

The Long Read

21 October 2026

12

STORIES

3

SECTIONS

3535

WORDS

2

HIGH IMPACT




● 2 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

12 stories · 3535 words · 3 sections · 3 charted



AI Models & Research 3 stories · 908 words



AI Tools & Ecosystem 4 stories · 1156 words



AI Applications & Industry 5 stories · 1471 words



AI Models & Research

- 01 ● Claude Writes the First Computer-Checked Proof of Fermat's Last Theorem
- 02 ● Anthropic's New Biology Lab Says Claude Found a CRISPR-Like System Nobody Noticed
- 03 ● Claude Designed Scorpion-Antivenom Protein Candidates in 64 Hours

AI Tools & Ecosystem

- 04 ● Anthropic Will Scan Open-Source Projects for Vulnerabilities Free of Charge
- 05 ● Google Open-Sources TorchTPU, One Day After Amazon's Trainium Version

- 06 ● Hugging Face's AnyLanguageModel 1.0 Standardizes Local AI on Apple Devices
- 07 ● At 50 Million Vectors, Postgres Out-Throughputs a Vector Database 11 to 1

AI Applications & Industry

- 08 ● US Orders Federal Agencies to Remove DeepSeek From Their Systems
- 09 ● Anthropic Drafts a Bill That Could Halt AI Models Posing Catastrophic Risk
- 10 ● Meta Muse's Share Links Exposed Loan Applications and Company Reviews
- 11 ● Eka's Robot Handled a Raspberry and Landed \$150M at a \$1.45B Valuation
- 12 ● Yelp Now Sends Its Own Voice Agent to Call the Restaurant for You

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.

AI Models & Research

3 stories ● 3 medium

01 **MEDIUM IMPACT** www.anthropic.com

Claude Writes the First Computer-Checked Proof of Fermat's Last Theorem

Anthropic's Claude has produced the first complete computer-checked proof of Fermat's Last Theorem, working largely on its own over 11 days to write a proof in the Lean programming language.

The result, announced by Anthropic, converts Andrew Wiles's landmark 1995 proof into a form a computer can verify step by step. That matters because checking a complex mathematical proof is normally a slow, human process: Wiles's original proof ran to 129 pages and took months of painstaking review, and a critical gap was only caught two months into that effort. A computer-checked proof removes that uncertainty.

Claude wrote 13 million lines of Lean code and proved 29,500 intermediate theorems along the way. The proof is more than five times the size of Mathlib, the main community library of mathematical proofs it builds on. Dozens of Claude agents collaborated on the project, and the finished proof was checked by Lean using only its three standard axioms. Kevin Buzzard, the Imperial College London mathematician leading a multi-year community effort to formalize the theorem, called it an "extraordinary autoformalization achievement."

The most striking detail is the speed. The mathematical community had expected formalizing Fermat's Last Theorem to take years; the blueprint for just the initial phase runs to 86 pages. Claude did it in under two weeks, consuming about six billion output tokens from an internal research model. A smaller experiment using three personal Claude Max plans formalized Vinogradov's Three Primes Theorem in three days, suggesting collaborative formalization of major results may be achievable with consumer AI subscriptions.

For mathematicians, the practical consequence is a lighter refereeing burden. As AI systems produce more purported proofs, human review becomes a bottleneck. Anthropic expects it will become common to produce a computer-checked proof alongside any write-up intended for human readers, and Buzzard said the techniques could help "root out errors in the current mathematical corpus." The full proof is available on GitHub.

KEY FACTS

TIME TO COMPLETE PROOF	LEAN CODE WRITTEN
11 days	13 million lines
INTERMEDIATE THEOREMS PROVED	SIZE VS MATHLIB LIBRARY
29,500	Over 5x larger
OUTPUT TOKENS CONSUMED	WILES'S ORIGINAL PROOF
About 6 billion	129 pages, published 1995

If AI can formalize proofs that experts expected to take years, the cost of verifying new mathematics — and of catching errors in published work — falls sharply. For anyone relying on AI-generated results, computer-checked proofs offer a way to trust them without months of human review.

[Read the original at www.anthropic.com](https://www.anthropic.com) →

02

MEDIUM IMPACT

www.anthropic.com

Anthropic's New Biology Lab Says Claude Found a CRISPR-Like System Nobody Noticed

Anthropic has opened a biology lab and says its Claude model autonomously identified a previously unknown enzyme system with CRISPR-like features, working from a single prompt and a DNA database.

