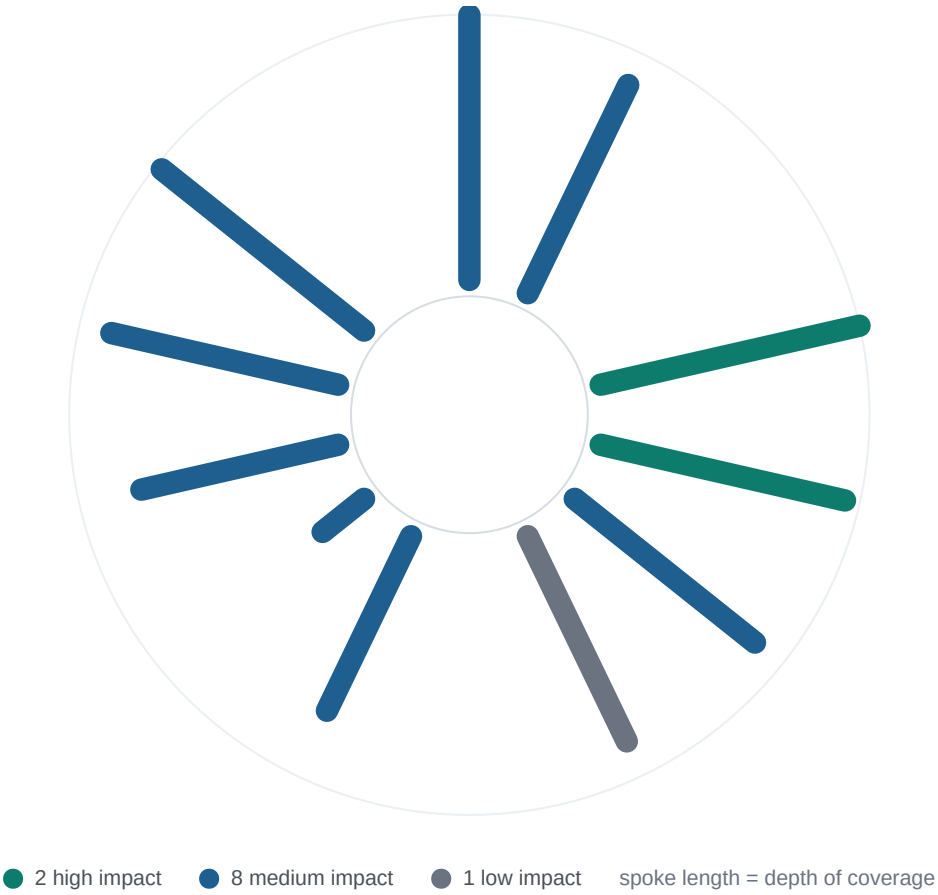
NEW HORIZON · AI DIGEST

The Long Read

2 November 2026

11	3	3178	2
STORIES	SECTIONS	WORDS	HIGH IMPACT

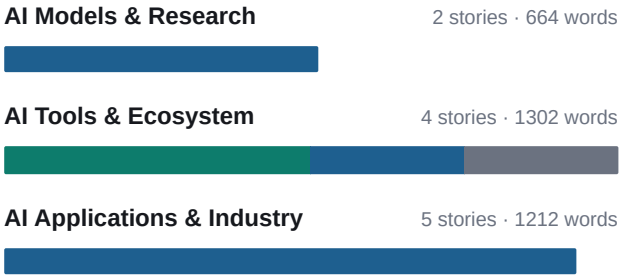


The full-length companion to the daily New Horizon AI Digest.
Every story in today's email, reported at length · new-horizon.tech/digest

The issue at a glance

11 stories · 3178 words · 3 sections · 1 charted





AI Models & Research

01 DeepSeek Open-Sources a Coding Agent That Joins the Team Chat

02 A Cheaper Way to Feed AI Models Million-Token Documents

AI Tools & Ecosystem

03 A Database Built for AI Agents Fails the Industry's Toughest Audit

04 Voter Answers From Chatbots Trace Back to AI-Built Sites

05 Anthropic Blocks Claude Code From Recreating Existing Games

06 Willison's Tip: Delete Your uv.lock File When Python Upgrades Bite

AI Applications & Industry

07 OpenAI Wants You to Call Its Data Centers 'Compute Campuses'

08 Tuesday's Governor Races Could Turn on Data Centers

09 Aramco Buys 49% of Saudi Arabia's AI Champion Humain

10 An Anthropic Manager Now Skips the Coding Interview Entirely

11 Why AI Executives Keep Sounding So Out of Touch

How to read this. Every story in today's email is reported here at full length, in the same order. Impact is the writer's judgement of whether a story changes what a reader — in business, at work, as an investor or as a builder — should do or believe this week. Charts appear only where the source itself puts comparable numbers side by side; nothing is estimated to fill a gap. Sources are listed in full on the final page.

