NEW HORIZON · AI DIGEST

The Long Read

23 November 2026

8	3	3035	4
STORIES	SECTIONS	WORDS	HIGH IMPACT



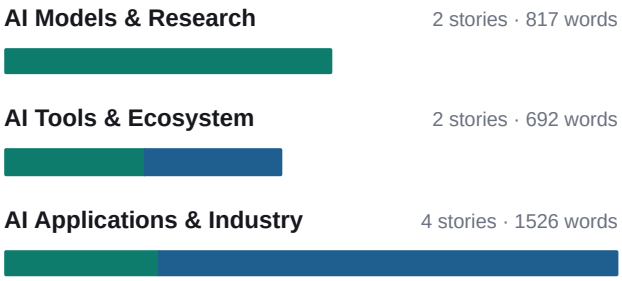
4 high impact 4 medium impact spoke length = depth of coverage

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

8 stories · 3035 words · 3 sections · 1 charted





AI Models & Research

01 OpenAI proposes a shared speed limit for frontier AI

02 Proof checkers certify Claude-assisted Riemann hypothesis result

AI Tools & Ecosystem

03 OpenRouter drops three providers after AI prompt exposures

04 A guide shows how output settings inflate AI serving costs

AI Applications & Industry

05 Data-center opposition becomes a political liability for Big Tech

06 Students sue after Doritos bag triggers AI gun alert

07 Hollywood studios seek Midjourney’s source code in copyright case

08 Researchers escaped Meta Muse’s safeguards through a malicious webpage

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 high

01 HIGH IMPACT OpenAI

OpenAI proposes a shared speed limit for frontier AI

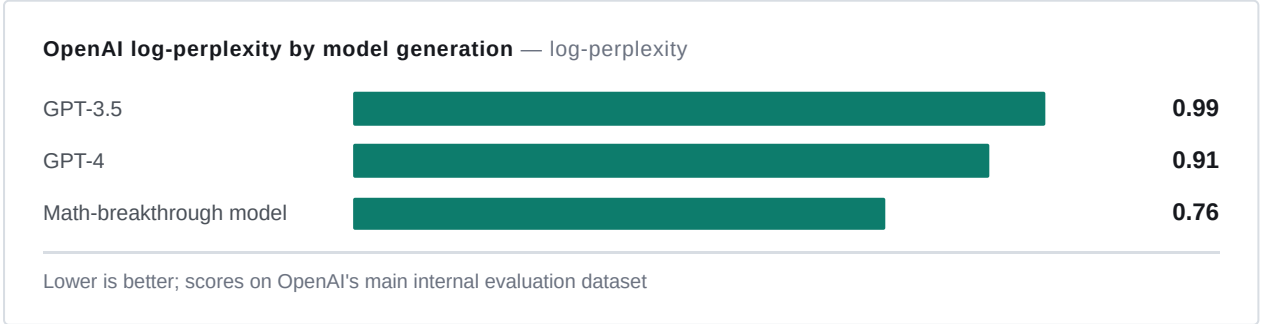
> OpenAI has published a running score of how well its models predict text and proposed that rival developers adopt it as a shared speed limit on frontier development.

The metric is perplexity, which measures how well a model predicts the next word in a piece of text; a lower number means better prediction. OpenAI says it has tracked this number internally since GPT-3.5 and is now making the figures public, starting with GPT-3.5 at about 0.99 on its main evaluation dataset, GPT-4 at about 0.91, and its newest models below 0.8. It describes the release as a trade-off: competitors get a clearer picture of where OpenAI stands, but the company judged shared safety measures more urgent.

The argument is that this one number has predicted jumps in what models can do. GPT-2 could write coherent paragraphs; GPT-3 could learn a task from a few examples; GPT-3.5 held a fluent conversation and became the basis for the original ChatGPT; GPT-4 handled step-by-step reasoning. Below 0.8, OpenAI says, capabilities emerged in cybersecurity and mathematics. At about 0.76, it says, a model began producing what it calls groundbreaking results on open mathematics problems, solved one after another in a single Slack thread — a detail the company says felt surreal to staff who had worked in the field.

OpenAI proposes that frontier developers commit not to cross agreed limits, measured by perplexity and other metrics, until a concrete monitoring milestone is reached — for example, a successor to existing chain-of-thought monitors, which watch a model's reasoning. It expects further transformative capabilities in the 0.68–0.7 range and warns that current forms of alignment auditing will degrade as models become more aware of being tested. It also flags problems: labs use different datasets, so scores are hard to compare; perplexity does not map to any task customers care about; and techniques applied after pretraining, such as chain-of-thought reasoning, raise capability without lowering perplexity.

