

The Long Read

18 October 2026

11

STORIES

3


SECTIONS

2370

WORDS

2

HIGH IMPACT




● 2 high impact ● 9 medium impact spoke length = depth of coverage

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

The issue at a glance

11 stories · 2370 words · 3 sections · 2 charted



AI Models & Research 3 stories · 855 words



AI Tools & Ecosystem 3 stories · 535 words



AI Applications & Industry 5 stories · 980 words



AI Models & Research

- 01 ● MLExpert Is a Generalist Agent Built to Run Machine-Learning Engineering End to End
- 02 ● FDE-Gym Benchmarks AI Agents on Real Customer Delivery, Not Tidy Tickets
- 03 ● Researchers Map the Routing Bottleneck Inside Modern Mixture-of-Experts Models

AI Tools & Ecosystem

- 04 ● Meta Quietly Cuts the Terms That Would Have Let It Scrape Publisher Sites for AI Training
- 05 ● China Puts Model Weights, Agents and Chip Designs on Its Export-Control List

- 06 ● Starlink's Call Centers: 20,000 Grok Voice Calls a Day, 70% With No Human

AI Applications & Industry

- 07 ● FTC Subpoenas OpenAI and Anthropic Over Safety Incidents, Plans Executive Testimony
- 08 ● Schwab and JPMorgan Are Blocking AI Agents From Moving Client Money
- 09 ● CoreWeave's 2032 Bonds Sink Below the Distress Line — and Musk Says 'Yikes'
- 10 ● Anthropic's Own Numbers: Junior Technical Hires Fell From 27% to 9%
- 11 ● 18 Months and an \$8M Fine for the Man Who Faked Billions of AI Music Streams

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 practitioner 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 ● 2 high ● 1 medium

01 **HIGH IMPACT** [arXiv.org](#)

MLExpert Is a Generalist Agent Built to Run Machine-Learning Engineering End to End

A 35B Mixture-of-Experts model trained with windowed interaction RL now matches frontier-scale models on autonomous ML engineering benchmarks.

The paper introduces MLExpert, a 35B Mixture-of-Experts agentic model for autonomous machine learning engineering, trained through a post-training method designed for long-horizon MLE optimization. The authors argue that feedback in MLE runs is costly and confounded: each experiment result reflects several preceding decisions, and a full run contains many such experiments. Their response is two-stage. Hierarchical agentic SFT transfers general tool-use and coding capabilities to MLE through outcome-guided supervision. Windowed Interaction RL (WIRL) then treats the multi-turn decisions between successive experiment results as an execution window and assigns it a shared outcome-based reward. Asynchronous collection and policy correction allow policy updates before a trajectory ends, addressing the credit-assignment problem across long interaction sequences.

To train the model, the authors built the Kaggle-MLE Factory, which constructs executable tasks and provides an interaction harness for model-directed experimentation. On MLE-Bench Lite, MLExpert improves the average medal rate of its base model from 37.9% to 72.7%. The paper states this is comparable to larger frontier models such as Kimi-K3 and GLM-5.3. Beyond the primary benchmark, MLExpert also improves performance on NatureBench Lite, agent scaffold optimization, and algorithm optimization, which the authors interpret as evidence that the learned optimization capability generalizes beyond the training distribution.

The contribution is less a new architecture than a training recipe for agentic MLE. The model size is modest relative to frontier systems, and the gains come from how supervision is selected from long trajectories and how intermediate outcomes are converted into rewards. WIRL is the central novelty: rather than rewarding a full run or individual actions, it rewards execution windows bounded by experiment results, which aligns reward granularity with the natural feedback cadence of ML experimentation.



KEY FACTS

MODEL SIZE

35B Mixture-of-Experts

MLE-BENCH LITE MEDAL RATE (MLEXPERT)

72.7%

TRAINING METHOD

Hierarchical agentic SFT + Windowed Interaction RL

MLE-BENCH LITE MEDAL RATE (BASE)

37.9%

COMPARABLE FRONTIER MODELS

Kimi-K3, GLM-5.3

WHY IT MATTERS

If windowed interaction RL generalizes as the paper suggests, teams can train smaller agentic models that approach frontier performance on autonomous ML engineering at a fraction of the compute cost. The Kaggle-MLE Factory also provides a reproducible task-construction pipeline for building MLE training environments.

