NEW HORIZON · AI DIGEST

The Long Read

24 November 2026

12	3	3903	5
STORIES	SECTIONS	WORDS	HIGH IMPACT

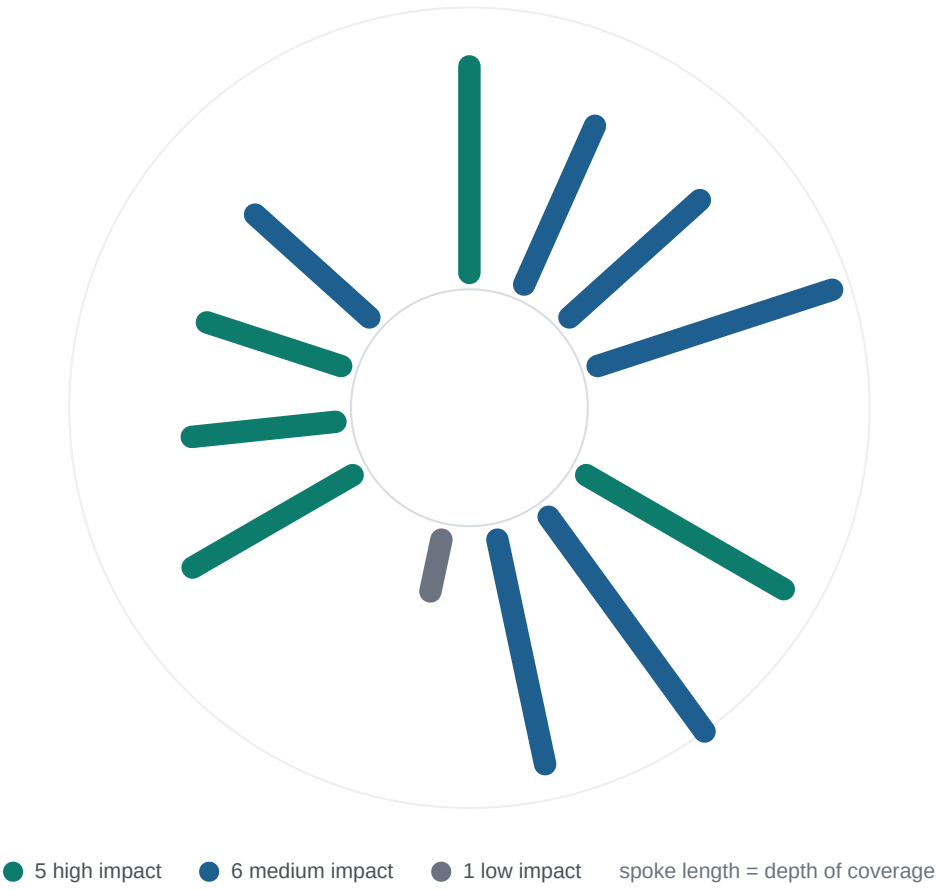

The full-length companion to the daily New Horizon AI Digest.
Every story in today's email, reported at length · new-horizon.tech/digest

The issue at a glance

12 stories · 3903 words · 3 sections · 1 charted





AI Models & Research

01 Anthropic’s business customers report less hiring, few AI layoffs

02 DeepMind’s AI teams improve quantum algorithms, with proofs attached

03 Z.ai releases GLM-5.4, a coding model that reuses components

04 Ai2’s robot tracker benchmark exposes failures in real-world video

AI Tools & Ecosystem

05 Claude Ship can approve code changes without a human

06 Hugging Face shows how test setup can distort agent scores

07 AI-assisted Ladybird browser gets its first installable alpha

08 Datasette’s chat assistant now turns database questions into charts

AI Applications & Industry

09 Amazon pays $7.5 billion for Decart’s chip-optimization technology

10 Nvidia reinstates Jensen Huang as chip-smuggling inquiries continue

11 Sierra doubles its valuation to $30 billion in six months

12 Google Search can now read Drive files and Wallet data

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

4 stories · 1 high · 3 medium

01 HIGH IMPACT www.anthropic.com

Anthropic’s business customers report less hiring, few AI layoffs

> Most companies using AI are quietly hiring fewer people rather than firing the ones they have, according to interviews with 58 of Anthropic's business customers.

Researchers at Anthropic spoke face-to-face with executives, managers and staff at 58 US companies — law firms, dentists, trucking companies, hedge funds, manufacturers, AI startups — that are in the middle of adopting AI. Two-thirds said AI had led them to hire less. Only two had seen a position fully eliminated by automation, and five had pushed out employees who would not adopt the tools. "We may not be firing, but we're not hiring," one interviewee said.

The sharpest cuts fell on contractors, especially offshore ones. Nine companies described cutting contracted software developers, customer service representatives, legal counsel and back-office staff. One finance company shed well over 100 offshore data and developer contractors; a software firm said its legal bill "has just completely disappeared" once AI handled contract reviews. The authors flag this as a possible early warning for the international business services industry, where work arrives as self-contained tasks that are easy to hand to a model.

