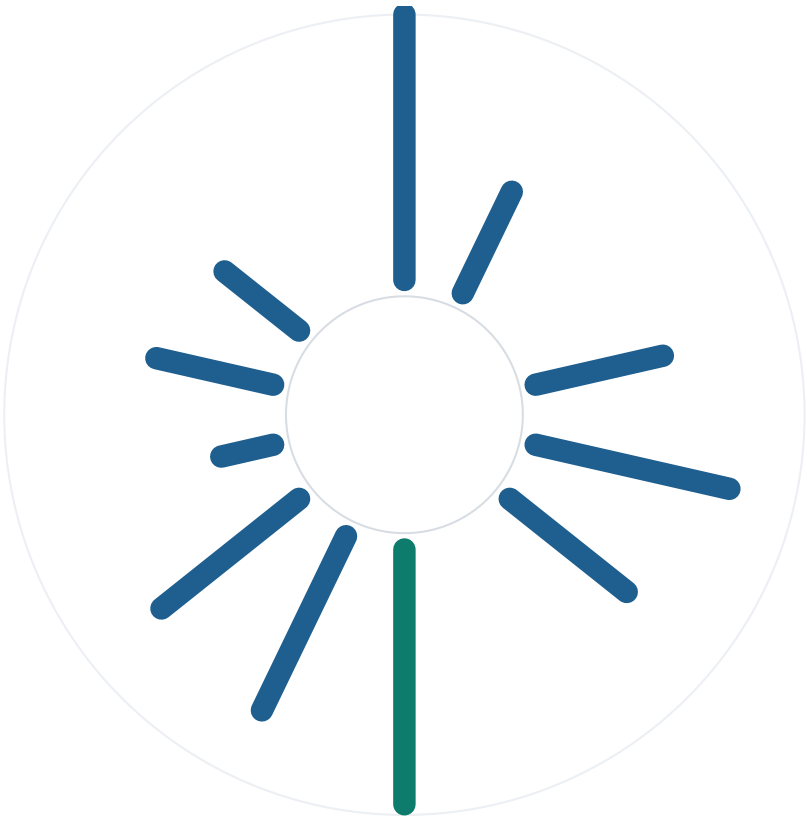
NEW HORIZON · AI DIGEST

The Long Read

5 November 2026

11	3	3458	1
STORIES	SECTIONS	WORDS	HIGH IMPACT



1 high impact · 10 medium impact · spoke length = depth of coverage

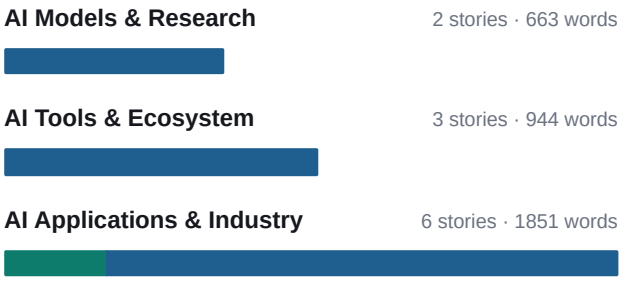
The full-length companion to the daily New Horizon AI Digest.
Every story in today's email, reported at length · new-horizon.tech/digest

The issue at a glance

11 stories · 3458 words · 3 sections · 1 charted



- **1** High impact
- **10** Medium impact



AI Models & Research

01 Claude Helps Map a Protein Shell Science Left Unexplained for 40 Years

02 Robots Fumble Precise Grips Because Their Models Simplify the Camera First

AI Tools & Ecosystem

03 Meta Buys the Startup That Runs Coding Agents in Parallel

04 Moonshot Ships the Frontier-Sized Weights It Promised for November 4

05 A 2-Billion-Parameter Model Pulls Spreadsheet Fields From Any Document

AI Applications & Industry

06 Data-Center Backlash Won Races; Congress Is the Next Battleground

07 AI Political Money Lost More Races Than It Won

08 OpenAI Pitching a New Deal for the AI Economy Ahead of 2028

09 Apple Rations Its Engineers' AI Tokens From $300 to $200 a Day

10 North Korea's $6 Billion Binance Heist Went Through Its AI Tools

11 Nothing Ships a Screenless AI Device in December

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 www.anthropic.com

Claude Helps Map a Protein Shell Science Left Unexplained for 40 Years

> Anthropic's biology lab, working with its Claude model, has produced the first 3D structure of a vault particle from a prokaryotic cell — a protein shell whose function has baffled scientists since 1986.