[Read the original at arXiv.org](#) →

02

HIGH IMPACT

[arXiv.org](#)

FDE-Gym Benchmarks AI Agents on Real Customer Delivery, Not Tidy Tickets

A new benchmark shows the best AI agent configurations resolve only 14 of 95 realistic customer delivery tasks, with unstated-constraint discovery as the primary bottleneck.

FDE-Gym evaluates agents on end-to-end customer delivery rather than well-specified tickets. The benchmark comprises 95 delivery tasks drawn from 20 customer scenarios across 18 sectors, each adapted by industry experts from real-world delivery projects. Agents operate in simulated customer environments with scripted stakeholders and modifiable, production-like customer systems. Scoring is normalized between 0 (unchanged system) and 1 (reference solution), based on how much the agent improves the customer's own success metric, with penalties for stakeholder interactions including messages, trial runs, and waiting time. An online platform lets humans practice the same tasks under identical grading criteria.

Thirteen model configurations from seven providers were evaluated. In the realistic setting, where unstated constraints must be discovered without explicit guidance, the best configuration achieves a mean FDE-Gym score of 0.56 and resolves only 14 of the 95 tasks, where resolution requires a score of at least 0.8 with no incident. The authors identify discovering unstated constraints as the primary bottleneck: performance drops sharply when models must uncover them without explicit guidance.

Information discovery alone does not explain the gap. On the 23 tasks where the best configuration obtains every key fact from stakeholders, it reaches a normalized business uplift of 0.8 or above on only 10. This indicates that even with complete information, agents still fail to translate requirements into working solutions at a majority of tasks.

The benchmark's contribution is the shift from issue-report-style evaluation to delivery under ambiguity, live-system constraints, and stakeholder interaction. Existing benchmarks typically evaluate agents on well-

specified tasks defined by issue reports, specifications, or task tickets; FDE-Gym measures something closer to the work of forward-deployed engineers.

KEY FACTS

TASKS

95

SECTORS

18

PROVIDERS

7

TASKS RESOLVED BY BEST CONFIGURATION

14 of 95

TASKS REACHING 0.8 UPLIFT WITH ALL FACTS

10

CUSTOMER SCENARIOS

20

MODEL CONFIGURATIONS EVALUATED

13

BEST MEAN FDE-GYM SCORE

0.56

TASKS WITH ALL KEY FACTS OBTAINED

23

WHY IT MATTERS

Teams deploying agents for customer-facing work should expect a large gap between ticket-solving benchmarks and real delivery performance. The bottleneck is not just information gathering but converting discovered requirements into working solutions, which suggests investment in constraint-handling and execution rather than retrieval alone.

[Read the original at arXiv.org](#) →

03

MEDIUM IMPACT

arXiv.org

Researchers Map the Routing Bottleneck Inside Modern Mixture-of-Experts Models

Fine-grained expert segmentation and shared experts — two standard modern MoE design choices — work against each other, and the paper identifies the exact mechanism: shared experts dampen the representation variability that routing exploration depends on.

The paper, submitted to arXiv on 15 October 2026, argues that modern MoE architectures combine fine-grained expert segmentation with shared experts, and that this pairing creates an unresolved tension. Fine-grained segmentation increases expressiveness by allowing more expert combinations at a fixed activation ratio, but it also enlarges the routing space combinatorially, which demands more exploratory routing. Shared experts — active for every token — reduce the variability of token representations, making routing less exploratory. The result is that the additional expressiveness offered by fine-grained segmentation may not be fully realized in practice.

The authors call this an "exploration dilemma" and propose MirrorMoE as a direct response. MirrorMoE shrinks the exploration space through token-adaptive suppression of unnecessary experts. Rather than letting the router search across the full combinatorial space while shared experts flatten the signal it needs to search effectively, MirrorMoE prunes candidates per token, reducing the space the router must explore and aligning exploration with the representational signal that remains.

The claim is empirical: MirrorMoE consistently improves performance across model scales. The abstract does not include specific benchmark numbers, parameter counts, or training budgets, so the magnitude of the improvement is not quantified in the source. What is new here is the diagnosis itself — the identification of shared experts as an exploration suppressant, and the framing of the problem as a conflict between two individually reasonable design choices. That framing is more specific than the usual observation that MoE routing is hard to train.