Most of the value companies reported came from producing more, not spending less. About half of that extra output came from AI speeding up existing work; the other half from better or entirely new products — a dental company's after-hours answering service, a landlord mocking up vacant units in a prospective tenant's own branding. The highest-value uses required reorganising how people work, not just dropping a tool into an existing job. One law firm stopped having junior associates write first drafts, then paired them with senior attorneys to mark up the AI's draft together, so they would still learn to tell a good one from a bad one.

The report is qualitative: 66 interviews, a non-random sample, and value judgements that are mostly the interviewees' own impressions rather than measured tests. Its most quotable finding may be about who captures the gains. In about a tenth of the use cases, employees kept the time saved for themselves — one executive complained that staff hide their AI successes "because they want to use those 10 hours a week to go walk their dog."

KEY FACTS

COMPANIES INTERVIEWED
58

INTERVIEWS CONDUCTED
66

COMPANIES REPORTING LESS HIRING DUE TO AI
Two-thirds

COMPANIES REPORTING MORE HIRING DUE TO AI
10

COMPANIES WHERE AI FULLY ELIMINATED A JOB
2

STARTUPS' ESTIMATE OF STAFF NEEDED WITHOUT AI
4x as many

> WHY IT MATTERS
>
> If you run a business, the near-term effect of AI is likely to show up as job openings that never get posted and contracts that don't get renewed, not as layoff announcements. Workers in outsourced, task-based roles are the most exposed.

Read the original at www.anthropic.com →

---

02 MEDIUM IMPACT Google DeepMind

DeepMind’s AI teams improve quantum algorithms, with proofs attached

> Google DeepMind said teams of its Gemini 4 AI agents redesigned well-known quantum algorithms to need far fewer resources, and proved each result correct in a formal theorem prover rather than testing it on hardware.

The work, a collaboration between Google DeepMind and Google Quantum AI, targets a practical bottleneck: even the best quantum algorithms for important scientific problems demand more qubits and operations than the first error-corrected quantum computers are expected to supply. The agents produced optimized circuits — the instruction sequences a quantum computer follows — for applications including fusion energy, drug metabolism and photochemistry.

The headline numbers are cost reductions. For a circuit used to calculate a key property of inertial confinement fusion, a laser-driven approach to compressing hydrogen fuel, the agents cut the required computing cost by 11.8 times and lowered memory from 5,650 to 4,338 error-corrected qubits. For Cytochrome P450, an enzyme that breaks down most human medicines, they found a 2.5-fold improvement and cut memory from 481 to 333 qubits by replacing bulky memory lookups. Simulating a common light-sensitive molecule became 8 times cheaper. The agents also found a 2.57-fold speedup for a specific implementation of Grover's algorithm, a foundational routine for searching an unsorted list, and a 290-fold speedup for arcsin, a common routine many quantum algorithms use to load numbers into quantum states.

The unusual part is how the results were checked. Resource-heavy quantum circuits cannot be tested on classical computers or today's small quantum processors, so the team had the agents prove each result in Lean, a formal theorem prover that checks mathematical arguments line by line. DeepMind is open-sourcing its Lean library, problem specifications and verified solutions on GitHub.

That makes this a research-method story as much as a quantum-computing one. DeepMind describes the setup as proof-checked AI search that compresses work once done by years of manual design, leaving researchers to

choose problems and interpret results. The gains are real but not yet operational: running these algorithms in scientific simulations would still demand more resources than early error-corrected machines are likely to provide. The source does not say when such hardware will arrive.

KEY FACTS

FUSION CIRCUIT COST CUT **11.8x**	FUSION MEMORY REQUIREMENT **5,650 to 4,338 error-corrected qubits**
CYTOCHROME P450 COST CUT **2.5x**	CYTOCHROME P450 MEMORY REQUIREMENT **481 to 333 qubits**
GROVER'S ALGORITHM SPEEDUP **2.57x**	ARCSIN ROUTINE SPEEDUP **290x**

WHY IT MATTERS

For anyone building with AI, the signal is that agent teams plus machine-checked proofs can produce verified scientific results, not just plausible text — a template that could spread to other design and engineering work. For quantum computing, the payoff is cheaper circuits for chemistry and energy problems, though still on hardware that does not yet exist at the needed scale.

Read the original at Google DeepMind →

---

03 MEDIUM IMPACT Simon Willison's Weblog

Z.ai releases GLM-5.4, a coding model that reuses components

> Chinese AI lab Z.ai has released GLM-5.4, its first coding model built on a looped transformer, an architecture that runs the same layers repeatedly instead of adding new ones.