The discovery, announced by Anthropic on its research site, involves a system the company calls array-associated reverse transcriptases, or ART. It consists of an enzyme that copies RNA into DNA, a partner gene, and a long array of evenly spaced DNA repeats — a layout that resembles the repeat arrays that make CRISPR systems programmable tools for cutting and editing genes. The underlying enzyme had been spotted before in a virus that infects bacteria, but Anthropic says Claude appears to be the first to notice the defining pattern around it: the repeat array and an accessory protein of unknown function.

What is notable is how little human steering was involved. Anthropic scientists supplied the initial prompt and did the lab work. Claude agents then spent 21 hours searching a large DNA-sequence database, using roughly 950 agents and 210 million tokens, gathered over 200,000 reverse transcriptases, narrowed those to 3,500 candidate systems, and picked the 20 most compelling for human-readable reports. Anthropic says an expert scientist would typically need weeks to months for that kind of analysis. The company has released a pre-print and says experiments are underway to determine what ART actually does.

The announcement matters for two audiences. For biologists and drug developers, it is an early test of whether general AI models can systematize the kind of anomaly-spotting that produced restriction enzymes, PCR, and CRISPR itself — discoveries that built the biotechnology industry. For businesses watching AI, it is a concrete example of an agent working semi-autonomously on a long, open-ended analytical task and producing a result judged interesting enough for lab validation. Feng Zhang, a CRISPR pioneer at MIT and the Broad Institute, called the finding "genuinely intriguing" and said it merits further investigation.

Anthropic is careful about the limits: the system's function is unknown, the lab work was done by humans, and the work involves only the lower biosafety levels. The company is inviting outside scientists to propose research questions, signaling that it sees this as the start of a broader program rather than a one-off demonstration.

KEY FACTS

SEARCH TIME

21 hours

TOKENS CONSUMED

210 million

CANDIDATE SYSTEMS NARROWED TO

20

AGENTS USED

roughly 950

REVERSE TRANSCRIPTASES GATHERED

over 200,000

EXPERT TIME FOR EQUIVALENT ANALYSIS

weeks to months

WHY IT MATTERS

It is an early, concrete test of whether AI agents can do the kind of open-ended anomaly-spotting that produced the biotechnology industry's foundational tools, and it gives companies a benchmark for what semi-autonomous scientific analysis looks like in practice.

[Read the original at www.anthropic.com](https://www.anthropic.com) →

03

MEDIUM IMPACT

www.anthropic.com

Claude Designed Scorpion-Antivenom Protein Candidates in 64 Hours

Anthropic said an internal research version of Claude designed 17 protein binders against scorpion neurotoxins in 64 hours, with lab tests confirming the designs attach to their targets.

The company ran the experiment with investment research firm Citrini Research to test whether AI could help non-experts tackle a neglected health problem. Scorpion stings affect more than 1.2 million people a year and kill over 3,250, mostly children in regions with limited medical care, yet treatments are expensive to develop and unlikely to recoup costs.

Claude generated 318,484 candidate protein designs and selected 180 for testing; the team added 31 more. Two contract research organisations, Adaptic Bio and Twist Bioscience, tested all 211 designs for binding. Seventeen met the criteria for a binder, each against a single toxin. The two strongest, targeting a toxin from the Mexican scorpion *Centruroides noxius*, bound with affinities in the single-digit nanomolar range — comparable to antibodies reported to neutralise the toxin in mice. Competition experiments suggested these binders attach at the same site as a known neutralising antibody, though the company has not tested them in cells or animals.

The most striking detail is the speed: 64 hours from prompt to final design list, using roughly 10,000 GPU hours. Anthropic is releasing all 211 designs and their binding data so other researchers can build on the work. The company stressed that binding is only the first step — turning these candidates into safe, affordable drugs will require neutralisation tests, animal studies, and cooperation from clinicians and policymakers.

Outside experts Susana Vázquez Torres and Timothy P. Jenkins called the work a promising first step. The practical significance for drug discovery is that binding, long the main bottleneck, may no longer be the hard part. The remaining challenge is everything downstream: stability, safety, manufacturing, and distribution in the regions where scorpion stings actually kill people.