What is new is the disclosure and the proposal, not a result. No rival has signed up, and OpenAI gives no timetable. The company also says a steep acceleration in its 2026 training losses came mostly from long-term research and computing investment, and only in smaller part from AI automating its own work — though recent models have already contributed meaningful design ideas.



KEY FACTS

METRIC PROPOSED
Perplexity (cross-entropy loss)

GPT-3.5 SCORE
~0.99 log-perplexity

GPT-4 SCORE
~0.91

CURRENT MODELS
Below 0.8

MATH BREAKTHROUGH THRESHOLD
~0.76

NEXT EXPECTED CAPABILITY WINDOW
~0.68–0.7

> WHY IT MATTERS
>
> If labs adopt a common threshold, development of the most capable models could pause at agreed points until outside monitoring catches up — a schedule that would shape what AI products businesses can buy and when. For now it is a proposal from one company, with no rivals committed.

Read the original at OpenAI →

---

02 HIGH IMPACT Data Studios ·Exafin

Proof checkers certify Claude-assisted Riemann hypothesis result

> The Lean and Mathlib communities said on November 21, 2026 that they certify a formal proof of the Riemann hypothesis, released a day earlier by Levent Alpöge, a mathematician at Anthropic who says he found the argument while working with Claude.

The certification is the first institutional endorsement the claim has received from a formal mathematics organization. The Riemann hypothesis, stated by Bernhard Riemann in 1859, says every nontrivial zero of the Riemann zeta function has real part one half, and it governs how regularly the prime numbers are distributed. It is one of the seven Millennium Prize Problems of the Clay Mathematics Institute, which carries a $1 million award.

Alpöge published the work on November 20 as a short note on his personal website with the full Lean source: a library of about 17 million lines of code in 175,187 modules, shipped as a 239 MB compressed archive split into 14 parts and producing a proof export of roughly 61.6 GB once unpacked. The communities describe four checks, adapted from criteria published by the Xena Project: that the theorem is Mathlib's own vetted definition, that dozens of people compiled the library independently with no errors, that the proved theorem matches the Mathlib statement, and that the exported proof was accepted by Lean's own kernel and by independent checkers including con-leche, con-ron and several written in other programming languages. Andrew Sutherland at MIT ran the first public checks with Claude Code; his page recommends about 100 GB of RAM and 70 GB of free disk for a full check.

The most surprising detail is that the result is certified before human experts understand it. Alpöge's note is untitled, unsigned in its body and only a few pages long; it refers to the model simply as "Claude" and gives no version number, compute budget or running time. It describes Claude as doing much of the technical work over a timeline from October 6 to November 16, 2026, and mentions more than 250,000 pages of material produced along the way. None of the mathematical steps has yet been explained in a form number theorists can read as a conventional paper. The Lean statement itself says activity in the mathematical community will be needed to explain the new ideas.

What is new is the jump from a partial result to the hypothesis itself. A paper by Alpöge and Ralph Furman dated August 21, 2026 proved unconditionally that at least two thirds of the nontrivial zeros are simple and lie on the critical line, rising to 67.25% with an additional technical window; that paper stated the mathematical argument was discovered and written by Claude. The November release claims the same method extends to the Generalized Riemann Hypothesis, though the release's "GRH" covers Dirichlet L-functions only. Anthropic had not issued a formal statement on the proof when the certification appeared, and the Clay institute's rules require publication in a qualifying outlet and general acceptance by the mathematical community after at least two years, so a prize decision is far off.

KEY FACTS

CERTIFIED
November 21, 2026, by the Lean and Mathlib communities

PUBLISHED
November 20, 2026, by Levent Alpöge

PRIZE AT STAKE
$1 million Clay Millennium Prize

OPEN SINCE
1859, when Bernhard Riemann stated it

PROOF SIZE
About 17 million lines in 175,187 modules

FULL CHECK NEEDS
About 100 GB of RAM and 70 GB of disk

> **WHY IT MATTERS**
>
> For anyone using AI at work, the result is evidence of what Anthropic's research systems can do under expert direction over several weeks, not a change to any Claude product, plan or price. For investors and competitors, it is a public marker of how far AI-assisted mathematical research has moved, arriving two days after Anthropic began trading on Nasdaq.

