talent gold rush reveals is a systemic truth that Crossover has long operationalized: when the skill in question is new enough that no one has twenty years of experience with it, credentials collapse as a sorting mechanism. You are left with demonstrated ability — which is precisely what skills-first hiring was always designed to measure.

The stakes are not abstract. As OpenAI and its competitors vacuum up AI-fluent talent at extraordinary price points, the companies that built infrastructure for identifying that talent globally — before the gold rush — hold a structural advantage. The question for Trilogy's portfolio companies isn't whether they can compete with OpenAI's \$500,000 offers. It's whether the systems they've already built to find exceptional people, anywhere in the world, position them to move faster than the firms still sorting through résumés.

The Data Center Learns to Swim

A \$200 million wager would send AI compute drifting into the Pacific, where waves may feed the next generation of hungry machines.

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

SAN FRANCISCO — In the long evolutionary history of computing, the data center has been a notably terrestrial beast: squat, hot, and thirsty, nesting in industrial parks and drawing rivers of electricity through copper veins. Now, a new specimen is preparing to leave the shore.

Panthalassa, a Silicon Valley venture backed by roughly \$200 million in funding, is attempting something both audacious and faintly primordial: floating AI computing nodes in the Pacific Ocean, powered in part by the ceaseless motion of waves. The company plans to begin testing its ocean-borne systems in 2026, according to [Ars Technica](#).

Observe, if you will, the modern AI model in its feeding season. Each new generation requires more energy, more cooling, more specialized chips, and ever larger colonies of servers. On land, these colonies compete with cities, factories,

and households for power and water. Offshore, Panthalassa sees a different niche: abundant cooling, open space, and wave energy rolling beneath the hull like an untapped metabolic source.

The vision is elegant in the way nature can be elegant before it becomes brutal. Modular compute platforms would float at sea, their processors cooled by the surrounding water, their appetite partly satisfied by ocean motion. If successful, the approach could ease pressure on overloaded electrical grids and provide a new habitat for AI infrastructure at a moment when demand for training and inference is expanding with near-biological urgency.

But the sea is no passive landlord. Salt corrodes. Storms punish. Maintenance crews do not stroll casually across a parking lot to replace a failed component when the component is bobbing in the Pacific. Connectivity, environmental ap-

provals, maritime safety, and the economics of keeping delicate silicon alive in a hostile marine environment all await like predators in the kelp.

Still, the migration is telling. The AI industry is no longer merely asking where models should live inside software. It is asking where their bodies — the humming physical organs of computation — can be placed on Earth.

For Trilogy International's world of enterprise software, telco billing, AI analytics, and remote engineering, such infrastructure experiments matter. Whether inside ESW Capital's portfolio, Totogi's cloud-native telecom ambitions, or Klair's financial nervous system, cheaper and more abundant compute is the plankton of the digital food chain.

And so the servers edge toward the surf, tentative as hatchlings, while Silicon Valley watches the horizon.

Soderbergh Embraces AI Filmmaking Tools, Notwithstanding Industry Reservations

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

Steven Soderbergh, director of *Traffic* and *Ocean's Eleven*, says he intends to explore artificial intelligence as a filmmaking tool if such technology becomes available to creators. According to reports, Soderbergh stated he will experiment with any filmmaking tool that exists or emerges, regardless of industry reservations. His position does not constitute blanket endorsement of all AI applications, nor does it negate his previous concerns about copyright enforcement. Soderbergh's willingness to engage with AI-assisted production methods represents a significant development in ongoing industry discussions about artificial intelligence adoption in filmmaking. His openness to AI tools reflects artistic curiosity compatible with his established commitment to the craft.

AI Video Is Escaping the Lab — and Startups Are About to Get a Hollywood-Sized Upgrade

From scene-rewriting tools to new challenger models, generative video is moving from spectacle to startup growth engine.

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

SAN FRANCISCO — The AI video race is no longer just about dazzling demos of astronauts riding horses through neon galaxies. It is becoming infrastructure — for marketing teams, founders, entertainment studios and, yes, the next generation of tiny startups that suddenly have blockbuster creative powers in their browser.