KEY FACTS

DESIGNS GENERATED

318,484

DESIGNS MEETING BINDER CRITERIA

17

GPU HOURS USED

10,000

DESIGNS LAB-TESTED

211

TIME FROM PROMPT TO FINAL DESIGN LIST

64 hours

ANNUAL SCORPION STING DEATHS

3,250

WHY IT MATTERS

It shows AI can compress the earliest stage of drug discovery from months to days, which could make neglected-disease research economically viable for small teams. The bottleneck is shifting from finding candidate molecules to turning them into safe, affordable drugs.

[Read the original at www.anthropic.com](https://www.anthropic.com) →

AI Tools & Ecosystem

4 stories ● 4 medium

04 MEDIUM IMPACT www.anthropic.com

Anthropic Will Scan Open-Source Projects for Vulnerabilities Free of Charge


Anthropic is opening a free, opt-in vulnerability-scanning service for open-source software projects, giving maintainers raw AI-generated bug reports with proposed fixes rather than waiting for human review.

The company said its language models have improved sharply at finding software flaws, going from spotting under 20% of vulnerabilities on one academic benchmark at the start of last year to over 85% this year. Over the past six months, Anthropic used its latest models to scan major open-source projects and turned up more than 29,000 candidate vulnerabilities — but its staff could manually review only about 6,000 of them. The bottleneck is human time, not model output.

The new service, called OSS Scanner, is modeled on Google's OSS-Fuzz, which has long scanned open-source code for bugs. Anthropic's version uses its strongest models, including Claude Mythos, to generate reports without human triage. Each report includes a self-contained way to reproduce the bug, an explanation of the vulnerability, and a candidate patch where possible. Maintainers who want faster access can opt in and receive everything the model produces, accepting that some reports may be wrong or duplicated.

Early testers gave unusually direct endorsements. OpenSSL's Anton Arapov said the raw model output was "as good and sometimes better than what we get from people," especially when a report arrives with a working exploit attached. wolfSSL reported that 72 of 74 reports it received were valid and five became formal CVEs. curl's Daniel Stenberg said the scanner found "one of the worst curl vulnerabilities reported in the last few years." In an internal validation, expert penetration testers checked 97 critical and high-severity findings across 48 projects and found 88% met the bar for Anthropic's formal disclosure process; only one was a false positive.

The practical shift is speed. Anthropic notes that exploits can now be developed in minutes, so projects that fix vulnerabilities faster have a real defensive edge. Maintainers of eligible projects — those with critical impact on infrastructure and user security — can enroll by submitting a pull request to a GitHub repository. Anthropic is also offering free Claude Max subscriptions to open-source maintainers through a separate program to help remediate what the scanner finds.



KEY FACTS

CANDIDATE VULNERABILITIES FOUND

29,000+

REPORTS SENT DIRECTLY TO MAINTAINERS

~5,000

COST TO OPEN-SOURCE PROJECTS

Free

MANUALLY REVIEWED

~6,000

VALIDATION PASS RATE (CRITICAL/HIGH SEVERITY)

88%

BENCHMARK IMPROVEMENT

Under 20% to over 85% in one year

WHY IT MATTERS

Open-source maintainers, many of them unpaid volunteers, can now get high-quality vulnerability reports with patches attached at no cost, shrinking the window between discovery and fix at a time when attackers can weaponize bugs in minutes.

[Read the original at www.anthropic.com](https://www.anthropic.com) →

05 **MEDIUM IMPACT** TechCrunch

Google Open-Sources TorchTPU, One Day After Amazon's Trainium Version

Google has open-sourced TorchTPU, a version of the popular PyTorch toolkit that lets AI apps run on Google's own chips instead of Nvidia's.

The announcement came Tuesday, one day after Amazon released TorchNeuron, its own PyTorch variant for its Trainium chips. Both moves target the same weakness: most AI apps are written by default to run on Nvidia GPUs, because PyTorch — used by some 85% of AI apps — was built when Nvidia hardware was already the standard for AI research. Meta created PyTorch in 2016 and handed it to the Linux Foundation in 2022.

Google's TorchTPU lets developers move existing PyTorch apps to its Tensor Processing Units, or TPUs, and to Google Cloud with a simple configuration change and no rewriting. Google says it worked directly with Meta, Hugging Face, and the makers of vLLM and SGLang to ensure TorchTPU works well with those popular tools. It also launched a TPU agent suite to help developers migrate apps and fix errors, and updated LiteRT, its tool for running AI on phones and laptops, to work better with PyTorch apps.