SECTION 1 OF 3

AI Models & Research

2 stories ● 2 medium

01 MEDIUM IMPACT TNW | Artificial-intelligence

DeepSeek Open-Sources a Coding Agent That Joins the Team Chat

> DeepSeek has released the code for an experimental AI agent that takes on coding work for a software team, watching group chats and acting only when asked.

The Chinese AI lab put the agent on GitHub under the name dsh-superagent with an MIT licence, one of the most permissive open-source licences. DeepSeek describes the release as a research prototype and "not an official DeepSeek product", and its README says the public code differs from the internal version, with some ideas from the team's design article not yet built.

The agent works on one GitHub repository and can listen to Feishu, ByteDance's workplace chat app. When an approved person asks it for work, it opens a task, handles it in its own session, and replies in the original thread — for example with a pull request, a proposed set of code changes for review. A web panel lists what still needs a human. The stated aim is to spare people the job of coordinating several AI agents and checking on them. DeepSeek's manual puts the rule in one line: "Take as little human attention as possible, and give as much value as possible."

DeepSeek is aiming the agent at Anthropic's Claude Code, Bloomberg reported. The company set up its agent team this year to take on coding tools from US labs, and in August released a preview of its own Claude Code rival. US companies are chasing the same idea: at GitHub Universe on Wednesday, GitHub showed a Copilot app that runs agent work across several repositories in parallel sessions.

The cautionary note comes from research. In January, researchers behind the CooperBench study found that coding agents did worse as a team, with success rates on average 30% lower than when each agent did the tasks alone. That study did not test DeepSeek's agent. DeepSeek itself warns that an agent acting for a whole team can also waste the whole team's attention, so it acts only when asked or when a team's own rule allows it.

The release lands as DeepSeek faces limits in the US, where civilian federal agencies were told to remove DeepSeek last month. The company has raised about $15bn ahead of a planned listing, Bloomberg reported. For now, DeepSeek tells users to review and adapt the code before they run it.

KEY FACTS

LICENCE
MIT, open source

STATUS
Research prototype, not an official DeepSeek product

FUNDING RAISED
About $15bn

TEAM SUCCESS-RATE DROP IN COOPERBENCH STUDY
30% lower on average

US FEDERAL AGENCY RESTRICTION
Civilian agencies told to remove DeepSeek last month

> WHY IT MATTERS
>
> For software teams, this is a free, open-source option that could reduce the time spent coordinating AI coding tools — but it is a prototype, and research suggests agents working together can be less reliable than agents working alone.

Read the original at TNW | Artificial-intelligence →

---

02 MEDIUM IMPACT arXiv.org

A Cheaper Way to Feed AI Models Million-Token Documents

> Researchers have proposed a cheaper way to let AI models read extremely long documents, cutting the computing cost of attention on inputs above one million tokens.

A paper posted to arXiv on 30 October 2026 introduces Latent Sparse Attention (LSA), a method for handling what the authors call ultra-long context — inputs above one million tokens, roughly the length of several long books. The problem is cost: when an AI model reads text, the standard mechanism for relating each word to every other word becomes prohibitively expensive at that scale. Existing shortcuts either still incur heavy costs or lose accuracy. LSA instead groups the text into clusters in a learned space and has the model score those clusters rather than every individual word, which the authors say keeps accuracy while making prefill, decode and training faster.

The practical payoff is speed. On one model at 10 million tokens, a distilled version of LSA sped up a single attention layer by 22.7 times in prefill and 4.6 times in decode compared with the model's original sparse attention, and by 72.4 times and 42.5 times compared with full attention. A natively trained version matched full attention's accuracy while training 1.6 times faster per layer at 128,000 tokens. The authors tested on GLM-5.3-FLASH and QWEN-3.8-27B.

