

The OpenAI bugging case leads to Saudi Arabia

SpaceX shares for US children, AI code nobody understands, Jack Clark at the IMF and the DeepSeek ban

Tuesday, 20 October 2026

IN THIS ISSUE

1. The OpenAI bugging case: prosecutors name Saudi Arabia as a supporter
2. Musk promises one SpaceX share to each child in the US after the Anthropic fund
3. Programmers and AI code that they do not understand
4. Jack Clark at the IMF: "the singularity has started"
5. The cheap tier of models: MiniMax M3.1 Flash and the diffusion model Mercury 3
6. The US removed DeepSeek from civilian federal systems
7. Mistral: a model for time series and a physics model with CERN
8. A Russian satellite came within hundreds of metres of a radar satellite that Ukraine uses

TOPIC 1

The OpenAI bugging case: prosecutors name Saudi Arabia as a supporter

SOURCES Ars Technica, 19.10 [Ars Technica](#) · BBC, 20.10 [BBC](#) · The New York Times, 19.10 [NYT](#) · The Wall Street Journal, 19.10 [WSJ](#) · Financial Times, 19.10 [FT](#)

On 14.10 this digest wrote about the arrest of the OpenAI engineer Abdulrahman Alfozan. Prosecutors accuse him of hiding a recorder in a meeting room. On Sunday, prosecutors filed new details in court. On Monday, Judge Alex Tse kept Alfozan in custody until the trial, because he can flee.

OpenAI found the device on 30.09. It was behind a TV and covered with toilet paper. Prosecutors write that the device stayed there for several days. In those days, OpenAI leaders met US government representatives in the same room. On the phone of Alfozan, the FBI found a photo of one more hidden recorder in another meeting room.

The hardest part is about Saudi Arabia. According to the prosecution, Alfozan wrote in WhatsApp with a member of the royal family from 2025 to 05.10. A message said "Saudi Arabia needs you more than the US", and he answered "100". On 04.10, he opened a folder "Saudi Gov Docs" on his phone.

Prosecutors say that the Saudi government paid for a one-way ticket to Riyadh for him, his wife and his child. After the arrest, a representative of the Saudi military bought business class tickets to Doha for his wife and son.

One more detail is about the OpenAI tools. Prosecutors write that Alfozan used an internal OpenAI agent to find convenient meeting rooms. He also asked ChatGPT from the account of his wife how to make a recorder start on voice. He used the same agent to search for "10 people from OpenAI for an informal meeting with ministers of Saudi Arabia".

Now the charge is illegal recording of conversations, up to 5 years in prison. Prosecutors say that they can add more serious charges. They accept that Alfozan could act for himself, and they say that this does not matter for the case.

Why it matters. Frontier labs became a target of state intelligence, like defence contractors.

This case shows a cheap attack path: an employee with access and a home-made recorder. Here, the logs of an internal agent became evidence against its user. A company that gives agents to its employees also gets a detailed record of their intentions.

TOPIC 2

Musk promises one SpaceX share to each child in the US after the Anthropic fund

SOURCES The New York Times, 19.10 [NYT](#) · Semafor, 19.10 [Semafor](#) · The Wall Street Journal, 19.10 [WSJ](#) · BBC, 19.10 [BBC](#)

Yesterday this digest wrote that Anthropic gives 5% of its shares to a public benefit fund. A few hours later, Elon Musk wrote on X that next year he will give one SpaceX share to "every child in America". Semafor writes that Musk did this in reply to the post of Anthropic. Musk did not explain how he will give the shares. Later, Trump thanked Musk and wrote that the shares will go to "Trump accounts". This is a government savings account program for children.

The NYT calculates it this way: the US has about 72 million children under 18, and one SpaceX share costs about \$148. The total is about \$10.5 billion. Semafor gives \$10.8 billion, which is about 1% of the SpaceX shares of Musk. For comparison, Michael and Susan Dell earlier pledged \$6.25 billion in cash for these accounts.

Musk has a history of promises that he did not keep. The NYT reminds the reader that the Musk Foundation had more than \$14 billion in assets, but for four years in a row it did not give the minimum that the law requires. In 2021, Musk pledged \$6 billion to the UN World Food Programme if it showed a plan. The programme showed a plan, and Musk did not give the money. Semafor raises one more question. SpaceX has two classes of shares, and Musk controls the votes. The voting rights of the children are still unknown.

Why it matters. Shares of AI labs become a political tool before the big IPOs. Anthropic tied its donation to the risks for the labor market. Musk tied his donation to a White House program.

In both schemes, a large block of shares ends up with an owner that is friendly to the company.

Semafor writes directly that a company can use this to build a loyal block of shareholders.

TOPIC 3

Programmers and AI code that they do not understand

SOURCES Ask HN, 19.10 [Hacker News](#) · Paul Battley, 18.10 [Po-ru](#) · discussion on HN, 19.10 [Hacker News](#) · Simon Willison, 19.10 [Simon Willison](#) · Greptile on X, 19.10 [@greptile](#)

Two threads about the same topic got the most comments on HN in one day. The first is the question "How do you live with shipping AI code that you do not understand?". The author quotes a post of a Meta engineer: "Code became a black box. If you check every line, everything takes forever". The question has 435 comments.