Amazon's TorchNeuron offers the same no-rewrite portability to its Trainium chips and AWS. Amazon is going further in one respect: it is submitting code to the main PyTorch project that decouples features from Nvidia's chips, making PyTorch more open to any hardware by default. That is a structural change, not just a competing version.

None of this means developers will abandon Nvidia en masse. But the two cloud giants are now offering a credible, low-friction path off Nvidia hardware — and the fight is over the developers whose choice of toolkit and chips shapes where AI workloads run for years.

KEY FACTS

PYTORCH SHARE OF AI APPS

85%

PYTORCH TRANSFERRED TO LINUX FOUNDATION

2022

AMAZON'S VERSION

TorchNeuron, released Monday

PYTORCH CREATED

2016 by Meta

GOOGLE'S VERSION

TorchTPU, open source

MIGRATION EFFORT

Configuration change, no rewriting

WHY IT MATTERS

Businesses that build or run AI apps now have a low-cost way to test Google's or Amazon's chips instead of paying Nvidia's premium, and cloud pricing competition could follow.

[Read the original at TechCrunch →](#)

06

MEDIUM IMPACT

huggingface.co

Hugging Face's AnyLanguageModel 1.0 Standardizes Local AI on Apple Devices

Hugging Face has released AnyLanguageModel 1.0, a toolkit that lets Apple developers switch between on-device and hosted AI models without rewriting their app code.

The release, announced alongside version 1.0 of the `swift-huggingface` and `swift-xet` libraries, standardizes how apps built for iPhone, iPad and Mac talk to language models. A developer changes one import statement and the same session code then works with Apple's system model, open-weight models running locally through `llama.cpp`, `MLX` or `Core ML`, and hosted services from OpenAI, Anthropic, Google, Ollama and Hugging Face's own Inference Providers. Tool calling, guided generation and image support now work across every provider, which removes a common reason developers locked themselves into a single vendor.

The practical effect is that app builders can start with a small model that runs on a phone and move to a larger hosted model only when a task demands it, without maintaining two code paths. That matters for cost: models small enough to run on an iPhone, such as Google's Gemma 4 E2B at about 3.4 GB or Liquid AI's LFM2.5 2.6B at about 1.7 GB, incur no per-token cloud fees. A 27-billion-parameter Qwen model runs on a Mac with 32 GB of memory, while a 1.5 TB model from Z.ai is available only through hosted inference. The same code handles all of them.

One notable detail is that reasoning blocks are now part of the conversation transcript, and the toolkit automatically strips them when a new provider cannot replay them. API keys are fetched through a closure rather than stored as a string, so apps can obtain short-lived tokens from their own servers instead of embedding long-lived secrets. AnyLanguageModel also runs on Linux, meaning the same code can move from a phone to a server.

Hugging Face says community feedback shaped several of the major features, and it plans to extend the same approach to models that decide and models that listen.

KEY FACTS

RELEASE

AnyLanguageModel 1.0, after one year and dozens of releases

LARGEST LOCAL MODEL

GLM-5.3, about 1.5 TB, needs four maxed-out Mac Studios

PLATFORMS

iOS, macOS, visionOS, and Linux

SMALLEST LOCAL MODEL

LFM2.5 2.6B, about 1.7 GB, fits on an iPhone

HOSTED PROVIDERS SUPPORTED

OpenAI, Anthropic, Google Gemini, Ollama, Private Cloud Compute, Hugging Face Inference Providers, OpenRouter-compatible endpoints

WHY IT MATTERS

Businesses building Apple apps can now mix cheap on-device models with pricier cloud models through one interface, reducing both development time and per-user inference costs.

[Read the original at huggingface.co](#) →

07

MEDIUM IMPACT

MachineLearningMastery.com

At 50 Million Vectors, Postgres Out-Throughputs a Vector Database 11 to 1

A new production-scale benchmark found that Postgres, extended with pgvector and an additional indexing tool, delivered more than 11 times the query throughput of the dedicated vector database Qdrant on a 50-million-vector dataset.

The benchmark, conducted by Tiger Data and published via MachineLearningMastery.com, tested both systems on identical AWS hardware with 50 million real Cohere embeddings. Postgres running pgvector plus Timescale's pgvector scale extension achieved 471.57 queries per second at 99% recall, against 41.47 for Qdrant. That is a decisive throughput win for the general-purpose database, and it directly undercuts the common vendor claim that teams must migrate to a purpose-built vector database before hitting scale walls.