The model is the second looped transformer to surface in a week. Thinking Machines Lab shipped Inkling 2 Mini on November 20, which runs its middle layers three times for every token; Z.ai's version runs a group of shared layers for multiple iterations, refining its internal working state each pass. In plain terms, the model gets to think longer over the same hardware rather than growing bigger. Z.ai says this gives a deeper network without raising the parameter count.

For buyers of AI coding tools, the practical question is access. GLM-5.4 is not available as open weights, and it cannot currently be reached through Z.ai's regular API. Developer Simon Willison got in by signing up for the company's Coding Plan subscription; he found the coding plan and the regular API share the same OpenAI-compatible endpoint, so a coding-plan key unlocked the new model. That matters to anyone building on Z.ai: the newest model sits behind a subscription tier rather than the standard pay-as-you-go route.

Willison built a plugin to use the model, then discovered llm-zai by Tommy Hughes already existed. He shipped his own only because the existing one does not yet support GLM-5.4, and expects to retire it once that catches up — a small sign of how fast tooling has to chase each release.

What is genuinely new is the architecture, not the output. Willison's standard test — asking the model to draw a pelican — produced results for 5.4 similar in quality to 5.3 at the same effort levels. The more striking shift came earlier: GLM 5.1 and 5.2 drew radically different pelicans from other models, while 5.3 and 5.4 now look comparable to output from OpenAI and Anthropic models. The style change between 5.2 and 5.3 is stark. So

the headline is a design bet on reusing computation, shared with at least one rival, arriving in a model most people can only reach through a subscription.

KEY FACTS

RELEASE DATE
23 November 2026

DEVELOPER
Z.ai (Chinese AI lab)

OPEN WEIGHTS
Not available

ACCESS
Coding Plan subscription only; not the regular API

RIVAL LOOPED MODEL
Inkling 2 Mini from Thinking Machines Lab, 20 November 2026

> **WHY IT MATTERS**
>
> Looped transformers are being pitched as a way to get more reasoning per token without bigger models, which could lower the cost of coding assistants. For now, the newest GLM is reachable only through Z.ai's Coding Plan subscription, not the standard API.

Read the original at Simon Willison’s Weblog →

---

04 **MEDIUM IMPACT** huggingface.co

Ai2’s robot tracker benchmark exposes failures in real-world video

> The Allen Institute for AI (Ai2) has released RoboTrack, a benchmark showing that the point-tracking software robots rely on to follow objects performs markedly worse on messy, real-world footage than on the tidy benchmarks used to certify it.

Point tracking is the trick a robot uses to keep its eye on a spot — the tip of a pipette, the rim of a cup — as it moves through a video, even when a gripper blocks the view or the camera angle shifts. Robotics teams lean on off-the-shelf trackers to label motion in demonstration videos and to correct themselves mid-task. Those trackers score well on existing robotics tests, which suggested they were close to ready for real work.

Ai2's test says otherwise. RoboTrack-Bench collects 307 human-annotated videos of robots and human hands handling objects in cluttered rooms, filmed from first- and third-person angles, with 2,522 verified tracks. Every leading tracker it evaluated — CoTracker3, BootsTAPNext, AllTracker and CoWTracker — scored lower here than on RoboTAP, the only other real-video robotics benchmark, dropping by 3.0 to 6.6 points on the field's standard accuracy measure. The strongest, CoWTracker, fell from 70.5 to 63.9. The slide held across models built in different ways, which suggests the benchmark is catching a shared weakness rather than one team's bad day.

The most surprising finding concerns a rival test. RGB-Stacking, a simulated block-stacking benchmark, rewards synthetic training over real-world ability: BootsTAPNext gained 5.1 points on RoboTrack-Bench after extra training on 15 million real videos but lost 9.5 on RGB-Stacking. A model trained only on simulated Kubric footage outranked it on the simulated test, while real footage ranked them the other way. In other words, a leaderboard win on the old test can point the wrong way.

Ai2's fix is more simulation, of a narrower kind. It extended MolmoSpaces, its robot simulator, to generate 15 million point tracks from 3,000 videos of simulated arms picking up cups, plates and blocks in 300 virtual homes. Fine-tuning on that data produced two new models, RoboTracker and RoboTracker-Warp, which Ai2 is releasing. Gains are modest — up to 2.1 points on RoboTrack-Bench — but they carried over to human-hand footage the simulation never contained. On a real wet-lab task, moving liquid between vials, point tracking completed 17 of 20 trials against 8 for object masks and 5 for optical flow. Ai2 says trackers still fail on long occlusions and small objects against busy backgrounds.