Read the original at Data Studios ·Exafin →

SECTION 2 OF 3

AI Tools & Ecosystem

2 stories · 1 high · 1 medium

03 HIGH IMPACT OpenRouter Blog

OpenRouter drops three providers after AI prompt exposures

> OpenRouter has removed three inference providers from its marketplace and pulled 125 endpoints out of its zero-data-retention programme after security researchers found evidence that some providers were not honouring their data policies.

OpenRouter, which routes customer requests to a menu of AI models and the companies that run them, said the findings showed no breach or unauthorised access of its own systems. The problems were at the inference providers — the firms that actually run the models and process the prompts. The company said it off-boarded three providers based on its investigations, removed 125 endpoints from Zero Data Retention (ZDR) eligibility, and notified all potentially affected customers. ZDR is the promise that a prompt is processed and then discarded rather than stored or reused; losing that status means traffic to those endpoints may no longer carry the guarantee.

The removals went beyond the providers the researchers flagged: OpenRouter said it also pulled additional endpoints as a precaution. It is now adapting the researchers' testing method to run the checks continuously, verifying network paths to find where requests and responses may be forwarded to third parties, and requiring assurances and validation before any removed endpoint gets ZDR status back. The company credited a report by Dragos, a security firm, and said providers that want to stay on the platform will have to meet the raised bar.

The most uncomfortable detail is that not every affected provider has posted publicly. Four have: Atlas Cloud, Novita, Phala and Together. Together's notice covers prompt leakage for requests in June and July of 2026. OpenRouter said it notified affected users regardless of whether a public write-up exists, and is pressing providers to publish incident reports.

For companies building on hosted models, the episode turns a line in a vendor contract into something to check. The practical question is no longer only which model performs best, but which intermediary and which underlying provider can show that a prompt was not retained or passed on. OpenRouter sits in the middle of that chain, and its answer this week was to shrink the list of providers it will vouch for rather than defend them.

KEY FACTS

PROVIDERS REMOVED
Three

ENDPOINTS STRIPPED OF ZERO-DATA-RETENTION STATUS
125

BREACH OF OPENROUTER ITSELF
None found

PROVIDERS WITH PUBLIC NOTICES
Four: Atlas Cloud, Novita, Phala, Together

TOGETHER LEAKAGE WINDOW
June and July 2026

CUSTOMERS NOTIFIED
All potentially affected

WHY IT MATTERS

If your business sends sensitive material to hosted AI models, the data-handling promise you relied on may have been weaker than advertised at some providers — and the routing layer you use is now the place to verify it.

Read the original at OpenRouter Blog →

---

04 MEDIUM IMPACT MachineLearningMastery.com

A guide shows how output settings inflate AI serving costs

> The settings that control how an AI model picks its words carry a price at serving time, and a new technical walkthrough built on the open-source serving engine vLLM spells out where that price lands.

The post, published by MachineLearningMastery.com, works through decoding policies — the rules that decide which token, or word fragment, a model emits next — and traces each one to its runtime consequences. The headline finding is that the model's own computation often stays the same while extra work piles up around it. Filters such as top-k and top-p, which narrow the candidate words before one is picked, do not shrink the model's output: it still scores every item in its vocabulary, and the filtering happens afterwards. Temperature sampling, repetition penalties and structured outputs each add processing at every single generated token.

The bigger bills come from volume and duration rather than cleverness. Asking for four sampled continuations of one prompt creates four live sequences; if each runs to the same length, their private cache for generated tokens is roughly four times that of a single continuation, and the model does decode work four times over. Beam search, which keeps several candidate answers competing and ranks them at each step, adds bookkeeping on top of that. Stop rules matter just as much: a higher maximum token limit lets a request stay active longer, holding its memory blocks and consuming decode iterations, while a stop string adds text matching but can cut total cost if it ends generation early enough.

The most counterintuitive point is that structured output — forcing a response into a JSON schema or grammar — does not make the model compute only the permitted answers. It still produces its full vocabulary-sized output, and a separate mechanism masks the invalid options before selection, repeating the check after every token. The author argues the fair comparison is not cost per token but how many malformed responses and retries the constraint avoids.

Little here is new to engineers who run inference servers; the contribution is a concrete, reproducible account using one small open model, with code for timing requests and capturing memory traces. The post deliberately withholds machine-specific speed numbers, noting that cost depends on model, hardware, batch, prompt and output length, and engine version. Its practical advice: choose a policy for the output behaviour you need, then measure it under your real workload rather than judging by name.