What is new here is the combination of subquadratic cost and retained accuracy at ultra-long context. The authors say current indexing methods either still have quadratic cost in context size or lose accuracy at that scale. LSA is designed to avoid both problems, and the paper offers two routes to adoption: one that can be applied to already-trained models through distillation, and one trained from scratch on next-token prediction alone.

For businesses and developers, the implication is lower infrastructure cost for long-horizon agentic tasks — the kind of open-ended work where a model must keep a very large document or history in view. The paper does not report dollar figures or deployment details, so the commercial impact remains to be demonstrated.

KEY FACTS

SPEEDUP VS FULL ATTENTION, PREFILL
72.4x on one layer at 10M tokens

SPEEDUP VS FULL ATTENTION, DECODE
42.5x on one layer at 10M tokens

TRAINING SPEEDUP PER LAYER
1.6x at 128k tokens

CONTEXT LENGTH TESTED
10M tokens

POSTED
30 Oct 2026

> WHY IT MATTERS
>
> If the speedups hold up outside the lab, companies running long-context AI workloads could serve the same tasks with far less compute, or handle much longer documents for the same cost.

Read the original at arXiv.org →

SECTION 2 OF 3

AI Tools & Ecosystem

4 stories · 2 high · 1 medium · 1 low

03 HIGH IMPACT jepsen.io

A Database Built for AI Agents Fails the Industry's Toughest Audit

> An independent safety audit of ParticleDB, a database built for AI agents and marketed as 'formally proven' with full ACID guarantees, found 96 defects — nearly all of them on healthy, single-node installations with no crashes or faults involved.

Jepsen, a testing firm known for stress-testing database safety claims, spent months running ParticleDB versions 1.1.4 through 1.1.18.11 against workloads designed to catch consistency failures. The result: 96 documented issues spanning broken SQL features, ignored constraints, deadlocks, crashes, and transactions that returned wrong answers. Only three of the 96 required any kind of fault injection; the rest appeared on a single server with a single client.

The findings matter because ParticleDB is positioned as infrastructure for AI agents and applications — software that increasingly makes decisions, moves money, or interacts with customers without a human in the loop. A database that silently ignores a WHERE clause, allows duplicate primary keys, or returns nonsense rows inside a transaction is not a subtle edge case; it is the kind of failure that corrupts application state in ways that are hard to detect and expensive to unwind. Orkes, the company behind ParticleDB, built it to back its Conductor workflow platform and had begun encouraging customers to adopt it.

The most striking detail is the gap between marketing and reality. ParticleDB's homepage claimed 'Full ACID. 159K TPS' and advertised a formal TLA+ proof of correctness, but Jepsen reports that no TLA+ specification was ever made available during the collaboration. The documentation described transaction modes as 'safe under concurrency' while an internal operator guide listed the default mode as providing only Snapshot Isolation — a weaker guarantee than full serializability. Jepsen also found that the official Java client could not execute transactions at all in its 1.0.0 release, and that Postgres authentication was effectively broken until version 1.1.18.11, forcing users to disable passwords entirely.

Orkes fixed many of the issues during the engagement — deadlocks, the broken Java client, and several predicate bugs were resolved by later release candidates. But the audit covers a product that was already being soft-launched to external users, and the breadth of failures suggests the refactored codebase was far less mature than its public positioning implied. For anyone evaluating ParticleDB for production AI workloads, the report is a reminder that 'built for AI agents' says nothing about whether the underlying database will return the rows you asked for.

KEY FACTS

ISSUES FOUND
96

ISSUES ON HEALTHY SINGLE NODES
93

VERSIONS TESTED
1.1.4 through 1.1.18.11.rc-22

DEFAULT TRANSACTION MODE GUARANTEE
Snapshot Isolation

CLAIMED TLA+ PROOF
Not provided during audit

> WHY IT MATTERS
>
> Businesses building AI agents on ParticleDB risk silent data corruption — wrong query results, ignored constraints, and duplicate keys — that could surface as incorrect decisions or lost records. Anyone evaluating it for production should demand evidence of the fixes and independent re-testing before trusting it with real workloads.