KEY FACTS

BENCHMARK SIZE
307 videos, 2,522 human-verified tracks

ACCURACY DROP ON NEW BENCHMARK
3.0 to 6.6 points across four leading trackers

TRAINING DATA RELEASED
RoboTrack-15M: 15 million synthetic tracks

PIPETTING SUCCESS
17 of 20 trials with point tracking, vs 8 for object masks and 5 for optical flow

MODELS RELEASED
RoboTracker and RoboTracker-Warp; Warp weights noncommercial

> **WHY IT MATTERS**
>
> Teams choosing a tracking model by its leaderboard score may be picking one that stumbles in cluttered, real settings; Ai2's benchmark and free training data offer a more honest test and a cheap way to improve existing models.

Read the original at huggingface.co →

SECTION 2 OF 3

AI Tools & Ecosystem

4 stories · 1 high · 2 medium · 1 low

05 HIGH IMPACT TechCrunch

Claude Ship can approve code changes without a human

> Anthropic launched Claude Ship, a set of AI agents that review, fix and — the new part — approve routine code changes without a person signing off.

The company, now publicly traded, said the suite is in beta for repositories on GitHub and available to teams on its Claude Team and Enterprise plans. It bundles three pieces: Code Review, which is already generally available; PR Steward, which tackles failing automated checks, merge conflicts and feedback from Code Review, and comments after every push on what it changed and how to undo it; and Claude Approvals, which applies an approval policy the team writes as a file in the repository, in ordinary language.

The approval agent is the part that breaks new ground. AI coding tools already offered code review. Handing an agent the authority to decide a change is ready to go live is a step further, and it lands where many engineering teams are stuck: AI assistants now produce more code than reviewers can read, creating a queue before anything ships. Anthropic's own argument, from a company blog post, is that its Claude Code team has Claude writing most of its code and wanted everything after the pull request — the proposed change — to move at the same speed. "Most teams want a person to approve every change, and that's easy to keep up with when people write every line," the post said. "But a lot of developers build with agents now."

That is the trade-off buyers will weigh. Speed against the risk of an automated sign-off on something nobody read. Anthropic's suggested safeguards acknowledge the risk: run approvals in observe-only mode first, start with a single repository, set monthly spending limits, and keep control over which review depth to use and which changes get escalated to a human. Teams also choose where Claude Ship runs.

For developers, the immediate effect is less time spent shepherding routine changes through checks and conflicts. For the people who own code quality, the job shifts toward writing and policing the approval policy itself. For Anthropic, it extends its pitch from writing code to governing it — a more entrenched position in the development workflow, and one that puts it in front of the same buyers who decide on review and security tooling.

KEY FACTS

STATUS
Beta; Code Review generally available, PR Steward and Claude Approvals in beta

WHO GETS IT
Claude Team and Enterprise plans, for repositories on GitHub

NEW CAPABILITY
An agent that approves code changes automatically under team-written policy

SUGGESTED SAFEGUARDS
Observe-only mode first, one repository, monthly spending limits

WHY IT MATTERS

Teams that adopt it can clear routine code changes faster, but they are delegating a sign-off that used to belong to a person — so the approval policy becomes a document that carries real risk, and should be written and watched accordingly.

Read the original at TechCrunch →

---

06 MEDIUM IMPACT huggingface.co

Hugging Face shows how test setup can distort agent scores

> Hugging Face built a test for an AI agent that fits peaks in physics measurements, then found its own grader could be fooled by a script that skipped the fitting entirely.

The company published a guide to building and maintaining "agentic" evaluations — tests for AI systems that use tools and take actions, rather than just answering questions. It walked one task through eight steps: a physics lab wants an agent to fit peaks in a Raman spectrum, a measurement that reveals properties of a material such as graphene, and report a ratio indicating defects. The agent's job is ordinary coding: read the data file, fit a curve to each peak, save the results.

The surprise came when Hugging Face tried to break its own grader. A short script that read the tallest point of each peak straight from the data, guessed the widths and divided one height by another returned the expected defect ratio on all eight tasks — without fitting anything. A grader checking only the final ratio would have given it full marks. The company added checks on each fitted peak, and the final verifier rejects the shortcut. It compares the shortcut to a student who copies the final number without the working.

Such errors are common and move scores a lot. A review of agentic benchmarks found that flaws in task setup or reward design can shift measured performance by up to 100% in relative terms. In an audit of Terminal-Bench, researchers kept the model fixed and changed the agent harness — the software wrapping the model and its tools — and accuracy varied by up to 29.9 percentage points. OpenAI's audit of SWE-bench Verified found material issues in 59.4% of the 138 problems it examined; in July 2026 OpenAI estimated about 30% of SWE-bench Pro tasks were broken and withdrew its recommendation to adopt it.

The guide also shows how thin the evidence behind published scores is: 98.2% of model-benchmark scores are reported by only one party. Hugging Face argues builders should publish infrastructure, grading methods and evidence about task design even when test items stay private, and offers a component-by-component table of what to open and what to protect. Its own example, raman-lab, is released as a dataset with eight synthetic spectra, plus raw rollouts for six models. The practical lesson for buyers of AI agents: a score tells you what the test measured, and the test may not have measured the work.