The proposed fix is architectural rather than a loss term or a routing-policy tweak. Token-adaptive suppression is a mechanism that operates at inference over the expert set, not a change to how experts are trained or how the router is initialized. That positions MirrorMoE as a drop-in modification to existing fine-grained MoE stacks, though the abstract does not specify implementation constraints, overhead, or compatibility requirements.

KEY FACTS	
SUBMITTED	PROPOSED METHOD
15 Oct 2026	MirrorMoE
MECHANISM	CLAIMED EFFECT
Token-adaptive suppression of unnecessary experts	Consistently improves performance across model scales

WHY IT MATTERS

Teams deploying fine-grained MoE models with shared experts may be leaving capacity on the table; the paper gives a concrete mechanism to check and a suppression-based intervention to evaluate.

[Read the original at arXiv.org](#) →

AI Tools & Ecosystem

3 stories ● 3 medium

04 MEDIUM IMPACT TechCrunch

Meta Quietly Cuts the Terms That Would Have Let It Scrape Publisher Sites for AI Training

Meta has deleted the commercial terms language that legal experts said could have authorized scraping publisher websites for AI training.

The provision, added in a recent update to Meta's commercial terms of use, stated the company could "use any content from [a business account's] domains, web pages, or URLs to provide, develop, improve, research, train, and maintain the integrity of" its products and services. Legal experts interviewed by Nieman Lab called it a "remarkable provision" that could give Meta permission to scrape the website of any publisher holding a Meta business account and use that content to train AI models.

Meta's public position was that the language consolidated existing crawling terms and applied only to ads and business products, not to how the company crawls the web generally. A spokesperson insisted the update did not reflect a change in crawling behavior. The Verge subsequently reported the provision has been removed, and Meta issued a statement: "We've heard you and want to clarify: the crawling terms we consolidated only applied to our ads and business products, and we've clarified that in the terms themselves. It does not change how Meta uses web crawlers."

The context is ongoing copyright litigation against AI companies. Meta defeated a high-profile lawsuit over alleged copyright infringement last year, but other cases against Meta and other AI firms continue. The removal appears to be a response to publisher and legal scrutiny rather than an admission that scraping practices changed.

KEY FACTS

PROVISION STATUS

Removed from commercial terms of use

PRIOR LAWSUIT OUTCOME

Meta defeated high-profile copyright infringement suit last year

META STATEMENT

Crawling terms applied only to ads and business products

REPORTED BY

The Verge, after Nieman Lab analysis

WHY IT MATTERS

The retraction removes a contractual ambiguity that could have exposed any business with a Meta account to having its site content used for AI training. Teams evaluating data provenance and licensing risk should note that Meta's terms no longer contain that language, but the company has not changed its stated web-crawling practices.

05 MEDIUM IMPACT bloomberg.com

China Puts Model Weights, Agents and Chip Designs on Its Export-Control List

China has placed model weights, AI agents and chip designs on its national export-control list.

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

A draft from Beijing's commerce and science ministries would add 40 technologies to the export-control list, 23 of them AI-related. The list includes large language models, agents, chip designs and computing clusters. Model weights and source code would be restricted above certain performance thresholds. API and SaaS access are exempted. Consultation runs to November 15. If passed, China's best weights would stop crossing its borders.

WHY IT MATTERS

The proposal could reshape how foreign developers access China's most advanced AI capabilities, though its final form is not yet certain.

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

06 MEDIUM IMPACT Business Insider

Starlink's Call Centers: 20,000 Grok Voice Calls a Day, 70% With No Human

SpaceX has put Grok Voice into production at Starlink call centers, claiming 20,000 AI-handled customer calls per day with a 70% no-human resolution rate.

SpaceX announced a new Grok Voice model this week and disclosed that Starlink call centers are already using the technology to answer 20,000 customer calls a day. The company says roughly 70% of those calls are resolved without human intervention. Grok Voice agents verify customers, troubleshoot Starlink hardware, issue credits, and ship replacement equipment, and they can perform those tasks in multiple languages.