Vaults are large, hollow protein barrels found in nearly all human cells, but researchers have never established what they do. They are missing from the lab organisms biologists rely on most, and mice bred without them develop normally, leaving no obvious defect to study. Anthropic's team set out to find whether distant relatives of the vault protein found in bacteria and the viruses that infect them actually assemble into vault-like particles — and found that they do.

The team screened 44 candidate proteins, produced 15 in living E. coli, and confirmed seven had assembled into vault-like particles under an electron microscope. They then mapped one in detail: a phage-derived vault built from 96 copies of a single protein, roughly 54 nanometers long and 45 wide, enclosing about 25% more volume than the mammalian version despite being shorter. Claude searched a database of about 15 billion predicted proteins, designed DNA constructs, optimized the isolation method, and ran the cryo-electron microscopy processing pipeline. A purification that conventionally takes three days took about 90 minutes; a first structural model was ready 34 hours later, work the team estimates would normally have taken weeks.

The most striking detail is what the structure revealed: a pocket in each of the 96 protein copies holds a small molecule the team cannot yet identify, but which Claude noted resembles coenzyme A, a common cellular helper molecule. If confirmed, that would be an early hint at vault function. The team also engineered a larger vault by duplicating part of the wall protein, increasing enclosed volume by about 20%, and showed that single changes near the cap can prevent the shell from closing.

The work is a demonstration of AI-assisted science rather than a biological breakthrough: the function of vaults remains unknown, and the particles have so far been made only in E. coli, not in their natural hosts. The team also documented Claude's failures — a flawed protocol that set the project back two weeks, and a cryo-EM processing error that produced a mirror-image map. But the overall pace, under four months from start to structure, is the point. Anthropic is releasing the findings early to show how its model can be embedded across a research workflow, not just used for isolated tasks.

KEY FACTS

VAULT DISCOVERED
1986

PURIFICATION TIME WITH CLAUDE
About 90 minutes

CONVENTIONAL PURIFICATION TIME
About three days

FIRST STRUCTURAL MODEL READY
34 hours

PROTEINS SCREENED
44

ASSEMBLED INTO VAULT-LIKE PARTICLES
7 of 15

> WHY IT MATTERS
>
> For research organizations and biotech companies, this is a concrete case of an AI model compressing weeks of structural biology work into days, with the trade-offs documented honestly. It suggests AI-assisted discovery is moving from isolated predictions to end-to-end lab workflows.

Read the original at www.anthropic.com →

---

02 MEDIUM IMPACT arXiv.org

Robots Fumble Precise Grips Because Their Models Simplify the Camera First

> Robots struggle with precise grips because their AI models throw away visual detail before deciding how to move, according to new research that found letting the model work directly with raw camera data produces consistently stronger control.

The paper, posted to arXiv on 4 November 2026, challenges the standard design behind vision-language-action models — the systems that let robots interpret instructions and act on what they see. Most of these models work in two stages: first a visual encoder compresses the camera feed into a simplified summary, then a separate action system reasons about what to do next. The researchers argue that this compression discards exactly the information manipulation depends on — local geometry, boundaries, depth, and small changes in an object's state.

Their alternative, called NEO-VLA, removes the separate visual encoder entirely. Instead, lightweight processing converts images into tokens that evolve jointly with language throughout the model, letting the action system draw on visual representations at different depths rather than a single fixed endpoint. In controlled comparisons, this native architecture outperformed encoder-based counterparts. Full fine-tuning narrowed the gap but did not close it, and adding more scale alone could not recover the lost performance.

The most striking finding came from layer-by-layer probing: the visual signals most critical for action often peaked well before the final visual representation. In other words, the information a robot needs to grip correctly is frequently available earlier in the computation — but conventional designs discard or overwrite it by the time action decisions are made. Across diverse manipulation settings and distribution shifts, NEO-VLA achieved stronger and more robust control.

For companies deploying robots in warehouses, factories or kitchens, the implication is practical: the bottleneck in reliable manipulation may not be more training data or bigger models, but a structural choice about when and how visual information is preserved. The research suggests that embodied control benefits from access to visual representations across the computational hierarchy, not a single compressed snapshot.