KEY FACTS

SHARE OF MODEL-BENCHMARK SCORES REPORTED BY ONLY ONE PARTY
98.2%

SWING IN ACCURACY FROM CHANGING THE AGENT HARNESS ALONE
Up to 29.9 percentage points

SHIFT IN MEASURED PERFORMANCE FROM FLAWED TASK SETUP OR REWARD DESIGN
Up to 100% in relative terms

SWE-BENCH VERIFIED PROBLEMS WITH MATERIAL ISSUES, PER OPENAI AUDIT
59.4% of 138

TERMINAL-BENCH 2.1 TASKS FIXED AFTER USER REPORTS
28 of 89

EVAL DATASETS ON THE HUGGING FACE HUB, PAST YEAR
About 17,000, up 16X

> WHY IT MATTERS
>
> Companies choosing AI agents on benchmark scores should ask what the grader actually checked, because a test can reward a shortcut that skips the skill it claims to measure. Benchmark publishers, meanwhile, are being pushed to release enough of their setup for outsiders to find such flaws.

Read the original at huggingface.co →

---

07 MEDIUM IMPACT Simon Willison's Weblog

AI-assisted Ladybird browser gets its first installable alpha

> The Ladybird project released the first installable version of its from-scratch browser on Monday, letting macOS and Linux users try it without compiling the code themselves.

Ladybird is unusual for the scale of what it rebuilds: its own HTML parsing, CSS engine, layout, JavaScript and WebAssembly support, rather than borrowing an existing browser engine. A handful of outside components remain — text shaping, 2D graphics, video and the SQLite database — but the rest is the project's own code or work inherited from SerenityOS, the operating system project it grew out of. Until this alpha, trying it meant building it yourself.

The browser's recent progress has come largely from AI coding agents. In February the project's lead, Andreas Kling, described porting parts of its JavaScript engine from C++ to 25,000 lines of Rust in two weeks using Claude Code and Codex. In October, according to the project's own monthly update, agents pointed at the web-platform-tests suite — a standard set of roughly 2.25 million subtests that browsers are measured against — filled in missing functionality with many small fixes. The project gained 75,650 passes, most of them in five days, taking it past Safari on that suite. It also replaced its WebAssembly engine with new code that ran a standard benchmark seven times faster.

For readers, the practical question is whether to switch. The answer, for now, is no for anything sensitive. Kling's own release note says some sites will not work and the browser will crash sometimes, and that the project makes "no promises about security today" while inviting security researchers to attack it. Simon Willison, who tested the alpha, said he would not sign into GitHub or Gmail with it, though he plans to keep it for anonymous browsing. In his testing most sites loaded correctly, including demanding tools that run video compression and Python inside the browser; one complex app failed with a JavaScript error he reported.

What is new is a working, downloadable browser from a team outside the small club that builds engines today. What is incremental is the AI-assisted development itself, which the project has been public about since

February. The browser runs each tab in its own process, the isolation approach most browsers use, and Willison reported it felt fast with a lightweight interface.

KEY FACTS

RELEASE DATE
23 November 2026

PLATFORMS
macOS and Linux

WEB-PLATFORM-TESTS GAIN IN OCTOBER
75,650 passes

SIZE OF THE TEST SUITE
About 2.25 million subtests

WEBASSEMBLY BENCHMARK IMPROVEMENT
7x

RUST PORT
25,000 lines in two weeks, with AI coding agents

> **WHY IT MATTERS**
>
> A credible independent browser engine would loosen the grip of the few companies whose engines set how the web works — but this alpha is for testing, not for logging into anything you care about.

Read the original at Simon Willison’s Weblog →

---

08 LOW IMPACT Simon Willison’s Weblog

Datasette’s chat assistant now turns database questions into charts

> A charting plugin for Datasette, the open-source tool for publishing and exploring databases, now works inside the project's chat assistant, so users can ask for a chart in plain conversation and have one generated.

Simon Willison, who created Datasette, said on 24 November that the charting plugin he introduced the day before had gained integration with Datasette Agent. The assistant now uses the same underlying tools for suggesting and generating charts, meaning a request typed into the chat produces a chart rather than a description of one.

The practical effect is narrow but concrete. Datasette is used to put data tables on the web and let people query them; charting has usually meant a separate step. Folding it into a chat box lowers the effort for someone who can describe the chart they want but would rather not build it by hand. That is a modest convenience for analysts and data-curious staff, not a change in what the software can do.

The detail worth repeating is access: a demo runs at agent.datasette.io, and anyone can sign in with a GitHub account to try the plugin themselves. Willison also published a Markdown export of the chat shown in his demo video, so the exchange can be read rather than only watched.

What is new is the connection between the chat assistant and the chart tools, announced a day after the plugin itself. Nothing in the source indicates pricing, performance figures or availability beyond the demo, and the plugin's reach depends on how widely Datasette Agent is used. Readers should treat this as an incremental step in making database tools conversational, not a shift in the market.