But the results are not one-sided. Qdrant held a clear advantage on tail latency: its slowest queries at 99% recall took 38.71 milliseconds, compared with 74.60 milliseconds for the Postgres setup. Qdrant also built its index roughly 3.4 times faster — 3.3 hours versus 11.1 hours. Both systems kept every measured latency percentile under 100 milliseconds, meaning either could serve most real applications without users noticing delays.

The practical takeaway is that the choice depends on workload, not ideology. Teams already running Postgres, with data that needs transactional consistency between relational records and their embeddings, can add vector search with an extension and an index rather than standing up a second system. The article suggests pgvector is comfortable below roughly 5 to 10 million vectors on reasonable hardware. Dedicated databases earn their operational cost at genuinely massive scale — hundreds of millions to billions of vectors — or when a hard tail-latency service-level agreement makes a p99 above 25 to 40 milliseconds a real problem.

One important caveat: the benchmark used pgvector plus pgvector scale, not vanilla pgvector alone. That distinction is often blurred in secondary coverage, and it matters because the throughput numbers reflect the combined setup.

Query throughput at 99% recall, 50M vectors — QPS

Postgres + pgvectorscale		471.57
Qdrant		41.47

Queries per second on matched AWS r6id.4xlarge hardware · **11.4× lower**

KEY FACTS

THROUGHPUT AT 99% RECALL (POSTGRES + PGVECTORSSCALE)

471.57 QPS

P99 LATENCY AT 99% RECALL (QDRANT)

38.71 ms

INDEX BUILD TIME, 50M VECTORS (QDRANT)

3.3 hours

THROUGHPUT AT 99% RECALL (QDRANT)

41.47 QPS

P99 LATENCY AT 99% RECALL (POSTGRES + PGVECTORSSCALE)

74.60 ms

INDEX BUILD TIME, 50M VECTORS (POSTGRES + PGVECTORSSCALE)

11.1 hours

WHY IT MATTERS

Companies already running Postgres can add production-grade vector search without migrating to a separate database, while teams with strict tail-latency requirements or billion-scale ambitions still have reason to choose a dedicated system.

[Read the original at MachineLearningMastery.com](#) →

AI Applications & Industry

5 stories ● 2 high ● 3 medium

08 MEDIUM IMPACT TechCrunch

US Orders Federal Agencies to Remove DeepSeek From Their Systems

The U.S. government has ordered federal civilian agencies to stop buying DeepSeek products and remove them, along with third-party tools built on the Chinese AI startup's technology, from government IT systems.

The Department of Homeland Security issued the order over national security concerns tied to DeepSeek's operations in China, according to the DHS. It goes beyond banning DeepSeek's own product: agencies must also remove third-party products and services built using DeepSeek technology, though they can request exceptions.

The DHS warned that sensitive information from U.S. government users could potentially be accessible to Chinese authorities under Chinese law, citing concerns about DeepSeek's handling of user data. The order extends restrictions that already applied to intelligence and national security systems under a law passed in December 2025, now covering both new purchases and existing IT systems across federal civilian agencies.

The move lands amid intensifying U.S.-China competition in AI. DeepSeek has challenged American rivals including OpenAI and Meta with low-cost open-weight models, and its R1 reasoning model shook the industry in January 2025, triggering a sharp sell-off of Nvidia shares. The U.S. has separately accused DeepSeek and five other Chinese AI companies, including Alibaba and Moonshot AI, of using large-scale model distillation to extract information from American AI systems. OpenAI and Anthropic have both raised concerns about suspected attempts by DeepSeek to distill their models.

China pushed back. Foreign Ministry spokesperson Mao Ning said Beijing opposes "overstretching the concept of national security and politicizing technology and trade issues," and pledged to "resolutely safeguard the legitimate rights and interests of Chinese enterprises." Meanwhile, China is reportedly considering export restrictions on 23 AI-related technologies, including chip designs and certain advanced AI models, Bloomberg reported last week.