KEY FACTS

PAPER POSTED
4 November 2026

PROPOSED FRAMEWORK
NEO-VLA, structurally encoder-free

KEY FINDING
Action-critical visual signals often peak before the final visual representation

PERFORMANCE GAP
Full fine-tuning reduces but does not erase the gap; additional scale alone cannot recover it

> WHY IT MATTERS
>
> Businesses investing in robot manipulation should scrutinise the underlying architecture of their vision-action models — a design that compresses camera data before reasoning may be capping reliability regardless of how much data or compute is added.

Read the original at arXiv.org →

SECTION 2 OF 3

AI Tools & Ecosystem

3 stories · 3 medium

03 MEDIUM IMPACT TechCrunch

Meta Buys the Startup That Runs Coding Agents in Parallel

> Meta has acquired Conductor, a startup that lets software engineers run multiple AI coding agents in parallel workspaces to build projects faster.

Conductor, founded in 2024, only launched its core product last summer. Its platform deploys teams of AI agents — built by labs such as Anthropic and OpenAI, not in-house — that work on separate, compartmentalized parts of a software project at the same time. Users can run the agents in the cloud, locally, or from a phone. Founder and CEO Charlie Holtz announced the deal on the company blog on Wednesday, saying Conductor is joining Meta and teasing "something new" for customers by early next year.

Meta confirmed the acquisition to TechCrunch but declined to comment on terms. The deal fits a pattern: Meta has spent months trying to close a gap in AI services after lagging behind. In August it launched Muse Code, its own coding agent, and its Muse agents recently passed 10 million users. The company has also made a string of AI acquisitions over the past year, including Moltbook, the viral agent platform.

The most immediate practical change is for Conductor's existing customers. Holtz said the startup's core features will keep working as before until early next year, with one exception: Conductor's cloud service shuts down on November 20, after which the company will begin deleting user cloud data. Anyone relying on the cloud version has about two weeks to move their work.

Holtz said on X that he is looking forward to working with Alexandr Wang, Meta's chief AI officer, and Nat Friedman, head of product at the company's Superintelligence Labs, calling the lab "the highest talent density of all the labs." For Meta, the acquisition looks like a way to bring in both a team and a product aimed at a fast-growing part of the developer market: tools that let multiple AI agents collaborate on real codebases rather than answer one-off prompts.

KEY FACTS

ACQUIRER
Meta

ACQUIRED COMPANY
Conductor

CONDUCTOR FOUNDED
2024

CLOUD SERVICE SHUTDOWN
November 20

MUSE AGENTS USERS
Over 10 million

META CODING AGENT LAUNCHED
Muse Code, August

WHY IT MATTERS

For Conductor's cloud customers, the clock is ticking: the service closes on November 20 and data will be deleted. For everyone else, the deal signals that Meta is buying its way into the competitive market for AI coding tools, where parallel multi-agent workflows are becoming the new battleground.

Read the original at TechCrunch →

---

04 MEDIUM IMPACT huggingface.co

Moonshot Ships the Frontier-Sized Weights It Promised for November 4

> Moonshot AI has published the open weights for Kimi K3.1 Pro, a frontier-scale multimodal model it had promised to ship on November 4.

The release makes good on a dated commitment: Moonshot AI said the model would arrive on November 4, and the weights are now live on Hugging Face, where developers can download them and run the model through several standard inference tools, including vLLM, SGLang, and Docker Model Runner. The company describes Kimi K3.1 Pro as its most capable model in the 3-trillion-parameter class, trained with large-scale reinforcement learning across hundreds of downstream domains. It handles both text and images and offers a context window of one million tokens — enough to take in a long book or a large codebase in a single request.

The practical significance is that a frontier-class model is now available for self-hosting rather than only through a paid API. Moonshot AI also offers the model through its own platform at platform.kimi.ai with OpenAI- and Anthropic-compatible interfaces, but the open weights mean companies with the hardware to run it can keep data in-house and avoid per-token fees. The model is released under a custom license, the Kimi K3.1 Pro License, which users will need to review before commercial deployment.

The most striking detail in the release is the breadth of the benchmark table. Moonshot AI reports scores across coding, agentic, reasoning, multimodal, and cybersecurity tests, often placing Kimi K3.1 Pro ahead of named competitors on tasks such as Terminal-Bench 4.0 (65.7 versus 57.9 for Claude Fable 5.1 and 58.2 for GPT-6 Astra) and BrowseComp (92.1 versus 91.5 for GPT-6 Astra). On other benchmarks, such as ExploitBench, it trails GPT-6 Astra by a wide margin (74.7 versus 100). The mixed results are typical of frontier releases and underline that no single model leads everywhere.