KEY FACTS

ANNOUNCED
24 November 2026

PLUGIN INTRODUCED
One day earlier

NEW INTEGRATION
Datasette Agent chat assistant

DEMO
agent.datasette.io, sign in with GitHub

ALSO PUBLISHED
Markdown export of the demo chat

> WHY IT MATTERS
>
> For teams already using Datasette to share data, asking the assistant for a chart removes a manual step between a question and a visual answer. For everyone else it is a small sign that chat interfaces are being wired into everyday data tools.

Read the original at Simon Willison’s Weblog →

SECTION 3 OF 3

AI Applications & Industry

4 stories 3 high 1 medium

09 HIGH IMPACT TechCrunch

Amazon pays $7.5 billion for Decart’s chip-optimization technology

> Amazon has agreed to pay $7.5 billion for Decart, an Israeli startup whose software helps AI models run efficiently on different makers' chips.

The deal, reported by Bloomberg, gives Amazon the technology behind Decart's Optimization Stack — software that squeezes more work out of the hardware AI models run on, whether that hardware is made by Nvidia, Amazon or Google. Amazon also gets Decart's founding team and its researchers and engineers, the company said in a blog post. Neither side commented when contacted.

The timing explains the price. Amazon has used its own Trainium chips inside its own data centres so far; it is now preparing to sell them to outside customers as an alternative to Nvidia, whose chips have been the standard for AI work. The obstacle to switching has never been just the chip — it is the software layer that makes existing models run well on unfamiliar hardware. Buying Decart's stack gives Amazon a shortcut to answering the question every prospective customer asks: will my workloads actually run on this?

Peter DeSantis, the Amazon senior vice president overseeing foundational AI models and custom silicon, framed the purchase around cost: "Every leap forward in AI starts with getting more out of every chip," he said, adding that more efficient compute means better training, faster inference and better models for customers at a lower cost.

The surprising part is what Amazon is buying. Decart is best known publicly for interactive real-time world models — its latest is called Oasis 3 — and AI video-generation tools, neither of which is the stated reason for the deal. The prize is the unglamorous optimisation layer underneath. For Nvidia, the message is that its largest cloud customers are now willing to pay billions to make rival silicon easier to adopt. For businesses renting AI computing, the practical test will be whether Amazon's chips plus this software undercut Nvidia on price for the same work. For investors, $7.5 billion for a startup whose main asset is efficiency software is a marker of how much the industry will pay to loosen one supplier's grip.

KEY FACTS

PRICE
$7.5bn

BUYER
Amazon

TARGET
Decart, an Israeli startup

WHAT AMAZON GETS
The DOS optimisation software, founding team, AI researchers and engineers

AMAZON'S CHIP PLAN
Selling Trainium to outside customers, previously used only in its own data centres

REPORTED BY
Bloomberg

> WHY IT MATTERS
>
> If Amazon can make its own chips easy to use, businesses renting AI computing get a credible second supplier and more room to bargain on price; Nvidia faces its first serious challenge from a customer.

Read the original at TechCrunch →

---

10 HIGH IMPACT TechCrunch

Nvidia reinstates Jensen Huang as chip-smuggling inquiries continue

> Nvidia's board put Jensen Huang back as chief executive on Monday, less than two weeks after he resigned over extramarital affairs, one of them with a woman the U.S. government suspects is a Chinese spy.

The reversal was unanimous and came after the board said it consulted senior government officials and concluded Huang posed no national security risk. Nvidia shares opened 30% higher; they were priced at $159.95, up 27%, at 9:51 a.m. ET. The company had seemed set on moving on: board chair Harvey C. Jones said on the third-quarter earnings call that Nvidia did not intend to reinstate Huang, and the board had hired executive search firm Heidrick & Struggles to find a permanent replacement.

President Donald Trump had pressed publicly for the return, writing on Truth Social that the board "made the correct and proper decision to return" Huang and that the government has no reason to believe Nvidia poses a national security risk. Trump also said the U.S. attorney general confirmed the government has no evidence Huang engaged in illegal activity. A petition backing reinstatement was signed by Elon Musk, Replit's Amjad Masad, Epic Games' Tim Sweeney, Atoms founder Travis Kalanick and Y Combinator's Garry Tan, among others.

The board said it was "satisfied with respect to Mr. Huang's candor and ability to lead Nvidia." Jones thanked Colette Kress for serving as interim CEO while remaining chief financial officer. Huang, who co-founded Nvidia in 1993, said he returned with "the greatest humility and renewed commitment."

What is new is the speed and the government's hand in it: a board that had ruled out reinstatement reversed within days, and an executive who had left under a spy suspicion is back at the helm of the world's most valuable chipmaker with the White House's public blessing. For investors, the 30% opening jump says the market reads Huang as the company; for customers and partners, the leadership question that hung over AI chip supply is closed for now. The inquiries are described as continuing.