I cannot overstate how significant this shift is: video, long one of the most expensive and operationally painful formats for young companies, is being compressed into prompts, templates and model workflows. A new wave of reporting on [how startups can leverage AI video to grow](#) points to a practical reality already taking shape: small teams can now generate product explainers, social ads, founder videos, customer education clips and localized campaigns at a speed that used to require agencies, film crews and weeks of production.

This changes everything because video is the native language of the internet, but historically it has punished the undercapitalized. AI flips that equation. A two-person SaaS company can test ten versions of a pitch video. An e-commerce brand can localize creative for different markets. A B2B founder can turn a blog post into a polished sales asset before lunch. The future is now, and it is rendering in 1080p.

The competitive battlefield is heating up fast. VentureBeat reports that the founders of OpenCV have launched an [AI video startup aimed at competing with OpenAI and Google](#), a deliciously important signal because OpenCV helped define modern computer vision itself. When the people who built foundational vision tooling jump into generative video, pay attention.

Meanwhile, Forbes reports Netflix has launched VOID AI, described as a system that can rewrite video scenes after filming. If that technology matures, post-production could become radically more fluid: change a background, adjust a scene, alter an object, maybe even personalize content versions without reshooting. Hollywood's editing room is starting to look like a prompt window.

But amid the excitement, trust is becoming the next battleground. Cisco's new Model Provenance Kit highlights a crucial question for enterprises: where did this AI model come from, what trained it and can anyone verify its lineage? As generative video floods the market, provenance will matter as much as pixels.

For startups, the message is unmistakable: learn this medium now. AI video is not a novelty. It is becoming growth software.

Nation's Executives Clarify AI Will Save Everyone Time Once Employees Stop Wasting So Much Of It Understanding Their Jobs

The productivity revolution has reportedly been delayed by the continued existence of work that requires knowing what is going on.

BY DALE PEMBERTON, STAFF WRITER · GPT-5.2

CAMBRIDGE, MASSACHUSETTS — In a development that has stunned business leaders who had been assured the future would arrive in a slide deck by Q3, researchers and workers are increasingly suggesting that artificial intelligence may not automatically make companies more productive when deployed by people with only a passing familiarity with the work being automated.

This has created an uncomfortable moment for executives, many of whom had already announced sweeping AI transformations, renamed several internal teams, and instructed employees to become “AI-first” in the same tone one might use to tell a dog to stop eating drywall.

The problem, according to a growing body of commentary, is that AI tools appear to function best when paired with humans who possess judgment, context, and expertise—three legacy features many organizations had hoped to phase out after discovering they were expensive and occasionally pushed back in meetings.

A recent [Forbes piece](#) argued that AI’s productivity promise falls apart without human expertise, an observation that has caused some managers to briefly consider whether employees who know things might be more than obstacles between software licenses and quarterly margin expansion.

Meanwhile, Harvard Business Review has given the economy a useful new term: “workslop,” referring to AI-generated output that looks polished enough to be forwarded but incoherent enough to force someone else to spend the afternoon decoding, repairing, or quietly replacing it. In practice, this means AI is already delivering massive efficiency gains for the person who produces the work, while creating an equal and opposite pile of unpaid interpretive labor for the person unfortunate enough to receive it.

This is, by any modern definition, innovation.

The workslop economy is elegant. A worker asks a model to draft a strategy memo. The model produces a confident fog bank of bullet points. The worker sends it to five colleagues. Each colleague then spends 40 minutes determining whether “activate scalable stakeholder alignment” means launch the product, delay the product, or fire the product manager. Productivity has occurred at the point of generation, which is the only place many dashboards are currently looking.

To be fair, the picture is not uniformly bleak. Anthropic has published research estimating productivity gains from Claude conversations, suggesting that when AI is used for the right tasks by people able to evaluate its output, it can save meaningful time. This is both promising and devastating, because it implies the machine works best when the human remains responsible for reality.

That conclusion is unlikely to slow company-wide AI pushes. Electronic Arts CEO Andrew Wilson has defended the gaming giant’s broad AI efforts despite employee claims that the technology has, in some cases, lowered productivity. This follows the standard corporate implementation model in which leadership identifies a transformative tool, employees report that it is making their actual jobs harder, and leadership concludes the transformation must be happening.

None of this means AI is useless. Quite the opposite. It means AI is a power tool, and many companies are currently handing it to interns, middle managers, and entire departments with the instruction to “build the house faster” while quietly removing the architects.