For developers, the release includes a notable operational requirement: in multi-turn conversations and tool calls, the full assistant message — including reasoning content and tool calls — must be passed back to the model unchanged, or it will lose track of its own prior thinking. That is a small but consequential detail for anyone integrating the model into an application.

KEY FACTS

RELEASE DATE
November 4

CONTEXT WINDOW
1 million tokens

MODALITIES
Text and image

LICENSE
Kimi K3.1 Pro License

API ACCESS
platform.kimi.ai

DOWNLOADS LAST MONTH
797

WHY IT MATTERS

Businesses that want frontier-level multimodal AI without sending data to a third-party API now have a self-hostable option, though they will need serious hardware and must review the custom license. Developers integrating the model need to preserve its full reasoning output between turns or risk breaking multi-step tasks.

Read the original at huggingface.co →

---

05 MEDIUM IMPACT huggingface.co

A 2-Billion-Parameter Model Pulls Spreadsheet Fields From Any Document

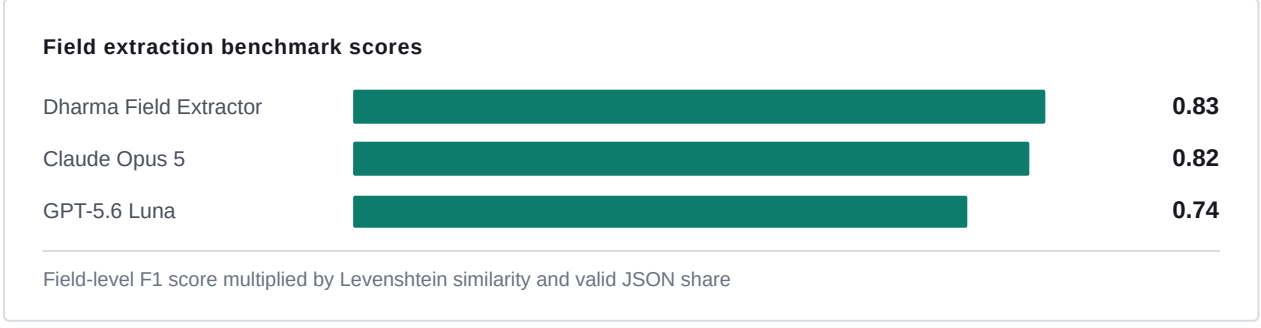
> Dharma-AI has released a small, self-hosted model that pulls specific fields out of documents and returns them as structured data, outperforming every frontier LLM it tested at a fraction of the cost.

The company behind DharmaOCR has built a field-extraction model that takes a document image and either a request for everything it can find, or an exact schema to fill, and returns structured JSON. An invoice comes back with every line item as an ordered list, each with its own quantity and value. An ID comes back as the fields you will store, under the key names your system already uses. The same checkpoint handles all document types with no per-type fine-tuning.

The model was trained on 43,384 examples covering receipts, invoices, forms, bank statements, ID documents and synthetic documents, in English, Portuguese and other languages. It is built on a 2-billion-parameter OCR model, small enough to run self-hosted on a single L40S GPU at $2.24 per hour. On a benchmark of 546 documents and 1,638 cases, it scored 0.834, above Claude Opus 5 at 0.816 and GPT-5.6 Luna at 0.741. Cost per extraction ran from roughly 6 times lower than GPT-5.6 Luna to roughly 147 times lower than GPT-5.6 Sol.

The most striking detail is the generalization test. On a SROIE receipt, the model captured all 18 printed fields correctly, returned null for the two blank fields, and matched every number to the cent. On a FUNSD form, it marked exactly the one checked box as true and the other eleven as false, read a handwritten ID number correctly, and returned null for two blank fields. Neither dataset was used in training.

This is the second time Dharma-AI has measured a small specialized model against frontier APIs and won. Its earlier OCR paper put a 3-billion-parameter model above every commercial API tested on Brazilian Portuguese, and the gap held when newer models arrived. The thesis is consistent: on a bounded task, a small model trained close to that task beats a general-purpose model asked to do the same thing.