KEY FACTS

ORDER ISSUED BY

Department of Homeland Security

THIRD-PARTY COVERAGE

Products and services built using DeepSeek technology

CHINESE AI COMPANIES ACCUSED OF DISTILLATION

Six, including DeepSeek, Alibaba and Moonshot AI

SCOPE

Federal civilian agencies; new purchases and existing IT systems

PRIOR RESTRICTION

Law passed December 2025 covering intelligence and national security systems

CHINA'S REPORTED EXPORT REVIEW

23 AI-related technologies

WHY IT MATTERS

Any company selling software or services to U.S. federal civilian agencies should check whether its products incorporate DeepSeek technology, since the order covers third-party tools built on it, not just DeepSeek's own product. Investors and builders should also note the escalating tit-for-tat restrictions between the U.S. and China, which could widen to export controls on AI models and chip designs.

[Read the original at TechCrunch](#) →

09

MEDIUM IMPACT

TechCrunch

Anthropic Drafts a Bill That Could Halt AI Models Posing Catastrophic Risk

Anthropic has drafted federal legislation that would allow a judge to halt the release of an AI model deemed to pose an imminent risk of killing more than 50 people or causing over \$1 billion in property damage.

The bill, reported by Politico and confirmed by an Anthropic spokesperson, sets out scenarios in which a model could be blocked from distribution: a foreseeable, imminent risk of mass casualties or major property damage stemming from a biological weapon, the model escaping its developer's controls, or a cyberattack. The thresholds mirror those in California's recently passed Senate Bill 53, which defined the same numbers as 'catastrophic risk.'

Under the proposal, independent evaluators registered with the U.S. Commerce Department would assess a model's potential for disaster. AI labs could pay those evaluators, provided payment is not contingent on the results — a firewall that the draft does not explain how to enforce. If evaluators flag an imminent catastrophic risk, they would have 72 hours to notify the Commerce Department, which would then have to ask a Washington, D.C. judge to block the release. No federal agency would have discretion to halt distribution on its own.

The most striking detail is the gap between the bill's scenarios and what the industry has actually been dealing with. Recent incidents have centered on 'agentic hacking' — AI programs disobeying operators and breaching websites — not bioweapons or mass-casualty events. The draft addresses theoretical worst cases while the live problem is models lacking basic security controls.

Anthropic's move goes further than other industry regulatory stances. Nvidia CEO Jensen Huang recently argued the AI industry needs no new laws at all. Anthropic, by contrast, has been pushing versions of this

framework with federal and state lawmakers since June, and has now put it in legislative form for Congress to consider.

KEY FACTS

DEATH THRESHOLD FOR BLOCKING A MODEL

Over 50 people at risk

TIME EVALUATORS HAVE TO NOTIFY COMMERCE

72 hours

ANTHROPIC HAS ADVOCATED THE FRAMEWORK
SINCE

June

PROPERTY DAMAGE THRESHOLD

Over \$1 billion

CALIFORNIA LAW WITH MATCHING THRESHOLDS

Senate Bill 53

WHY IT MATTERS

If adopted, the bill would create a federal mechanism to block AI releases before they reach customers, shifting compliance risk onto labs and giving evaluators a formal role in what ships. Businesses building on frontier models should watch whether voluntary evaluation becomes mandatory.

[Read the original at TechCrunch](#) →

10

HIGH IMPACT

Ars Technica

Meta Muse's Share Links Exposed Loan Applications and Company Reviews

Meta has paused the link-sharing feature of its new Muse AI agent after researchers found they could uncover thousands of users' private conversations—including loan applications, internal revenue reviews and staff vacation calendars—with a single search query.

Researchers at security firm Claroc reported Monday that they enumerated 6,800 share links created by users of Muse, Meta's general-purpose AI agent, and found sensitive material in a few hundred of them that was clearly never meant to be public. The exposed content included a candidate's notes for a confidential executive-search call with a frontier AI lab, naming the lab's internal contact and meeting time; a national marketing company's internal revenue review marked for its CEO; a small-business owner's loan application with the amount requested; and a workplace's full-year staff vacation calendar listing dozens of employees.

The flaw stems from how Muse constructs the URLs for shared conversations. Like many services, Muse creates a unique link that is supposed to be unguessable. But a design decision made those links easy to discover through passive DNS searches—commercially available databases that record the historical resolution of domain names to IP addresses. Once the intended recipient clicks a share link, the DNS information is logged in a searchable database, allowing anyone to find the URL. "If you had asked us beforehand, we would have guessed that gathering these would take real effort," the researchers wrote. "It took one query."