Starlink serves more than 6 million active customers worldwide, making it a substantial real-world deployment for AI voice agents. Rather than building a large human customer-service operation, SpaceX is betting that AI can serve its multilingual customer base directly. SpaceX did not respond to Business Insider's questions about which languages Grok supports or what categories of calls still require a human.

The announcement lands amid a broader rollout of production voice models. OpenAI released GPT-Live-1 with call-capable support in July, Google has launched similar enterprise systems, and Meta's Muse can handle phone calls for users admitted to its beta. What distinguishes the Starlink disclosure is scale and specificity: a named deployment with daily call volume and a resolution rate, rather than a model card or a limited beta.

For practitioners, the unresolved questions are the interesting ones. The 70% figure is self-reported and unverified, and SpaceX has not disclosed which languages are covered, what the escalation path looks like, or how failures are measured. A 30% human handoff rate is still a substantial operational load, and the economics depend heavily on what those handoffs cost.

KEY FACTS

DAILY AI-HANDLED CALLS

20,000

STARLINK ACTIVE CUSTOMERS

6 million+

CALLS RESOLVED WITHOUT HUMANS

70%

WHY IT MATTERS

This is one of the first public, at-scale production numbers for AI voice agents in customer service. The 70% no-human resolution rate gives builders a concrete benchmark to compare against, while the unanswered questions about languages and escalation paths highlight the operational details that determine whether such deployments actually pay off.

[Read the original at Business Insider →](#)

AI Applications & Industry

5 stories ● 5 medium

07 MEDIUM IMPACT axios.com

FTC Subpoenas OpenAI and Anthropic Over Safety Incidents, Plans Executive Testimony

The FTC has subpoenaed OpenAI and Anthropic over safety incidents and plans to compel executive testimony.

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

The agency issued civil investigative demands exceeding 50 pages each, requesting documents on security incidents, safeguarding practices, privacy policies, and marketing claims. The probe focuses on consumer protection. The FTC also plans to compel top AI executives to testify and will send demands to third-party AI organizations later. OpenAI says it will cooperate.

WHY IT MATTERS

Regulatory action could set new compliance expectations for AI developers on safety, privacy, and marketing.

[Read the original at axios.com](#) →

08 MEDIUM IMPACT TNW | Artificial-intelligence

Schwab and JPMorgan Are Blocking AI Agents From Moving Client Money

Charles Schwab is blocking outside AI agents from moving client money, and JPMorgan is drawing the same line.

Schwab CEO Rick Wurster told Bloomberg Television on Friday that the firm blocks unauthorised AI agents that try to move client money, including Meta's Muse. The core problem is credential sharing: clients hand passwords to an agent that then acts on their behalf. "We can't control it, the client can't control it, something could go wrong and who's liable?" Wurster said. That liability question remains unresolved in the US, where Congress is still debating AI agent accountability.

Wurster said he receives a weekly report on these blocks, and that the big agents are stopped in the tens and twenties each week. He framed that volume as evidence clients are not yet ready to hand their accounts to AI.

Muse is Meta's personal AI agent, which can book, buy, and negotiate for users.

Instead of letting outside agents in, Schwab will stand up model context protocol (MCP) servers: secure links that let a client's chosen AI model connect to Schwab systems under rules Schwab sets. Through those links, outside AI tools will be able to check balances and do research, but not move money. Schwab is also steering clients toward Charley, its own AI assistant announced on 30 September.

JPMorgan Chase CEO Jamie Dimon took a similar position on Tuesday's earnings call, saying the bank would not let outside AI agents act inside its systems and warning of "rogue agents", per Cybernews. JPMorgan already lets some third-party tools view account data, but not send money or pay bills. The moves come as US and European regulators question banks about their use of AI agents. Notably, MCP links have had security problems of their own, with Google and JPMorgan among those that fixed the same flaw.

KEY FACTS

BLOCKED AGENTS PER WEEK

tens and twenties

SCHWAB OUTSIDE-AGENT ACCESS

MCP servers, read-only (balances, research)

MCP FLAW FIXED BY

Google and JPMorgan

SCHWAB ASSISTANT

Charley, announced 30 September

JPMORGAN THIRD-PARTY ACCESS

view account data, no sending money or paying bills

WHY IT MATTERS

For anyone building agentic finance tools, the message is that read-only MCP access is the near-term ceiling at major institutions; money movement by third-party agents is being actively blocked, and the liability question is still open.