KEY FACTS

BENCHMARK SCORE
0.834

CLOSEST FRONTIER MODEL
Claude Opus 5 at 0.816

COST VS FRONTIER APIS
6 to 147 times lower

TRAINING EXAMPLES
43,384

BENCHMARK SIZE
546 documents, 1,638 cases

SELF-HOSTED GPU COST
$2.24 per hour on a single L40S

> WHY IT MATTERS
>
> Businesses that extract fields from high volumes of documents — invoices, IDs, medical requisitions, contracts — can get better accuracy than frontier APIs at a fraction of the cost, on their own infrastructure.

Read the original at huggingface.co →

SECTION 3 OF 3

AI Applications & Industry

6 stories 1 high 5 medium

06 HIGH IMPACT TechCrunch

Data-Center Backlash Won Races; Congress Is the Next Battleground

> Anger over data centers became a decisive midterm issue, and with Democrats now in control of Congress, the AI industry faces its first serious chance of federal limits on where and how it builds.

The November 3 midterm results handed Democrats full control of Congress, and the clearest signal from the vote is that data centers — the large facilities housing the computers that power AI — have become a political liability. An NBC exit poll found more than half of voters believe AI will do more harm than good, almost 80% oppose data centers in their own neighborhoods, and nearly 60% said data centers were a crucial or major factor in their vote. Candidates who campaigned against the buildout won seats in Texas, Maine, Michigan, Nevada, Georgia and South Carolina, while an Ohio county where Meta had planned a 700,000-square-foot facility passed an outright ban.

The practical consequences are already being priced in. Goldman Sachs predicts construction delays in Georgia and Texas because of Democratic gains, and Morgan Stanley analysts say Ohio's Democratic governor-elect Amy Acton poses a major risk to the state's data center buildout — including the $500 billion facility OpenAI has eyed there, backed by more than $100 billion in financing from Nvidia. Likely next steps include construction moratoriums and requirements that data centers pay for their own grid upgrades rather than passing costs to nearby residents.

In Congress, Democratic control improves the odds for bills like the one introduced by Senator Bernie Sanders and Representative Alexandria Ocasio-Cortez, which proposes a nationwide moratorium on new data centers until guardrails address electricity costs, environmental harm and community approval. A separate bill, the GRID Savings Act from New Mexico Senator Martin Heinrich, would stop grid infrastructure costs from being passed to communities during large AI buildouts. Any such legislation would still face a possible presidential veto.

The most striking detail is how little money moved the outcome. Pro-AI super PACs spent millions this cycle, yet voters in state after state rejected the buildout. "It is clear that money does not buy outcomes," said Anshi Bhatt, a tech and policy researcher. Nadine Smith of Color of Change framed the stakes bluntly: "If AI infrastructure is truly a national priority, then the companies building it can afford to pay for their own power, meet clean air standards, disclose what they use, and build with the consent of the people who live next door."

KEY FACTS

VOTERS WHO OPPOSE DATA CENTERS IN THEIR NEIGHBORHOOD
Almost 80%

VOTERS WHO SAID DATA CENTERS WERE A CRUCIAL OR MAJOR VOTE FACTOR
Almost 60%

OPENAI'S EYED OHIO DATA CENTER COST
$500 billion

NVIDIA FINANCING FOR THE OHIO PROJECT
More than $100 billion

META'S PLANNED OHIO FACILITY SIZE
700,000 square feet

> WHY IT MATTERS
>
> Companies planning data center projects should expect longer timelines, higher local costs and possible federal restrictions. Investors in AI infrastructure face new regulatory risk in states that were previously welcoming.

Read the original at TechCrunch →

---

07 MEDIUM IMPACT TechCrunch

AI Political Money Lost More Races Than It Won

> Political groups on both sides of the AI regulation fight spent millions trying to shape Tuesday's midterm elections, and many of their most heavily backed candidates lost.

Two industry-aligned political action committees, American Mission and Building the Future, together spent $4.8 million backing South Carolina Senate candidate Darline Graham, who lost. American Mission — the Republican spending arm of Leading the Future, a group backed by OpenAI president Greg Brockman and the venture firm Andreessen Horowitz — also spent more than $500,000 each on two losing House candidates. Leading the Future's Democratic arm, Think Big, had better results backing Florida Rep. Jared Moskowitz, who won reelection.

Groups pushing for stronger AI safeguards saw a similar split. Public First, backed by Anthropic, contributed $3 million to a PAC supporting California state Sen. Scott Wiener, who was leading his race at publication. Guardrails Alliance, which wants to make taking AI-industry money a political liability, spent more than $1 million each opposing two candidates who lost, and $300,000 supporting Angela Gonzales-Torres, who defeated Rep. Jimmy Gomez.