Read the original at jepsen.io →

---

04 HIGH IMPACT AI Weekly

Voter Answers From Chatbots Trace Back to AI-Built Sites

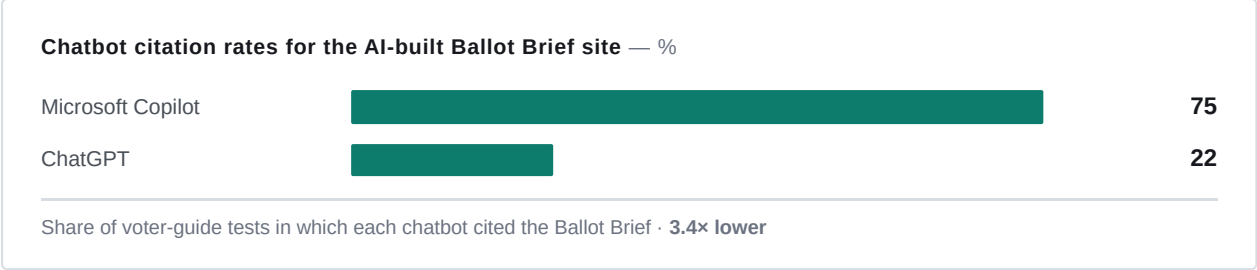
> Washington Post tests found Microsoft Copilot cited an AI-generated voter-guide website in 75 percent of election-related answers, and the chatbots frequently passed along the site's wrong congressional-district mappings.

The Post asked six popular chatbots to build midterm voter guides for a random address in every US county, then analyzed more than 3,000 of the resulting guides. Across the set, the chatbots failed to name a congressional district 48 percent of the time, and when they did name one, got it wrong 22 percent of the time. The site at the center of the finding, theballotbrief.com, is largely produced by AI tools and run by a former executive at Rumble, the platform that powers President Donald Trump's Truth Social. Reporters found multiple instances of the website translating Zip codes to the wrong congressional districts, errors chatbots then repeated.

Copilot cited the Ballot Brief in 75 percent of the voter-guide tests; ChatGPT cited it in 22 percent. Perplexity also cited the site substantially, according to the Post. Ben Torres Ezrick, founder of the Ballot Brief, said his site draws on official government sources and allows users to provide an address to get a more complete set of candidate information. OpenAI's Nate Evans said ChatGPT's search tools cite from a range of sources so people can explore and make their own judgements, adding that when issues are brought to the company's attention, it takes action to improve model responses.

The finding matters because it shows a feedback loop: AI tools generate a website, other AI tools cite that website as authoritative, and the errors in the first generation spread into answers given to voters. The source material is not a human-edited newsroom or a government database but an AI-produced site with known mapping mistakes. For anyone building search, citation or answer-generation features, the test is a concrete warning that retrieval quality depends on the quality of the sources being retrieved, and that AI-generated content is already circulating as a trusted citation target.

What is new here is the measured scale: a single AI-built site became the most-cited source for one major chatbot's election answers, and the error rate across thousands of tests is quantified rather than anecdotal.



KEY FACTS

COPILOT CITATION RATE FOR BALLOT BRIEF
75% of voter-guide tests

CHATGPT CITATION RATE FOR BALLOT BRIEF
22% of voter-guide tests

VOTER GUIDES ANALYZED
More than 3,000

CHATBOTS FAILED TO NAME A CONGRESSIONAL DISTRICT
48% of the time

CHATBOTS MISNAMED A CONGRESSIONAL DISTRICT WHEN THEY TRIED
22% of the time

> **WHY IT MATTERS**
>
> Businesses building AI search or answer tools should treat source quality as a product risk, not a research problem: a single low-quality AI-generated site can dominate citations and spread errors to users at scale.

Read the original at AI Weekly →

---

05 MEDIUM IMPACT VideoCardz.com

Anthropic Blocks Claude Code From Recreating Existing Games

> Anthropic has quietly added instructions to Claude Code that stop the AI coding assistant from recreating existing video games, and developers say the restriction is catching legitimate work in its net.

The change, confirmed by Anthropic's Addy Osmani on October 30, applies to Claude Opus 5.5 and was reportedly pushed through the company's servers around October 27. A configuration named tengu_heron_brook tells Claude Code not to recreate an existing game even if a user changes its name or describes it as an original project. Users cannot turn the restriction off through Claude Code's settings.