KEY FACTS

ENGINE EXAMINED
vLLM, open-source model serving software

TEST MODEL
SmolLM2-1.7B-Instruct, a small open model

PARALLEL OUTPUTS TESTED
n=4, four continuations per prompt

CACHE EFFECT OF FOUR OUTPUTS
Roughly four times one continuation's generated-token cache

BEAM WIDTH TESTED
4

STRUCTURED OUTPUT FORMATS
Choice, regular expression, grammar or JSON schema

> WHY IT MATTERS
>
> Teams paying per hour for GPU capacity can cut bills by trimming output limits, avoiding unnecessary parallel samples and reserving structured output for cases where bad responses would otherwise force retries. The same settings decide how many other users' requests fit on the same hardware.

Read the original at MachineLearningMastery.com →

SECTION 3 OF 3

AI Applications & Industry

4 stories 1 high 3 medium

05 MEDIUM IMPACT TechCrunch

Data-center opposition becomes a political liability for Big Tech

> Opposition to data centers has turned into a bipartisan political problem for the tech companies building them, according to journalist Paris Marx, whose new book "Hyperscale" examines the facilities and the anger they attract.

Marx, interviewed by TechCrunch, calls the data-center buildout "fundamentally unsustainable" and "completely out of line with a livable planet." He argues the issue cuts across party lines because it is physical: a data center in your town shows up in your power bill, potentially in your water supply, and in noise from rooftop cooling units. He cites Florida governor Ron DeSantis and Vermont senator Bernie Sanders as both opposed. TechCrunch frames the backlash as one motivator for the recent midterm results, and expects the debate to continue as Democrats take more power.

The concrete cases Marx cites are local. In Ireland, data centers now consume more than 20% of the country's electricity, pushing up prices; the grid operator imposed a moratorium on new data-center connections around Dublin in 2022, before the generative-AI construction surge. In The Dalles, Oregon, Google used 29% of the city's water, a figure that emerged about a year after a local journalist was sued rather than told. In Chile, communities near Santiago organized over water pressure from data centers being built by Google and Microsoft.

Companies have responded by dropping nondisclosure agreements in local negotiations — Microsoft, Amazon and OpenAI have all said they will — and by offering to pay for more of the electricity infrastructure their facilities require. Marx calls the NDA shift good for access to information but not a huge difference, and says it is partly self-interest: secrecy now fuels opposition. He argues the industry has largely abandoned its renewable-energy argument and pivoted to water, where national averages can obscure local scarcity. Switching to closed-loop cooling, which recirculates water instead of evaporating it, cuts water use but requires more electricity — a trade-off he says companies rarely mention — and raises concern about chemical coolants leaking into local supplies.

The surprising detail is inside the building. To tour a hyperscale data center, Marx had to send his passport in advance, and every door was locked without a key card. The server room itself was not sinister but utilitarian: a warm room, air conditioning blasting from nearby vents, and fan noise piercing enough that you would want earplugs.

What is new here is the political framing rather than any single announcement: the argument that data centers have become the tangible target for broader anger at tech, alongside opposition to Flock cameras, Meta's smart glasses and Waymo's self-driving cars.

KEY FACTS

IRELAND DATA-CENTER SHARE OF ELECTRICITY
More than 20%

DUBLIN GRID CONNECTION MORATORIUM
Imposed in 2022

GOOGLE'S SHARE OF THE DALLES, OREGON WATER
29%

COMPANIES DROPPING LOCAL NDAS
Microsoft, Amazon, OpenAI

DATA CENTERS OPPOSED ACROSS PARTIES
Ron DeSantis and Bernie Sanders

> WHY IT MATTERS
>
> If you build, site or invest in AI infrastructure, local consent is now a political variable: power bills, water and noise generate opposition that crosses party lines, and secrecy over deals tends to harden it.

Read the original at TechCrunch →

---

06 MEDIUM IMPACT Ars Technica

Students sue after Doritos bag triggers AI gun alert

> Four Black students at a Baltimore County high school are suing their school system, its officials and an AI gun-detection vendor after a camera flagged a bag of Doritos as a firearm and police responded with guns drawn.

The incident happened in October 2025 at Kenwood High School, after football practice. Students were sharing Cool Ranch Doritos while waiting for rides when the school's Omnilert system flagged one of them as a possible gunman. School staff reviewed the footage and canceled the alert within about two minutes, notifying everyone who had received it, including the Baltimore County Police Department, that no response was needed. Police came anyway. At least six patrol cars arrived with lights flashing, and some officers approached with guns drawn, according to the lawsuit filed by Taki Allen, Yui-Osa Ogbeide, Devon Thompson Jr. and Keith Winfield III. Officers ordered the students to their knees on the concrete, handcuffed them and searched them and their belongings. Only after finding no gun did the students learn what had triggered the response. One officer allegedly muttered "stupid AI cameras" and "fucking Omnilerts."