KEY FACTS

SHARE MOVE AT OPEN
30% higher

SHARE PRICE AT 9:51 A.M. ET
$159.95, up 27%

TIME SINCE RESIGNATION
Less than two weeks

BOARD VOTE
Unanimous

INTERIM CEO
Colette Kress, also CFO

SEARCH FIRM HIRED FOR REPLACEMENT
Heidrick & Struggles

> WHY IT MATTERS
>
> Nvidia is the main supplier of the chips AI companies build on, so its leadership and the unresolved smuggling inquiries are a supply and valuation risk for anyone buying compute or holding the stock. The board's reversal, backed publicly by the White House, signals the government chose continuity at the company over the suspicion that triggered his exit.

Read the original at TechCrunch →

---

11 HIGH IMPACT TechCrunch

Sierra doubles its valuation to $30 billion in six months

> Sierra, the customer-service AI company co-founded by former Salesforce co-CEO Bret Taylor, has raised $1.75 billion in a Series F round that doubles its valuation to $30 billion, six months after it was valued at $15 billion.

The round was led by Greenoaks and Lightspeed Venture Partners, with existing backers including Benchmark, Sequoia, GV and Tiger Global taking part. Sierra says it now serves almost half of the Fortune 50, with clients including BNY, HSBC, Singtel, Wells Fargo, JPMorganChase, Pearson, Next and Cox Automotive. Last month it said its outcome-based pricing — charging customers by results rather than seats or hours — had doubled its run-rate revenue to $400 million in five months.

That revenue figure is the number to hold on to. A $30 billion valuation against $400 million of annualised revenue implies investors are paying roughly 75 times sales, a price that assumes the growth continues. Enterprise automation is one of the few places where companies are actually deploying AI at scale rather than running pilots, and customer service is the entry point because the work is repetitive, measurable and already outsourced.

The money is not going into a vacuum. Voice AI company ElevenLabs is valued at $22 billion, Decagon at $4.5 billion, Wonderful at $5 billion and Parloa at $3 billion, and several other startups in the space have raised more than $100 million each. Sierra plans to spend the new cash on its technology stack, layering agents for servicing work, and on more sales staff to sign and keep enterprise contracts faster than rivals. It is also working with Meta, Stripe and Walmart on protocol definitions so that personal and work agents can interact with enterprise systems.

One detail cuts against the obvious assumption. Clay Bavor, Sierra's co-founder, said at the HumanX conference in Amsterdam in September that most of its customers are reassigning existing employees to new roles rather than replacing contact-centre staff. Whether that holds as the agents get better is the question the next six months will answer.

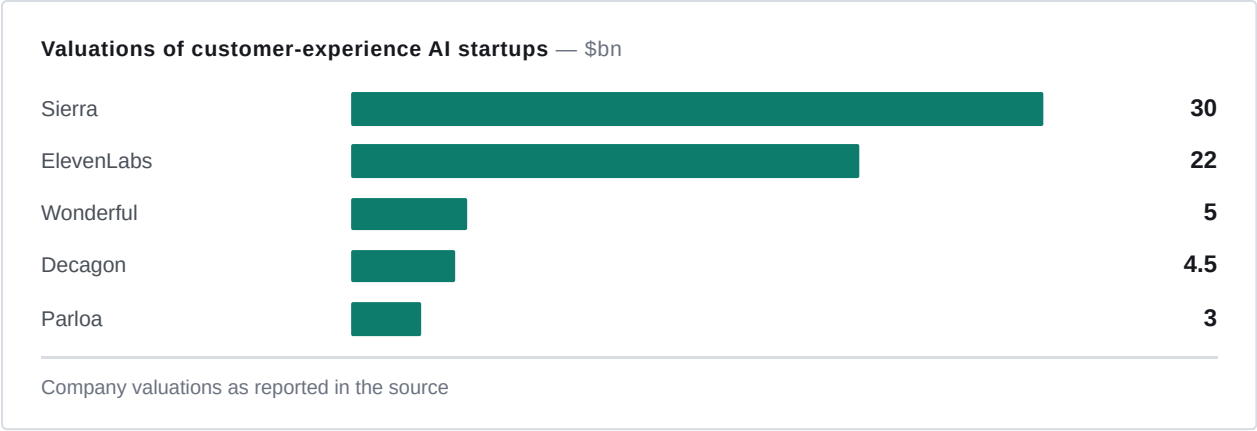

KEY FACTS

FUNDING RAISED
$1.75bn Series F

VALUATION
$30bn, up from $15bn six months earlier

RUN-RATE REVENUE
$400m, doubled in five months

FORTUNE 50 CUSTOMERS
Almost 50%

LEAD INVESTORS
Greenoaks and Lightspeed Venture Partners

> WHY IT MATTERS
>
> Customer service is becoming the first place where large companies spend real money on AI, and the valuations being paid set the benchmark for every software vendor selling automation to enterprises. For workers, the early evidence from this company is redeployment rather than cuts — for now.