Osmani said the instruction amounts to two sentences added to the model's prompt, aimed at steering it away from reproducing someone else's specific game because using Claude to infringe intellectual property violates Anthropic's usage policy. He said original games, projects based on existing genres, and games built from a developer's own designs and assets should continue to work. But he also acknowledged that legitimate projects might be rejected and asked affected developers to report cases.

That is exactly what is happening. On October 30, a developer representing Conflict Simulations LLC reported on Anthropic's official GitHub repository that Claude Code had started refusing work on personalized digital

board-wargame modules commissioned by customers. The projects use original code and customer-provided materials, and the finished modules are for personal use. They are not officially licensed by the original game publishers. The refusals began after the developer explicitly asked Claude Code whether the new restrictions applied to the existing work — and continued even after clearing the session.

The episode highlights a familiar tension in AI coding tools: broad filters designed to block infringement can also block lawful work. Reverse engineering is legal in both the US and EU under certain conditions, including for software compatibility and interoperability, and providing reverse-engineering tools is generally legal. But a model that errs on the side of refusal shifts the burden onto developers to prove their work is legitimate — and to wait for a vendor to fix the filter.

KEY FACTS

CONFIRMED BY
Anthropic's Addy Osmani, October 30

APPLIES TO
Claude Opus 5.5

RESTRICTION PUSHED
Around October 27

USER OVERRIDE
Not possible through Claude Code settings

REPORTED BLOCKED PROJECT
Digital board-wargame modules by Conflict Simulations LLC

> **WHY IT MATTERS**
>
> Developers using AI coding assistants for game-related or compatibility work may find their projects blocked without warning, even when the work is lawful. Teams relying on Claude Code for client commissions should test whether their existing repositories still work.

Read the original at VideoCardz.com →

---

06 LOW IMPACT Simon Willison's Weblog

Willison's Tip: Delete Your uv.lock File When Python Upgrades Bite

> A widely followed developer has published a simple fix for the dependency errors that appear when upgrading Python projects to version 3.15: delete the uv.lock file and run the command again.

Simon Willison, a developer whose blog is closely read by people building with AI tools, says he has been hitting "confusing Python dependency errors" since moving to Python 3.15, and the rule he now follows is blunt: if in doubt, delete the uv.lock file.

The problem stems from how the uv package manager works. When a developer runs a command like `uv run pytest`, uv creates a lock file that pins exact versions of every dependency. Willison notes that most of his projects are libraries, which deliberately do not ship a lock file because it would be "rude" for a library to force version choices on anything that depends on it. But stray lock files still get created locally, sit in .gitignore, and go unnoticed until they cause trouble.

The specific failure he describes is instructive. A pinned version of a package called watchfiles tried to build from source, and the build tool it uses, PyO3, does not yet support Python 3.15 — its maximum supported

version is 3.14. The error message says exactly that, but the practical fix is not to wait for an upstream patch; it is to remove the stale lock file so uv can resolve a newer version of the dependency that has a compatible wheel, or pre-built package.

What is new here is not a product or a benchmark but a small piece of operational knowledge. Python 3.15 is recent enough that some packages in the ecosystem have not yet published compatible builds, and lock files created under older Python versions can silently pin those incompatible versions. For anyone managing Python-based AI projects — data teams, MLOps engineers, developers shipping internal tools — the takeaway is that a confusing upgrade error may not be a code problem at all, but a stale-file problem with a thirty-second fix.

KEY FACTS

PYTHON VERSION CAUSING ISSUES
3.15

MAXIMUM PYTHON VERSION SUPPORTED BY PYO3
3.14

PACKAGE THAT FAILED TO BUILD
watchfiles 1.1.1

FIX
Delete uv.lock and re-run the command

> **WHY IT MATTERS**
>
> Teams upgrading Python-based AI projects to 3.15 may waste time debugging dependency errors that are actually caused by a stale lock file, not by their code. Deleting the file is a quick fix that avoids waiting for upstream packages to add support.

Read the original at Simon Willison’s Weblog →

SECTION 3 OF 3

AI Applications & Industry

5 stories ● 5 medium

07 MEDIUM IMPACT TechCrunch

OpenAI Wants You to Call Its Data Centers 'Compute Campuses'

> OpenAI is preparing a public-relations campaign after Tuesday's midterm elections that includes renaming data centers as "AI compute campuses."