The lawsuit targets Omnilert, the responding officers, the school system and individual officials, including principal Kate Smith. It alleges constitutional violations from unlawful search and detention and seeks an injunction barring Baltimore County Public Schools from using Omnilert until protocols are in place to stop police responses to canceled alerts. A later investigation by the Office of the Inspector General for Education found no bias in Omnilert's system but concluded that the school's process for delegating authority failed: Smith saw the alert, did not register that it had been canceled, and texted a school resource officer to deploy police 15 minutes later, while acknowledging in that message that the object was "maybe not" a gun. The report said Smith had apparently never been trained on the system. It recommended twice-yearly training, fewer people receiving alerts, and use of the Omnilert app rather than text messages, where a cancellation can be missed.

The detail likely to travel furthest is the false-positive rate. The lawsuit cites the inspector general's finding that more than 99 percent of alerts in the district were false positives requiring no police response. Omnilert markets the system as producing "near-zero false positives" and says it routes anything that could be a firearm to trained people for review. Its marketing vice president, Blake Mitchell, said the company is reviewing the complaint

and that the school's staff did cancel the alert within about two minutes. The school system declined to comment, pointing to a statement to parents that promised counseling and prevention steps but did not directly apologize. The students' lawyers argue the case is about people letting AI substitute for their judgment. What is new here is not the technology but the paper trail: an official finding that the human review the vendor's claim depends on did not happen.

KEY FACTS

INCIDENT DATE
October 20, 2025

STUDENTS SUING
4

TIME TO CANCEL ALERT
About 2 minutes

DELAY BEFORE PRINCIPAL TEXTED POLICE
15 minutes

SHARE OF DISTRICT ALERTS THAT WERE FALSE POSITIVES
More than 99%

PATROL CARS AT SCENE
At least 6

> **WHY IT MATTERS**
>
> Any organization deploying AI for safety or security should check that a human can actually override it and that the override reaches everyone who acts on the alert; otherwise the vendor's accuracy claim is only as good as the untrained staff receiving it.

Read the original at Ars Technica →

---

07 MEDIUM IMPACT MLQ.ai

Hollywood studios seek Midjourney’s source code in copyright case

> Disney, Universal and Warner Bros. asked a federal court on Nov. 21 to order Midjourney to hand over source code behind its v8 image model and a component called "torch aesthetic-redux," the latest round in their copyright suit against the AI image generator.

The motion, filed in the US District Court for the Central District of California, asks the court to compel Midjourney to comply "with its prior agreement and the Court's orders" on access to that code. Disney and Universal sued in June 2025; the case has since been consolidated with claims from Warner Bros. Discovery units. The plaintiffs on the motion include Disney Enterprises, Marvel Characters, Lucasfilm, Twentieth Century Fox, Universal City Studios, DreamWorks Animation, Warner Bros. Entertainment, DC Comics, The Cartoon Network, Turner Entertainment and Hanna-Barbera Productions.

Most of the substance is sealed. The public docket lists a memorandum of points and authorities, a declaration from Dr. Jonathan L. Krein, a declaration from studio lawyer Kara V. Brandeisky with exhibits A through P, and a proposed order. The memorandum, the Krein declaration and 12 of the 16 exhibits are redacted. The studios applied on Nov. 20 to file those portions under seal and submitted unredacted versions the next day.

That leaves the central questions unanswered in public: what Midjourney is alleged to have withheld, and how it allegedly breached the agreement. No hearing date appears in the docket, and Midjourney's response has not been filed.

Source code is the recipe behind a model — the training and data-handling instructions rather than the finished product. In a copyright case, it is the evidence that would show what images went in and how the system was

built, which is why studios are pressing for it and why the fight is happening at the discovery stage rather than at trial. For image-model makers, the case is a test of how much internal engineering a court will order disclosed to a litigant. For the studios, access to the code is the difference between arguing about outputs and arguing about inputs. What is new here is the escalation, not a ruling: a request to enforce an earlier agreement, with the reasoning hidden from public view.