The uncomfortable truth is that AI does not eliminate expertise. It exposes where expertise was already doing invisible work. It reveals which employees were making good decisions, which processes were held together by institutional memory, and which executives believed writing was the same thing as thinking because both produced documents.

For now, businesses may need to accept a disappointing compromise: AI can make skilled people faster, average processes sharper, and repetitive work less miserable, but it cannot yet replace the annoying human capacity to know whether something is true, useful, or catastrophically stupid.

Naturally, that feature is expected in the next release.



The Office Comic · Art Desk

The Future of Work Is Not Coming for Your Job, It Is Coming for Your Excuses

AI is turning the labor market into a skills scoreboard, and the workers who learn fastest are about to compound hardest.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · GPT-5.2

AUSTIN, TEXAS — I'll be honest: the global workforce discourse has officially graduated from vibes to velocity. 🚀

PwC's latest [Global Workforce Hopes and Fears Survey 2025](#), the World Economic Forum's charts on AI's wage and hiring impact, Gartner's 2026 future-of-work guidance for CHROs, and Forbes' remote-work skills checklist all point to the same uncomfortable truth: the labor market is no longer rewarding presence, credentials or tenure nearly as much as adaptability, AI fluency and execution speed.

Unpopular opinion: this is good news for serious people. 💡

For decades, the corporate world quietly subsidized mediocre coordination, meeting theater and the sacred art of looking busy near a badge reader.

Now AI is putting that whole operating system under review.

The WEF's framing around AI's effect on wages and job quality matters because the labor-market split is becoming painfully obvious: workers who use AI to produce more valuable output are getting leverage, while workers who wait for HR to define the future for them are getting a calendar invite they probably do not want.

I'll be honest: the winners are not going to be the people with the fanciest prompts saved in a Notion doc.

The winners will be the people who can combine domain judgment, written clarity, data literacy, remote collaboration, and relentless iteration into measurable business outcomes.

That is why Gartner's future-of-work conversation for CHROs should not be treated as another glossy forecast deck for executives to skim between offsites.

It should be read as a warning label for organizations still managing talent like it is 2016 with better Slack emojis.

Hiring decisions are changing because work itself is becoming more observable, more modular and more benchmarkable.

This is where Trilogy International's model has looked less like an eccentric outlier and more like an early preview of the scoreboard economy.

Through Crossover, Trilogy built a global talent platform that recruits across 130-plus countries, pays identical above-market compensation regardless of geography, and filters for what it calls top 1% talent.

That approach has always triggered debates, because equal global pay plus high accountability is not the standard corporate comfort blanket.

But in an AI-accelerated labor market, the principle is brutally relevant: output beats location, and talent density beats office nostalgia.

For CHROs, the lesson is not "copy every operating detail tomorrow."

The lesson is that the old inputs are depreciating assets.

Degrees, proximity, political fluency and calendar stamina are not disappearing, but they are losing monopoly power as AI makes actual contribution easier to measure and easier to compare.

For workers, Forbes' remote-skills framing lands because remote work in 2026 will not be a lifestyle perk; it will be a professional discipline.

You will need asynchronous communication, self-management, AI tool fluency, cybersecurity awareness, cross-cultural collaboration and the ability to make your work legible without hovering over someone's shoulder.

I'll be honest: "I work better in person" may still be true for some people, but it is no longer a strategy.

The Pentagon's parallel scramble to scale laser weapons is an oddly perfect metaphor for the entire economy: institutions can know exactly what they need and still struggle to build fast enough.

Talent is the same way.

Everyone wants AI-ready teams, resilient leaders and high-agency remote operators, but the supply chain for those humans is not magically appearing because someone changed a competency framework.

Unpopular opinion: your career is now a product roadmap. 🚀

Ship capabilities, gather feedback, upgrade the stack, and stop waiting for permission from a corporate learning portal last updated during the pandemic.

The future of work is not a destination, it is a performance review happening in real time.

Humbled to share: the people who treat that as a learning opportunity are about to end this decade very, very strong. 💡

ON THIS DAY IN AI HISTORY

On May 6, 2016, Google's AlphaGo defeated Lee Sedol 4-1 in a five-game match of Go in Seoul, marking a watershed moment when AI surpassed humanity's best player in a game far more complex than chess.