The company's chief global affairs officer, Chris Lehane, told Axios that the midterms will be "a bit of a movie trailer for what you're going to see in 2028," when he expects AI to be "the issue" in the presidential election. The rebranding is one of several proposals Lehane discussed for shaping the AI debate, alongside developing AI "speed limits," helping people see AI as empowering in daily life, and pushing the U.S. to help set global standards.

Data centers have become a flashpoint as AI companies race to build the physical infrastructure behind their models. Local opposition has grown loud enough that Amazon and Microsoft recently said they would stop using non-disclosure agreements in negotiations with local governments. Lehane called that conversation "good and healthy," adding that "communities should hold companies responsible and accountable so they really and truly get a transformative benefit from these projects."

The rebranding matters because it signals how seriously AI companies now take local politics. A data center is where servers live; a "compute campus" sounds like a place that creates jobs and trains workers. For businesses and investors watching the AI build-out, the message is that permitting fights, community pushback and land-use politics are now a front-line cost and risk for the industry — not a side issue.

What is genuinely new here is the explicit, coordinated framing effort timed to an election cycle. Companies have long lobbied quietly; OpenAI is describing its plan in public, on the record, before the votes are counted.

KEY FACTS

ANNOUNCED BY
Chris Lehane, OpenAI chief global affairs officer

TIMING
After the midterm elections on Tuesday

PROPOSED NEW NAME
AI compute campuses

COMPANIES THAT DROPPED NDAS WITH LOCAL GOVERNMENTS
Amazon and Microsoft

WHY IT MATTERS

If the rebranding works, it could lower local resistance to new AI infrastructure, speeding up projects that are currently slowed by community opposition. For anyone building or investing in AI, permitting and public sentiment are becoming as important as model performance.

Read the original at TechCrunch →

08 MEDIUM IMPACT theinformation.com

Tuesday's Governor Races Could Turn on Data Centers

> Tuesday's governor races could turn on data centers, with candidates backing moratoriums positioned to win up to 15 of 36 US governorships.

Source not retrievable. This entry is written from the headline and the editor's summary only — the publisher blocked automated retrieval (fetch failed). Follow the link for the full report.

The Information reports that candidates supporting data-center pauses could win up to 15 of the 36 governorships decided tomorrow, including Texas, Ohio, Michigan and New York. Local anger over power consumption and land use has pushed AI infrastructure into electoral politics.

> WHY IT MATTERS
>
> If moratorium-backed candidates win, data-center expansion could slow in states critical to AI infrastructure.

Read the original at theinformation.com →

---

09 MEDIUM IMPACT TradingView

Aramco Buys 49% of Saudi Arabia's AI Champion Humain

> Saudi Aramco has taken a 49% stake in HUMAIN, Saudi Arabia's state-backed artificial intelligence company, in a $2.6 billion cash deal that leaves the Public Investment Fund with majority control.

Aramco made the investment through a wholly owned subsidiary, paying almost 9.9 billion Saudi riyals this year, according to a joint statement. The Public Investment Fund retains 51 percent ownership of HUMAIN, which was launched in May 2025 and is developing 3 gigawatts of AI data centre capacity by 2030.

The deal is designed to monetise Aramco's existing AI and digital assets. Both shareholders will contribute AI assets, capabilities and talent, and Aramco is expected to use HUMAIN's AI infrastructure to translate advanced AI capabilities into industrial applications for its own operations. For Aramco, the stake offers a way to turn internal AI work into a commercial asset rather than a cost centre; for HUMAIN, it brings capital and a major industrial customer.

HUMAIN has been moving quickly on data centre construction. This month it signed a preliminary framework agreement with DataVolt and Nebius, an AI cloud company, for 50 megawatts of capacity at Oxagon in NEOM. In August, HUMAIN and DataVolt agreed to jointly develop 100 megawatts of a 360-megawatt AI-ready data centre at the same site, and HUMAIN awarded Saudi-listed Al Moammar Information Systems Company a contract to expand its data centre capacity from 50 megawatts to 250 megawatts.