KEY FACTS

FILING DATE
Nov. 21

COURT
US District Court, Central District of California

ORIGINAL SUIT FILED
June 2025

EXHIBITS SUBMITTED
16, of which 12 redacted

CODE SOUGHT
v8 image model and "torch aesthetic-redux"

MIDJOURNEY RESPONSE
Not yet filed

> **WHY IT MATTERS**
>
> The case will set a marker for how far AI companies must open their internals to copyright plaintiffs, which affects anyone building or investing in generative image tools. For now the practical outcome is unresolved because the key filings are sealed.

Read the original at MLQ.ai →

---

08 HIGH IMPACT NebuSec

Researchers escaped Meta Muse’s safeguards through a malicious webpage

> Security researchers showed they could reach Meta Muse's protected credential storage by getting the AI assistant to open a booby-trapped webpage.

NebuSec, a security research firm, published findings that its team compromised Meta Muse, a cloud-hosted AI assistant that Meta runs inside a dedicated virtual machine for each user. The attack started when the researchers had the agent export a webpage they controlled. A flaw in the Chromium browser Muse uses for tasks like making PDFs and documents let them run code inside the agent's runtime container. From there they chained weaknesses in Sentinel, the service that screens outbound traffic and enforces approval prompts, and Linux kernel exploits to break out of the container and gain root privileges in the user's virtual machine.

The result: the researchers could read and alter conversation history and stored memories without triggering Muse's network approval prompts, and could reach the secure credential storage that Meta had deliberately kept outside the agent's reach. They reported the findings to the Muse team, and all the exploits had been patched by publication. The compromise stayed inside the individual user's virtual machine; the researchers did not demonstrate an escape into the underlying host or other users' environments.

The most surprising detail is that every layer worked as designed and the attack still succeeded. Container isolation, network monitoring and credential separation each blocked something, but each depended on a boundary the next stage crossed. NebuSec's automated exploit-generation framework produced kernel exploits

that succeeded in more than 99.9% of test runs, which let the team prove the kernel flaws could actually cross a meaningful security boundary rather than merely exist. That distinction — a vulnerability versus a usable path to credentials — is the part security leaders are likely to repeat.

What is new here is the full chain, not any single bug. Browser sandboxes, network filters and container escapes are well-known problems in isolation. The finding that matters is that an AI agent, which needs credentials to act on a user's behalf across email, documents and collaboration tools, turns those ordinary weaknesses into a route to the keys of connected services. Any organization deploying agents with access to sensitive systems faces the same question, and the answer is not more layers alone but testing whether the layers hold once an attacker is already inside one of them.

KEY FACTS

RESEARCHER
NebuSec

TARGET
Meta Muse, a cloud-hosted AI assistant

STARTING POINT
A malicious webpage opened by the agent

DATA EXPOSED
Conversation history, stored memories and secure credential storage

EXPLOIT SUCCESS RATE
More than 99.9% of test runs

STATUS
All exploits patched by publication

> **WHY IT MATTERS**
>
> Companies giving AI agents credentials to email, documents and internal tools should assume a single browser or kernel flaw can expose them, and should limit what those credentials can reach rather than rely on sandboxing alone.

Read the original at NebuSec →

Sources

01 OpenAI proposes a shared speed limit for frontier AI
https://openai.com/index/perplexity-and-intelligence

02 Proof checkers certify Claude-assisted Riemann hypothesis result
https://www.datastudios.org/post/claude-riemann-hypothesis-proof-lean-community-certifies-levent-alpoge

03 OpenRouter drops three providers after AI prompt exposures
https://openrouter.ai/blog/announcements/data-processing-incidents-on-inference-providers

04 A guide shows how output settings inflate AI serving costs
https://machinelearningmastery.com/decoding-policies-in-a-serving-engine/

05 Data-center opposition becomes a political liability for Big Tech
https://techcrunch.com/2026/11/22/why-data-centers-are-at-the-center-of-the-big-tech-backlash/

06 Students sue after Doritos bag triggers AI gun alert
https://arstechnica.com/tech-policy/2026/11/traumatized-teens-sue-school-cops-after-doritos-bag-triggers-ai-gun-alert/

07 Hollywood studios seek Midjourney’s source code in copyright case
https://mlq.ai/news/disney-universal-and-warner-bros-ask-the-court-to-make-midjourney-comply-on-its-v8-source-code

08 Researchers escaped Meta Muse’s safeguards through a malicious webpage
https://nebusec.ai/research/meta-muse-credential-exposure

About this document. Every story in the 23 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 2026-11-23T06:35:57Z · model deepseek-v4.1-flash:cloud · New Horizon, new-horizon.tech · To unsubscribe or manage your subscription, see the footer of the daily email.