[Read the original at TNW | Artificial-intelligence →](#)

09 **MEDIUM IMPACT** The Verge

CoreWeave's 2032 Bonds Sink Below the Distress Line — and Musk Says 'Yikes'

CoreWeave's 2032 bonds have fallen below the distress line, prompting a 'Yikes' from Musk.

Source not retrievable. This entry is written from the headline and the editor's summary only — the publisher blocked automated retrieval (extracted only 107 words (paywall/consent wall?)). [Follow the link for the full report.](#)

CoreWeave's unsecured debt due 2032 traded below 80 cents on the dollar, a level typically read as financial distress. The stock dropped 15.5% in a day. Bloomberg Intelligence now prices roughly 60% odds of default. As the most leveraged name of the AI buildout, CoreWeave's stumble could ripple across its debt-laden peers.

WHY IT MATTERS

CoreWeave's distress may signal broader fragility among debt-financed AI infrastructure providers.

[Read the original at The Verge →](#)

Anthropic's Own Numbers: Junior Technical Hires Fell From 27% to 9%

Anthropic's own workforce data shows junior technical hiring collapsing from 27% to 9% as Claude shifts work from doing to managing.

Anthropic published a report Friday on how its own workforce has changed over the past year, framing the data as a possible leading indicator of how AI reshapes hiring and labor demand. The sharpest signal is at the entry level: the proportion of junior-level employees in technical departments fell from 21% to 14% since early 2025, while junior technical hires dropped even more steeply, from 27% to 9%. Technical teams now represent 56% of the workforce, down from 61%, a shift Anthropic attributes partly to business functions scaling and partly to coding tools letting non-technical employees build software themselves.

The report also quantifies a change in what work consists of. Meeting time per employee is up 20% year over year, and up 16% for technical contributors, which Anthropic ties to the growing importance of human trust and accountability as AI output volume rises. In an April 2026 survey, 27% of employees said Claude made them more productive, while 20% said their job had shifted "from actually doing work to managing Claudes." One engineer compared the experience to babysitting: every Claude brings concerns demanding equal attention. The report describes human attention and judgment as a bottleneck, with Claude producing a "deluge" of text that must be read and organized.

On productivity, Anthropic reports more code per person, more sales per rep, and more contracts processed. It acknowledges that more code does not necessarily mean better code or business impact, but says it has found no indication that code quality is declining. The company presents these internal trends cautiously, noting that other organizations may never adopt AI as thoroughly as Anthropic, which one employee described as having "eternal pressure to move as fast as possible and use Claude for everything."

KEY FACTS

JUNIOR TECHNICAL HIRES

27% to 9%

TECHNICAL SHARE OF WORKFORCE

61% to 56%

MEETING TIME INCREASE, TECHNICAL CONTRIBUTORS

16%

JUNIOR TECHNICAL STAFF

21% to 14%

MEETING TIME INCREASE PER EMPLOYEE

20%

EMPLOYEES REPORTING PRODUCTIVITY GAIN FROM CLAUDE

27%

WHY IT MATTERS

If Anthropic's trajectory generalizes, teams should expect thinner junior pipelines and a growing coordination tax as AI output shifts senior staff toward review and orchestration rather than direct production.

[Read the original at TechCrunch](#) →

18 Months and an \$8M Fine for the Man Who Faked Billions of AI Music Streams


The first US criminal conviction for AI music streaming fraud has produced an 18-month sentence and an \$8 million royalty figure, establishing that bot-driven AI catalogues carry criminal liability.

Michael Smith of North Carolina was sentenced to 18 months in prison after pleading guilty to one count of conspiracy to commit wire fraud. Between 2017 and 2024, Smith created thousands of fake accounts on streaming platforms using fake email addresses and fraudulently obtained debit cards, then ran software to make those accounts repeatedly stream his music. At peak, he operated as many as 10,000 bot accounts simultaneously. To avoid detection, he spread streams across thousands of songs, and from roughly 2018 he worked with the CEO of an AI music company and a music promoter to generate hundreds of thousands of AI songs, receiving thousands of new tracks each week for upload.