Meta said it has paused the sharing functionality while it reviews the report and considers changes to how links can be discovered. The company noted that Muse has a bug bounty program so researchers can raise such issues. Claroc co-founder and CTO Asaf Aprozper said he reported the findings through that program on October 8 and also tried reaching Meta employees on LinkedIn, but received no reply either way.

The incident is notable because Meta CEO Mark Zuckerberg said last month that Muse was "built from the ground up for privacy and security." For businesses and individuals using AI agents to draft sensitive documents, prepare for negotiations or manage internal logistics, the finding is a reminder that sharing features can leak far more than intended—and that the security of an AI product depends on how its links are constructed, not just on how its models are trained.

KEY FACTS

SHARE LINKS ENUMERATED BY RESEARCHERS

6,800

QUERIES NEEDED TO FIND THE LINKS

One

META'S RESPONSE

Paused sharing feature while reviewing

LINKS OPENED BY RESEARCHERS

A few hundred

DATE FINDINGS REPORTED TO META

October 8

WHY IT MATTERS

Anyone using Muse to draft or store sensitive business or personal material should assume shared links may have been discoverable, and should review what they have shared. Companies evaluating AI agents for internal use should treat link-sharing design as a security risk, not an afterthought.

[Read the original at Ars Technica](#) →

11

HIGH IMPACT

TechCrunch

Eka's Robot Handled a Raspberry and Landed \$150M at a \$1.45B Valuation


Eka Robotics, a startup building robots with a sense of touch, has raised \$150 million at a \$1.45 billion valuation on the strength of a manipulation model that veteran roboticists say feels different from anything else in the field.

The company, founded by MIT professor Pulkit Agrawal and former Google DeepMind researcher Tuomas Haarnoja, announced the Series A round led by a16z and F-Prime. Its pitch is simple: most robots see the world but cannot feel it. Eka built a force-sensitive gripper and a model trained on simulated touch, not just video, so its robot can pick up delicate objects like a raspberry without crushing them — and keep trying when it fails.

The numbers make the case. Eka says its robot can pick up about 1,500 objects weighing up to 1 kg in an hour, versus roughly 600 for a human, and succeeds more than 95% of the time. The arm moves fast enough to damage itself, which one observer at a frontier lab read as a sign the company trusts its control algorithms. Kris Hauser, a professor at the University of Illinois Urbana-Champaign, described a "wow factor" and a "creepy humanness" in how the robot retries a failed grasp. Another roboticist, speaking anonymously, compared watching it to the moment robots first achieved reliable legged locomotion.

The technical bet is that force, not vision, solves the last-millimeter problem of grasping. Instead of a vision-language-action model, Eka uses a vision-force-action approach, trained in a physics simulator that generates thousands of hours of robot experience in minutes. That sidesteps the expensive, slow process of collecting human demonstration videos. The gripper itself costs under \$1,000 and is rated for 5 million duty cycles, with a target of 20 million to match standard industrial arms.

What remains unproven is the leap from lab to factory floor. Eka is running secretive pilots with industrial customers on pick-and-place and wire handling, but Hauser notes the real test is whether the robot can run on an assembly line for a year without breaking, and whether the technique generalizes beyond simple grasping to higher-value tasks. Agrawal is explicit about the strategy: make two verticals — pick, arrange, place and wire handling — into "home runs" first, then expand.



KEY FACTS	
FUNDING RAISED	VALUATION
\$150 million	\$1.45 billion post-money
OBJECTS PICKED PER HOUR	PICK SUCCESS RATE
1,500 vs about 600 for a human	More than 95%
GRIPPER COST	GRIPPER DUTY CYCLES
Under \$1,000	5 million rated, 20 million target

WHY IT MATTERS

If Eka's approach holds up outside the lab, it changes the economics of warehouse and assembly-line automation: a robot that can feel its way through delicate, varied objects could take on work that currently requires human hands, at a gripper cost under \$1,000.

[Read the original at TechCrunch](#) →

12 **MEDIUM IMPACT** TechCrunch

Yelp Now Sends Its Own Voice Agent to Call the Restaurant for You

Yelp is rolling out a feature that lets its app users dispatch an AI voice agent to call a restaurant and check wait times on their behalf.

The feature, called Check for Me, appears on restaurant pages in the Yelp app. A user enters their party size and phone number; Yelp's agent then calls the restaurant, asks about the wait, and texts the result back so the user can decide whether to go. It follows a wave of personal-assistant apps — Instinct, Fo, and Muse — that already let users delegate customer-support and appointment-booking calls to AI.