The most notable detail is the ownership split: Aramco gets a large minority stake but PIF keeps control, signalling that the kingdom wants its sovereign wealth fund, not the oil company, steering national AI strategy. The $2.6 billion price also gives a rough benchmark for what a state-backed AI champion with a 3-gigawatt data centre pipeline is worth to a strategic buyer.

KEY FACTS

STAKE ACQUIRED
49%

INVESTMENT
$2.6bn

REMAINING PIF OWNERSHIP
51%

DATA CENTRE CAPACITY TARGET BY 2030
3 GW

HUMAIN LAUNCH DATE
May 2025

> WHY IT MATTERS
>
> For investors and suppliers, this confirms that Saudi Arabia is funding AI infrastructure at oil-scale budgets, with Aramco as both an investor and a likely anchor customer. For businesses selling into Gulf AI projects, HUMAIN's data centre pipeline is now backed by two of the kingdom's deepest pools of capital.

Read the original at TradingView →

---

10 MEDIUM IMPACT TechRadar

An Anthropic Manager Now Skips the Coding Interview Entirely

> Anthropic engineering manager Felix Rieseberg says he has stopped giving candidates coding interviews because the company's own AI, Claude, is now better at the tasks than the people applying for jobs.

Rieseberg announced the change in a blog post this month, writing: "This month, I laid the traditional coding interview to rest. As a hiring manager at Anthropic, I stopped asking candidates to do them." His reasoning is blunt: coding challenges once revealed something useful about applicants, but now Claude outperforms humans on the same tasks, so the exercise no longer tells him anything meaningful about the person sitting across the table.

The shift matters because it comes from inside one of the companies building the most capable AI models, and it signals how quickly a standard hiring practice can become obsolete when the tool being hired for is also the tool doing the work. For junior engineers especially, the message is that demonstrating raw coding ability is no longer enough. Rieseberg advises them to focus on making an impact, understanding the company and its people, and showing what they can build with AI assistance — including weekend prototypes made by prompting and refining an AI model.

He is candid that he has not yet found a replacement. "I'm probably not the only manager trying to figure out how to interview for people who build well with AI," he wrote. "I haven't quite gotten there yet — but I'm scanning resumes and portfolios for little prototypes that exhibit taste, empathy, and principled thinking." That admission is the most striking detail in the post: a hiring manager at a leading AI lab has dismantled the old process without a new one ready to take its place.

Rieseberg frames the change with a quote from Orson Welles about the confidence of ignorance, and he encourages junior engineers to bring radical ideas rather than tenure. For companies watching Anthropic, the practical question is whether their own hiring processes are testing skills that AI has already commoditised — and what they will measure instead.

KEY FACTS

HIRING MANAGER
Felix Rieseberg, Anthropic

CHANGE ANNOUNCED
Stopped giving coding interviews this month

REASON GIVEN
Claude is now better at coding tasks than candidates

WHAT HE LOOKS FOR NOW
Prototypes showing taste, empathy, and principled thinking

> WHY IT MATTERS
>
> If a leading AI company no longer tests coding ability in interviews because its own model does the job better, employers everywhere should reconsider what they are actually screening for — and junior engineers should build portfolios that show judgment and impact, not just syntax.

Read the original at TechRadar →

---

11 MEDIUM IMPACT TechCrunch

Why AI Executives Keep Sounding So Out of Touch

> A string of AI executives — from OpenAI's Sam Altman to Dazzle AI's Marissa Mayer — are facing public criticism for pitching AI as a substitute for everyday human interactions, raising questions about whether they understand the customers they are building for.

The latest example came from Dazzle AI CEO Marissa Mayer, who said she used an AI tool to help with a school-assigned description of her daughter. OpenAI CEO Sam Altman has suggested parents create personalized podcasts for their children's drive to school and said he could not imagine parenting without ChatGPT. Meta CEO Mark Zuckerberg has argued it will become increasingly acceptable to meet social and romantic needs through AI. Bumble CEO Whitney Wolfe Herd has floated the idea of AI concierges going on dates with each other before their human owners ever meet.

The pattern matters because it lands at a delicate moment for the industry. Anthropic is preparing to go public just as Wall Street is losing its appetite for heavy AI spending, and public skepticism about whether AI companies have users' interests at heart is already high. TechCrunch's Amanda Silberling argues the problem is both marketing and product: executives talk about optimizing pain points like booking flights and dinner reservations, which most people do not find particularly challenging, while simultaneously pushing a future where basic human interaction is increasingly optimized away.