The results are hard to separate from the broader Democratic wave on Tuesday — Elon Musk reportedly spent half a billion dollars backing Republicans and also saw many candidates lose. Leading the Future told TechCrunch it considers primaries a better use of its money than general elections, and spent $27.38 million in the primaries, including $8 million opposing a New York Assembly member who lost his congressional bid in June. Charlie Blaettler of Guardrails Alliance offered a different explanation for the industry group's lighter general-election spending: "They don't want their money to become an issue in these elections." He cited a Michigan Senate primary where a pro-Israel lobbying group's $30 million in spending became the story and its candidate lost. "We know people hate AI billionaires," Blaettler said.

The underlying fight is whether Congress should create a single national AI framework that overrides state laws, or whether states keep their own authority to impose safeguards. Leading the Future favors a national

framework and plans to keep spending in 2028; Guardrails Alliance says it will keep opposing AI-industry-backed candidates, with the stated goal of making "AI money politically toxic."

KEY FACTS

SPENT ON LOSING SOUTH CAROLINA SENATE CANDIDATE
$4.8 million

LEADING THE FUTURE PRIMARY SPENDING
$27.38 million

SPENT OPPOSING NEW YORK ASSEMBLY MEMBER ALEX BORES
$8 million

PUBLIC FIRST CONTRIBUTION TO SCOTT WIENER PAC
$3 million

GUARDRAILS ALLIANCE SPENDING AGAINST JIMMY GOMEZ
More than $1 million

GUARDRAILS ALLIANCE SPENDING SUPPORTING ANGELA GONZALES-TORRES
$300,000

> **WHY IT MATTERS**
>
> For companies and investors watching AI regulation, the results suggest that AI-linked political money does not reliably buy election outcomes — and that being publicly tied to AI-industry spending may even hurt candidates in competitive races.

Read the original at TechCrunch →

---

08 MEDIUM IMPACT TechCrunch

OpenAI Pitching a New Deal for the AI Economy Ahead of 2028

> OpenAI's top policy executive is urging politicians to craft a New Deal-style agenda for AI ahead of the 2028 election, framing the technology as a path to shared prosperity while critics say his proposals do more to advance the company's interests than to hold it accountable.

Chris Lehane, OpenAI's chief global affairs officer, published an essay comparing AI's economic potential to Franklin D. Roosevelt's New Deal and laying out five policy priorities spanning worker protections, education, AI safety, and data center development. The essay lands just after Democrats flipped control of Congress in the midterms amid voter frustration over the cost of living, and Lehane argues that Americans' anxiety about losing control of their economic futures demands a "compelling third path" that lets people share in AI's benefits.

Kate Brennan, who leads AI and corporate power policy at the AI Now Institute, told TechCrunch the essay marks a shift from OpenAI's agenda last year, which focused on accelerating AI development and reducing regulatory barriers. She called the New Deal framing "ahistorical at best and harmful at worst," noting that the original New Deal constrained corporate power while Lehane's vision offers few binding requirements. His tax proposals — incentives for workers to learn AI skills, easier capital access for entrepreneurs, new approaches to employee ownership — stop short of ensuring workers share in AI-generated wealth or that companies pay the costs of displacing workers.

On data centers, Lehane acknowledges concerns about electricity costs but emphasizes construction jobs and local tax revenue, and warns that politicians backing a moratorium could jeopardize America's lead over China — a comparison he extends to the 2002 vote authorizing the Iraq war. He also calls for enforceable AI safety standards, independent evaluations, and kill-switch capabilities, positions Brennan called "disingenuous" given

OpenAI's 2024 opposition to California's SB 1047 safety bill and its earlier push to limit state regulation. The company has since backed some state safety laws.

The most striking gap, according to critics, is accountability. Lehane ties middle-class success to participation in an economy increasingly shaped by companies like OpenAI, but offers no concrete mechanism to prevent AI's gains from concentrating among a handful of wealthy firms. OpenAI did not respond to TechCrunch's request for specifics.

KEY FACTS

ELECTION TARGETED
2028 presidential election

POLICY PRIORITIES OUTLINED
Five, spanning worker protections, education, AI safety, and data centers

MIDTERM RESULT CITED
Democrats flipped control of Congress

OPENAI'S 2024 POSITION ON CALIFORNIA SB 1047
Opposed the AI safety bill

> **WHY IT MATTERS**
>
> Businesses and workers should watch whether these proposals translate into binding rules — like revised corporate tax rates or incentives to hire humans over automation — or remain voluntary incentives that leave AI companies free to capture most of the economic upside.