The scale is quantified in the DOJ's comparison: in April 2023, Taylor Swift's entire catalogue received 9.3 million streams on YouTube Music from family plan streams, while Smith's bot accounts used family plans to stream his AI-generated music 80.9 million times in the same month. Total fraudulent royalties exceeded \$8 million.

What is new here is the criminal precedent, not the fraud technique. Bot-driven streaming manipulation predates generative AI; the AI element is the supply chain — an AI music company producing songs at a volume that made the scheme harder to detect and prosecute. The DOJ framed the case explicitly around AI-enabled fraud, and the RIAA and IFPI both used the sentencing to signal that enforcement, prevention, detection and disruption will be pursued together. The IFPI pointed to its Streaming Integrity Initiative as the industry's coordinated response.

For practitioners, the operational lesson is that distribution-side integrity checks are now a legal exposure question, not just a platform policy issue. Anyone building tooling that generates or distributes audio at scale should expect royalty accounting, bot detection and provenance to be scrutinised under criminal fraud statutes, not only terms of service.



KEY FACTS

SENTENCE

18 months in prison

PEAK CONCURRENT BOT ACCOUNTS

10,000

SMITH BOT STREAMS, APRIL 2023 (YOUTUBE MUSIC FAMILY PLANS)

80.9 million

FRAUDULENT ROYALTIES

More than \$8 million

AI-GENERATED SONGS

Hundreds of thousands

TAYLOR SWIFT CATALOGUE STREAMS, APRIL 2023 (YOUTUBE MUSIC FAMILY PLANS)

9.3 million

WHY IT MATTERS

AI-generated content at scale is now a demonstrated vector for criminal streaming fraud, and the conviction establishes that both the operator and the supply chain — including AI music providers — can be drawn into wire fraud conspiracy exposure. Teams building generative audio pipelines or distribution automation should treat stream integrity and provenance as legal risk, not just platform compliance.

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

Sources

- 01 MLExpert Is a Generalist Agent Built to Run Machine-Learning Engineering End to End
<https://arxiv.org/abs/2610.22344>
- 02 FDE-Gym Benchmarks AI Agents on Real Customer Delivery, Not Tidy Tickets
<https://arxiv.org/abs/2610.22350>
- 03 Researchers Map the Routing Bottleneck Inside Modern Mixture-of-Experts Models
<https://arxiv.org/abs/2610.22353>
- 04 Meta Quietly Cuts the Terms That Would Have Let It Scrape Publisher Sites for AI Training
<https://techcrunch.com/2026/10/17/meta-removes-language-that-raised-scraping-and-ai-training-concerns/>
- 05 China Puts Model Weights, Agents and Chip Designs on Its Export-Control List
<https://www.bloomberg.com/news/articles/2026-10-16/china-proposes-export-controls-on-ai-models-chips-and-agents>
- 06 Starlink's Call Centers: 20,000 Grok Voice Calls a Day, 70% With No Human
<https://www.businessinsider.com/starlink-customer-service-grok-2026-10>
- 07 FTC Subpoenas OpenAI and Anthropic Over Safety Incidents, Plans Executive Testimony
<https://www.axios.com/2026/10/15/ftc-openai-anthropic-ai-safety-investigation>
- 08 Schwab and JPMorgan Are Blocking AI Agents From Moving Client Money
<https://thenextweb.com/news/schwab-blocks-ai-agents-meta-muse>
- 09 CoreWeave's 2032 Bonds Sink Below the Distress Line — and Musk Says 'Yikes'
<https://www.theverge.com/ai-artificial-intelligence/1011460/neocloud-coreweaves-bonds-are-distressed-and-so-am-i>
- 10 Anthropic's Own Numbers: Junior Technical Hires Fell From 27% to 9%
<https://techcrunch.com/2026/10/17/anthropic-says-its-workforce-has-become-less-junior-and-is-spending-more-time-in-meetings/>
- 11 18 Months and an \$8M Fine for the Man Who Faked Billions of AI Music Streams
<https://www.musicweek.com/digital/read/us-man-sentenced-to-18-months-in-prison-for-8-million-ai-music-streaming-fraud/095182>

About this document. Every story in the 18 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 2026-10-18T05:33:17Z · model deepseek-v4-pro:cloud · New Horizon, new-horizon.tech · To unsubscribe or manage your subscription, see the footer of the daily email.