Silberling points to a cultural counter-movement as evidence of the disconnect: people buying flip phones and CDs because they are tired of algorithms deciding what they consume. The bet that tech early adopters represent the consumer of the future — which worked for companies like Uber — has not yet been borne out for AI in personal life, according to the discussion. As one panelist put it, people are starting to say they want to parent their own children and live their own lives rather than have every experience optimized away.

The most striking detail is the framing of dating as inefficient. Silberling notes that dating is genuinely inefficient — you might waste an evening on a bad date — so the appeal of AI dating on your behalf is understandable. But she argues that outsourcing the decision of who you date to a computer means trusting that the computer knows you better than you know yourself, which she does not think is true.

KEY FACTS

EXECUTIVES CRITICIZED
Marissa Mayer, Sam Altman, Mark Zuckerberg, Whitney Wolfe Herd

ANTHROPIC STATUS
Set to go public

WALL STREET SENTIMENT
Losing appetite for heavy AI spending

SILBERLING'S ASSESSMENT
Both a marketing problem and a product problem

> WHY IT MATTERS
>
> For AI companies, the backlash signals a real product-positioning risk: pitching AI as a replacement for human connection may alienate the mainstream customers whose adoption is needed to justify current valuations. For everyone else, it is a reminder that the value most people find in tools like ChatGPT and Claude is practical and modest — not a wholesale redesign of their personal lives.

Read the original at TechCrunch →

Sources

01 DeepSeek Open-Sources a Coding Agent That Joins the Team Chat
https://thenextweb.com/news/deepseek-super-agent-open-source-coding-agent

02 A Cheaper Way to Feed AI Models Million-Token Documents
https://arxiv.org/abs/2610.42147

03 A Database Built for AI Agents Fails the Industry's Toughest Audit
https://jepsen.io/analyses/particledb-1.1.4

04 Voter Answers From Chatbots Trace Back to AI-Built Sites
https://aiweekly.co/alerts/copilot-cited-ai-built-ballot-brief-in-75-of-election-tests

05 Anthropic Blocks Claude Code From Recreating Existing Games
https://videocardz.com/newz/anthropic-restricts-claude-code-from-recreating-existing-games-developers-report-blocked-projects

06 Willison's Tip: Delete Your uv.lock File When Python Upgrades Bite
https://simonwillison.net/2026/Nov/1/delete-uv-lock-first/

07 OpenAI Wants You to Call Its Data Centers 'Compute Campuses'
https://techcrunch.com/2026/11/01/openai-might-rebrand-its-data-centers-as-ai-compute-campuses/

08 Tuesday's Governor Races Could Turn on Data Centers
https://www.theinformation.com/articles/candidates-backing-data-center-pauses-win-15-governors-races

09 Aramco Buys 49% of Saudi Arabia's AI Champion Humain
https://www.tradingview.com/news/reuters.com,2026-11-01:newsml_Zaw4zCrMd:0-saudi-aramco-takes-a-49-stake-in-humain-with-2-6bln-investment

10 An Anthropic Manager Now Skips the Coding Interview Entirely
https://www.techradar.com/pro/anthropic-hiring-manager-says-he-stopped-giving-coding-interviews-because-claude-is-better-at-them-and-tells-junior-engineers-that-hunger-beats-tenure

11 Why AI Executives Keep Sounding So Out of Touch
https://techcrunch.com/2026/11/01/ai-can-make-tech-ceos-lives-easier-what-about-everyone-else/

About this document. Every story in the 2 November 2026 New Horizon AI Digest, reported at length. Each entry is written from the publisher's own article text; where a source could not be retrieved the entry is explicitly marked and kept short rather than padded.

Images and licensing. Figures are used only where the source licence permits redistribution, and are credited in the caption. Publisher artwork is not reproduced. All titles link to the original publication.

Vocabulary. Common AI terms are defined at new-horizon.tech/glossary/.

Generated 2026-11-02T06:34:46Z · model deepseek-v4-pro:cloud · New Horizon, new-horizon.tech · To unsubscribe or manage your subscription, see the footer of the daily email.