Read the original at TechCrunch →

---

09 MEDIUM IMPACT TechCrunch

Apple Rations Its Engineers' AI Tokens From $300 to $200 a Day

> Apple has cut the daily amount its engineering teams can spend on AI from $300 to $200, with further reductions already being planned.

Apple lowered the daily AI spending cap for its engineering teams in October, reducing it from $300 to $200, Bloomberg reported on Wednesday. More cuts are in the works, with some allowances falling to $100 a day, according to sources cited by the outlet. Outside engineering, Apple is reportedly weighing a daily limit as low as $25 a day. TechCrunch has reached out to Apple for comment.

The move marks a sharp reversal from earlier this year, when companies pushed staff to use AI as much as possible in the name of efficiency gains. That enthusiasm was short-lived: as the bills for AI usage came due, companies began deciding that rationing access was the more fiscally sensible approach.

The cuts reflect a broader concern across the industry about a gap between what AI costs to build and run and the productivity gains it is supposed to generate — gains that remain, in many cases, unclear. For Apple, which has worked to integrate the technology into its engineering processes, the reduced allowances signal that even well-resourced companies are now scrutinising the return on every dollar spent on AI tools.

What is genuinely new here is not that AI costs money, but that a company of Apple's scale is imposing hard daily limits on its own engineers. The $300-to-$200 reduction is a concrete number that puts the industry's cost anxiety into sharp relief: even the most valuable company in the world is treating AI access as a budget line to be managed, not an unlimited resource.

KEY FACTS

PREVIOUS DAILY CAP
$300

NEW DAILY CAP
$200

SOME TEAMS' PLANNED CAP
$100

POSSIBLE NON-ENGINEERING CAP
$25

TIMING OF CUT
October

> WHY IT MATTERS
>
> If Apple is rationing AI access for its own engineers, businesses of all sizes should expect AI tool budgets to face the same scrutiny — and should measure whether the productivity gains justify the spend before bills arrive.

Read the original at TechCrunch →

---

10 MEDIUM IMPACT Ars Technica

North Korea's $6 Billion Binance Heist Went Through Its AI Tools

> Binance said on Tuesday that hackers who stole billions of dollars in cryptocurrency spent months infiltrating the exchange, using fake recruitment interviews, malicious software and compromised internal AI tools.

The theft, which the FBI and security firm TRM have attributed to North Korean actors, initially targeted $6.34 billion in cryptocurrency from Binance wallets — the largest sum ever targeted in a cryptocurrency theft. Binance has since frozen or intercepted $2.85 billion of that amount and said it has set aside funds to cover the loss.

The attackers began their campaign no later than June and continued until last Thursday, when they compromised an engineer's laptop and withdrew assets from several wallets. Binance described the operation as a multi-stage campaign rather than a single breach or simple phishing incident. The first known compromise occurred in June, when an engineer downloaded a tampered version of Homebrew, a widely used software installation tool. Binance's security team detected the compromised laptop, but only after the malware operators had already gathered information about the engineer's role and those of colleagues on the Binance wallet infrastructure team.

The most striking detail is the use of compromised internal AI tooling as part of the attack chain. Binance listed it alongside social engineering, malicious downloads and endpoint compromise as techniques the attackers combined. The company did not specify which AI tools were involved or how they were exploited, but the inclusion signals that internal machine-learning systems — the kind many large companies now rely on for security, fraud detection and operations — can themselves become an attack surface.

For businesses, the incident is a reminder that sophisticated attackers now treat corporate AI infrastructure the same way they treat any other software: as something to be studied, poisoned or turned against its owners. The months-long timeline also underscores that the most damaging breaches are rarely the result of a single mistake, but of patient, layered campaigns that survive initial detection.

KEY FACTS

AMOUNT TARGETED
$6.34 billion

AMOUNT FROZEN OR INTERCEPTED
$2.85 billion

CAMPAIGN START
No later than June

ATTRIBUTION
North Korean actors, per FBI and TRM

FIRST KNOWN COMPROMISE
Tampered Homebrew download in June

> WHY IT MATTERS
>
> Companies that have integrated AI tools into internal operations should treat those systems as part of their attack surface, not just as defensive or productivity assets. The months-long timeline shows that detecting one compromised device may not stop a campaign that has already mapped the organisation.