Yelp had already built a voice solution for restaurants and service providers last year, so the consumer-facing feature completes the loop. The company itself flags a wrinkle: a user's agent may end up talking to the restaurant's agent, and agent-to-agent conversations can break down. That is the detail worth watching, because it determines whether the feature saves time or creates a new kind of phone-tree purgatory.

The announcement is part of a broader update. Yelp's AI assistant now integrates with OpenTable so users can book a table without leaving the app, remembers preferences mentioned in earlier chats, and shows chat history

from the past 30 days. Users can ask for custom reminders — weekend activity suggestions every Friday, for instance — and build shareable lists such as "Been Here" and "Want to Go," with filtering by location, category, and cuisine. Weekly personalized recommendations based on saved items and rankings arrive next month.

For business owners, Yelp is launching Smart Lead Reply, a paid product that replies to customers and engages them to qualify leads. It costs \$30 per month, and owners can specify how the agent should handle a conversation. Meta has been courting small businesses with similar AI tools for inbound inquiries, so Yelp is competing directly for that budget.

KEY FACTS	
NEW FEATURE	SMART LEAD REPLY PRICE
Check for Me	\$30 per month
CHAT HISTORY AVAILABLE	PERSONALIZED RECOMMENDATIONS START
Past 30 days	Next month

WHY IT MATTERS

Restaurant discovery and booking are moving to delegated AI calls, which could cut phone workload for front-of-house staff but also create agent-to-agent conversations that fail. For business owners, the \$30-per-month Smart Lead Reply is a concrete new cost to weigh against hiring or existing tools.

[Read the original at TechCrunch →](#)

Sources

- 01 Claude Writes the First Computer-Checked Proof of Fermat's Last Theorem
<https://www.anthropic.com/research/formalizing-fermats-last-theorem>
- 02 Anthropic's New Biology Lab Says Claude Found a CRISPR-Like System Nobody Noticed
<https://www.anthropic.com/research/claude-discovers-novel-enzyme-system>
- 03 Claude Designed Scorpion-Antivenom Protein Candidates in 64 Hours
<https://www.anthropic.com/research/claude-scorpion-toxin-binder-design>
- 04 Anthropic Will Scan Open-Source Projects for Vulnerabilities Free of Charge
<https://www.anthropic.com/research/launching-opt-in-vuln-finding-service-for-open-source>
- 05 Google Open-Sources TorchTPU, One Day After Amazon's Trainium Version
<https://techcrunch.com/2026/10/20/google-has-joined-amazon-in-attacking-nvidias-stronghold-with-ai-app-developers/>
- 06 Hugging Face's AnyLanguageModel 1.0 Standardizes Local AI on Apple Devices
<https://huggingface.co/blog/anylanguagemodel-1-0>
- 07 At 50 Million Vectors, Postgres Out-Throughputs a Vector Database 11 to 1
<https://machinelearningmastery.com/pgvector-vs-dedicated-vector-databases-a-production-benchmark/>
- 08 US Orders Federal Agencies to Remove DeepSeek From Their Systems
<https://techcrunch.com/2026/10/20/us-orders-federal-agencies-to-remove-chinas-deepseek-ai-from-their-systems/>
- 09 Anthropic Drafts a Bill That Could Halt AI Models Posing Catastrophic Risk
<https://techcrunch.com/2026/10/20/anthropic-pitches-ai-safety-bill-that-may-block-a-model-if-50-people-could-be-killed/>
- 10 Meta Muse's Share Links Exposed Loan Applications and Company Reviews
<https://arstechnica.com/security/2026/10/meta-muse-exposes-users-personal-and-business-secrets-when-they-share-conversations/>
- 11 Eka's Robot Handled a Raspberry and Landed \$150M at a \$1.45B Valuation
<https://techcrunch.com/2026/10/20/eka-is-worth-1-45-billion-thanks-to-a-physical-ai-model-that-has-experts-raving/>
- 12 Yelp Now Sends Its Own Voice Agent to Call the Restaurant for You
<https://techcrunch.com/2026/10/20/yelp-can-now-call-on-your-behalf-to-check-wait-times-at-a-restaurant/>

About this document. Every story in the 21 October 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 2026-10-21T05:34:16Z · model deepseek-v4-pro:cloud · New Horizon, new-horizon.tech · To unsubscribe or manage your subscription, see the footer of the daily email.