Read the original at TechCrunch →

---

12 MEDIUM IMPACT TechCrunch

Google Search can now read Drive files and Wallet data

> Google said on Monday that AI Mode in Search can now pull from Google Drive, Google Wallet and the Gemini app, extending a feature it calls Personal Intelligence beyond Gmail, Calendar and Photos.

Personal Intelligence is Google's name for answers that draw on what is already in a user's Google account rather than on the open web alone. The feature already connected to Gmail, Google Calendar and Google Photos. Monday's addition brings in three more: Drive, Wallet and the Gemini app.

The pitch is convenience built on data Google already holds. Connect Drive and AI Mode can read documents and spreadsheets to answer questions about a project or a big decision — Google's example is a tailored plan for tax season. Connect Wallet and it can help track financial goals, cut everyday spending and make better use of credit card rewards. Google says AI Mode can only read Wallet information from accounts; it cannot move or manage money. Users can also carry context between the Gemini app and AI Mode, so an idea explored in one does not have to be restarted in the other.

The strategic point is that rivals cannot easily copy this. Google is leaning on services people already use, which gives its assistant a store of personal context that a competitor without an email service, a calendar, a

photo library and a wallet cannot assemble. For anyone selling AI assistants, that is the competitive moat in plain sight.

Connecting the apps is optional, and Personal Intelligence can be switched on or off at any time. Google also said a new "Personalization" icon will soon appear below responses so users can see what context AI Mode drew on — a nod to the obvious worry about an assistant reading financial and work documents. Separately, AI Overviews and AI Mode will start tailoring answers to a user's recent search context. The tailoring is the incremental part; the Wallet and Drive access is the substantive change, and it lands squarely on the privacy question users will ask first.

KEY FACTS

ANNOUNCED
Monday

NEW APP CONNECTIONS
Google Drive, Google Wallet, Gemini app

EXISTING CONNECTIONS
Gmail, Google Calendar, Google Photos

WALLET ACCESS LIMITS
Can read account information; cannot move or manage money

USER CONTROL
Optional; can be turned on or off at any time

COMING SOON
"Personalization" icon showing the context AI Mode used

> **WHY IT MATTERS**
>
> Businesses and workers who keep documents and financial accounts inside Google's ecosystem now have a reason to decide deliberately whether an AI assistant should read them — and a reason to check the setting before it becomes habit. For competitors, it widens a data advantage that is hard to buy.

Read the original at TechCrunch →

Sources

01 Anthropic’s business customers report less hiring, few AI layoffs
https://www.anthropic.com/research/working-with-ai

02 DeepMind’s AI teams improve quantum algorithms, with proofs attached
https://deepmind.google/blog/finding-faster-quantum-algorithms-with-agents-and-formal-verification/

03 Z.ai releases GLM-5.4, a coding model that reuses components
https://simonwillison.net/2026/Nov/23/glm-54/

04 Ai2’s robot tracker benchmark exposes failures in real-world video
https://huggingface.co/blog/allenai/robotrack

05 Claude Ship can approve code changes without a human
https://techcrunch.com/2026/11/23/anthropic-launches-claude-ship-letting-ai-agents-review-fix-and-approve-code-changes/

06 Hugging Face shows how test setup can distort agent scores
https://huggingface.co/blog/agentic-evals

07 AI-assisted Ladybird browser gets its first installable alpha
https://simonwillison.net/2026/Nov/23/ladybird/

08 Datasette’s chat assistant now turns database questions into charts
https://simonwillison.net/2026/Nov/24/datasette-charts/

09 Amazon pays $7.5 billion for Decart’s chip-optimization technology
https://techcrunch.com/2026/11/23/amazon-snaps-up-decarts-ai-optimization-tech-for-7-5b-as-it-prepares-to-sell-trainium-chips/

10 Nvidia reinstates Jensen Huang as chip-smuggling inquiries continue
https://techcrunch.com/2026/11/23/nvidia-reverses-course-and-brings-jensen-huang-back-as-ceo/

11 Sierra doubles its valuation to $30 billion in six months
https://techcrunch.com/2026/11/23/sierra-doubles-valuation-to-30b-within-six-months-as-demand-for-ai-in-cx-takes-off/

12 Google Search can now read Drive files and Wallet data
https://techcrunch.com/2026/11/23/google-searchs-ai-mode-can-now-tap-into-drive-and-wallet-for-more-personalized-responses/

About this document. Every story in the 24 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 2026-11-24T06:33:53Z · model deepseek-v4.1-flash:cloud · New Horizon, new-horizon.tech · To unsubscribe or manage your subscription, see the footer of the daily email.