The most popular answer is simple. The person designs the abstractions and the links between them, and the model only fills in the implementation. Then the author knows how the system works, even if he did not read each function. Other commenters write in a darker tone. One senior with 10 years of experience says that he just approves the PRs of his colleagues, because he physically cannot read thousands of changed lines per week.

The second thread is about Paul Battley, who archived all his Ruby gems. He writes that agents now develop Ruby to a large degree, and Matz calls Ruby "a language for the AI era". Battley does not want to add to the shared code that companies collect to train models. The post has 418 comments.

On the same day, Willison quoted the "obituary for programming" by Jeremy Scheff. The main idea:

if you liked to solve puzzles in code, the loss is real. If you like to build systems, AI is a gift.

Meanwhile, the market pays for review. Greptile, which makes AI review of pull requests, raised a \$50 million Series B at a valuation of \$750 million. The company writes that it checks more than 1% of all PRs that people open on GitHub in one day. The team has fewer than 50 people.

Why it matters. Code became cheap to write, but not cheap to understand. The bottleneck moved to review, and investor money goes there. The practical advice from the thread is cheap: keep the architecture in the head of a person and give the lines to the model. Without this, the team gets a system that it cannot explain.

TOPIC 4

Jack Clark at the IMF: "the singularity has started"

SOURCES Import AI 477, 19.10 [Substack](#) · The Rundown AI, 19.10 [@TheRundownAI](#)

Jack Clark, a co-founder of Anthropic, gave the Per Jacobsson Lecture on 17.10 at the IMF meeting in Bangkok. He published the full text in his newsletter. The main thesis: "this year we got convincing evidence that the singularity has started".

Clark takes his arguments from science. On Humanity's Last Exam, the best result at the start in January 2025 was 8%, and now Claude Opus 5.5 has 67.7%. On FrontierMath, the result grew from 2% to 94%. One more example is the TheoPhysEval test with 50 unpublished problems of theoretical physics. GPT-6 Astra first scored a little below 90%. Then the model insisted that the experts were wrong, and in several problems the experts really found errors in their own answers. After the fixes, Astra has 49 of 50.

Clark says that Anthropic researchers write almost 10 times more code than in the first four years of the company. According to the Anthropic model, in a fast scenario the US GDP can reach \$44 trillion by 2030. He asks governments to build labs, robots and science equipment, as before he asked them to build data centres. And second, to limit the development of the technology itself.

The moderator Raghuram Rajan asked a question, and Clark reminded him that Dario Amodei wrote about possible "limits on the ceiling of intelligence" of the systems. The window to act, in his words, is "a few years".

Why it matters. A co-founder of a lab speaks about the singularity, and he speaks to central bankers and ministers. Yesterday this digest gave the number of Michael Sellito that Claude can run 80-90% of model training. Clark gives a wider picture from the same company. The benchmark numbers in the speech can be checked. The GDP forecast is a model of Anthropic itself.

TOPIC 5

The cheap tier of models: MiniMax M3.1 Flash and the diffusion model Mercury 3

SOURCES MiniMax, 19.10 [Minimax](#) · discussion on HN, 19.10 [Hacker News](#) · Inception, 19.10 [Inceptionlabs](#) · discussion on HN, 19.10 [Hacker News](#)

Yesterday this digest wrote that a better model no longer guarantees developer traffic. In one day, two models came out that compete on price and speed.

MiniMax released M3.1 Flash. The price is \$0.15 per million input tokens and \$0.60 per million output tokens. The company writes that on agent tasks the model costs 70% less than the previous one and works almost three times faster. The results, according to MiniMax itself: 73.7% on DeepSWE v1.1, 61.7 on OSWorld V2, 82% on CyberGym. The anonymous model Space Bunny, which processed more than 200 trillion tokens, was an early version of M3.1 Flash. The company promises to publish the weights later.

Inception released Mercury 3. It is a diffusion model. It refines a whole piece of text in a few passes, while a usual model writes one token at a time. Inception writes that Mercury 3 is more than 40 times larger than Mercury 1 and gives 1068 tokens per second on an NVIDIA

GPU. The context is 1 million tokens. The price is \$0.20 and \$0.70 per million, with a 50% discount at launch. According to the company, the model is not worse than GPT-6 Luna and Claude Haiku 5.5.

Why it matters. For agents, the speed of each step is multiplied by the number of steps. Thus in the cheap tier, latency now matters as much as price. The numbers of both companies are still their own. The HN threads have no independent measurements yet, so compare the models on your own tasks.

TOPIC 6

The US removed DeepSeek from civilian federal systems

SOURCES [BBC](#), 19.10 [BBC](#) · [The Wall Street Journal](#), 19.10 [WSJ](#)

On 18.10 this digest wrote that DeepSeek published V4.1 Flash under MIT. Now the US Department of Homeland Security ordered to remove DeepSeek technology from civilian federal agencies. These include the Treasury, the State Department and the Justice Department. The order does not cover the Department of Defense. The order became public on Monday, but it has been in force since the end of July.