Read the original at Ars Technica →

---

11 MEDIUM IMPACT TechCrunch

Nothing Ships a Screenless AI Device in December

> Smartphone maker Nothing says its first AI-native product — a screenless device CEO Carl Pei describes as a productivity tool — will ship in December.

Nothing, the startup behind the distinctive Nothing Phone, is about to enter the AI hardware race, but it is refusing to say exactly what it is building. CEO Carl Pei posted on X on Wednesday that the company's first "AI-native product" ships in December. "It's a tool, but one built with the personality and care that have come to define a Nothing product. It has no new screen, and it couldn't have been an app," Pei wrote.

The announcement is deliberately thin on detail. Pei has previously teased an AI-first device, and the company has said it plans to release AI-linked smart glasses and earbuds with AI integrations. But Pei appears to be distancing this product from the wave of AI wearables — pendants, pins and companion devices — that other startups have pushed as the consumer breakthrough for generative AI. "The whole industry is hunting for its iPhone moment," Pei wrote. "I admire the boldness, but I read this moment differently."

For anyone watching the consumer AI market, the timing matters. Hardware startups such as Humane and Rabbit shipped AI devices to lukewarm reviews, and the category has yet to produce a product that justifies replacing a smartphone app. Nothing is betting that a screenless, non-app form factor can succeed where those devices struggled — a contrarian position from a company that has built its brand on distinctive industrial design rather than raw specifications.

No firm release date has been announced beyond December, and Nothing has not disclosed pricing or distribution. The company's fans will not have to wait long to see whether Pei's "different" reading of the AI moment translates into a product people actually buy.

KEY FACTS

PRODUCT TYPE
Screenless AI-native productivity tool

SHIPPING
December (no firm date announced)

COMPANY
Nothing

ANNOUNCED BY
CEO Carl Pei on X

> WHY IT MATTERS
>
> Nothing is entering a consumer AI hardware market where earlier screenless devices from Humane and Rabbit failed to gain traction. Its December launch will test whether design-led hardware can succeed where utility-first AI gadgets have not.

Read the original at TechCrunch →

Sources

01 Claude Helps Map a Protein Shell Science Left Unexplained for 40 Years
https://www.anthropic.com/research/investigating-vault-particles-with-claude

02 Robots Fumble Precise Grips Because Their Models Simplify the Camera First
https://arxiv.org/abs/2611.06918

03 Meta Buys the Startup That Runs Coding Agents in Parallel
https://techcrunch.com/2026/11/04/meta-acquires-ai-coding-startup-conductor/

04 Moonshot Ships the Frontier-Sized Weights It Promised for November 4
https://huggingface.co/moonshotai/Kimi-K3.1-Pro

05 A 2-Billion-Parameter Model Pulls Spreadsheet Fields From Any Document
https://huggingface.co/blog/Dharma-AI/dharma-field-extractor

06 Data-Center Backlash Won Races; Congress Is the Next Battleground
https://techcrunch.com/2026/11/04/after-the-midterms-a-bigger-clash-over-data-centers-could-be-coming/

07 AI Political Money Lost More Races Than It Won
https://techcrunch.com/2026/11/04/ai-pacs-spent-millions-trying-to-shape-congress-heres-how-they-fared/

08 OpenAI Pitching a New Deal for the AI Economy Ahead of 2028
https://techcrunch.com/2026/11/04/openai-is-already-trying-to-shape-the-2028-election/

09 Apple Rations Its Engineers' AI Tokens From $300 to $200 a Day
https://techcrunch.com/2026/11/04/apple-is-reportedly-cutting-budgets-for-its-internal-ai-usage/

10 North Korea's $6 Billion Binance Heist Went Through Its AI Tools
https://arstechnica.com/information-technology/2026/11/binance-says-hackers-who-stole-billions-in-crypto-spent-months-planning-heist/

11 Nothing Ships a Screenless AI Device in December
https://techcrunch.com/2026/11/04/nothing-ceo-says-companys-mysterious-ai-first-device-is-dropping-next-month/

About this document. Every story in the 5 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 2026-11-05T06:34:48Z · model deepseek-v4-pro:cloud · New Horizon, new-horizon.tech · To unsubscribe or manage your subscription, see the footer of the daily email.