The WSJ writes that the ban also covers future purchases. It covers products of companies that are affiliated with DeepSeek and third-party products that are built on DeepSeek models. The order does not explain which threat the department sees. The BBC reminds the reader that Snowflake, Vercel and OpenRouter use DeepSeek.

Why it matters. The phrase "built on DeepSeek models" is the widest part of the order. It is easy to put open weights inside a product of another company, and the vendor does not always say which model is inside. For companies that sell to the US government, the origin of a model becomes a check item, like the origin of a chip.

TOPIC 7

Mistral: a model for time series and a physics model with CERN

SOURCES [Mistral](#), 15.10 [Mistral](#) · [discussion on HN](#), 19.10 [Hacker News](#) · [CERN](#), 15.10 [Home](#) · [discussion on HN](#), 19.10 [Hacker News](#)

Both announcements came out on 15.10. But HN discussed them yesterday, and they were not in these issues before.

Timestrail is the first family of Mistral models for time series forecasts: sensor readings, metrics, financial flows. There are four models, with 5, 28, 123 and 209 million parameters.

Mistral writes that the largest one is first on three public benchmarks: GIFT-Eval, TIME and fev-bench. According to the company, the model with 28 million parameters beats Chronos-2, which is about four times larger. A GPU is not necessary. On 8 CPU threads, a forecast for one day ahead takes from 10 to 140 ms. Most of the training data is synthetic.

CERN and Mistral signed a statement of intent to work together. The goal is a "Large Physics Model", the first foundation model for collider physics. They will train it on billions of simulated proton collisions at the Large Hadron Collider. CERN has about one exabyte of data, and after the start of the HL-LHC around 2030 there will be ten times more collisions.

Why it matters. Mistral builds a niche that large American labs almost do not take: small specialised models for industry and science that people can run on their own hardware. Time series forecasts without a GPU are a direct product for a factory or a power company. The partnership with CERN is only a statement of intent so far.

TOPIC 8

A Russian satellite came within hundreds of metres of a radar satellite that Ukraine uses

SOURCES Ars Technica, 19.10 [Ars Technica](#) [one source]

The company LeoLabs, which tracks orbits with radars, reported a maneuver of the Russian military satellite Kosmos 2612. It came close to ICEYE-X36, a commercial radar satellite from the American subsidiary of the Finnish ICEYE. Last week, Kosmos 2612 started a fast approach. On Friday it stayed within 2-3 km, and on the weekend the satellites were closer than 500 m. LeoLabs calls this the first known unauthorised approach in low orbit.

Kosmos 2612 launched in April together with five other satellites. At least four of them adjusted their orbits to ICEYE-X36. LeoLabs writes that this is not a random flyby: there was a controlled approach, a fly-around, a long stay nearby and a controlled departure. ICEYE satellites see the surface through clouds and at night. Ukraine has received intelligence images from ICEYE since

1. In the same year, a Russian official called the civilian satellites that help the Ukrainian army a "legitimate target".

Why it matters. Commercial satellites became part of the war, and now people chase them physically in orbit. For Ukraine, this is a risk to an intelligence source that does not depend on the decisions of one company or one government. After the story with the Starshield limits, this dependence is especially visible.

misc

- **Anthropic and Nscale** negotiate a second multi-billion compute deal in British data centres, the FT writes. This continues item 4 of yesterday: Microsoft left the Nscale project because of a lack of power. The details are behind a paywall.

[FT](#)

- **Deutsche Telekom, Orange, Telefónica and Vodafone** consider a joint bid for spectrum for mobile satellite communication, to compete with Starlink, the FT writes.

FT

- **Disney and Character.AI:** from the bankruptcy documents, Ars Technica explains that Disney demanded that the deal go through bankruptcy. This way Disney gets the people and the technology, and the lawsuits stay with what remains of Character.AI. This continues item 6 of 14.10.

Ars Technica

- **Linux 7.3** came out with a record: 18,765 commits from 3,021 developers, 807 of them new. Linux 7.2 had 16,418 commits from 2,652 people. [Lwn](#)
- **@OpenAI** writes that the new models write and port code in COBOL, Fortran and Visual Basic better. The company gave no numbers. [one source] [@OpenAI](#)
- **@emollick:** the new Claude Projects are a convenient interface for several agents. A coordinator model talks to the human, and the agents work in their own branches. The old Projects share only the name with them. [one source] [@emollick](#)
- **Plex:** attackers exploit a chain of vulnerabilities in Plex Media Server and install a crypto miner, and later a proxy. Versions 1.43.2 and older are vulnerable.

Boredhackerblog

- **GitHub Actions, Pages and webhooks** had outages. GitHub promises a root cause analysis.

Githubstatus

- **Waymo** wants to grow its fleet from 4,000 to 20,000 cars by the end of next year, Semafor writes. [Semafor](#)
- **Your own LLM proxy:** the author of an essay advises to send all requests to models through your own proxy. This shows token costs over time, and the keys do not spread over the agents.

209 comments on HN. [Hacker News](#)