

Ukraine gets Starlink speeds back as Musk blames Kyiv

Apollo sees no AI payoff outside tech, scanners miss Brotli in Python, AI text is hard to read

Sunday, 25 October 2026

IN THIS ISSUE

1. Starlink in Ukraine works at full speed again. Musk says that Ukraine broke the rules
2. Europe after Starlink: "who holds the kill switch"
3. Apollo: outside the tech sector, AI gave nothing yet
4. Joshua Saxe: where are the hordes of AI attackers?
5. Brotli in Python: dependency scanners do not see the vulnerability yet
6. Why AI text is hard to read: researchers measured it
7. The FBI can add an economic espionage charge to the case of an OpenAI engineer
8. Spain: the opposition made an ad with a deepfake of the prime minister's wife

TOPIC 1

Starlink in Ukraine works at full speed again. Musk says that Ukraine broke the rules

SOURCES The New York Times, 24.10 [NYT](#) · [@ZelenskyyUa](#), 24.10 [@ZelenskyyUa](#) · [@elonmusk](#), 24.10 [@elonmusk](#) · [@sternenko](#), 24.10 [@sternenko](#)

On 23.10 this digest wrote that Trump talked with Musk about the Starlink speed limit in Ukraine. On 24.10 Volodymyr Zelensky said that Starlink and Starshield again give the speed from the time before the limits. The recovery started at night. Defense Minister Yevhen Khmara said that drones on Starshield flew at about 250 miles per hour for the first time since the start of the limits.

According to the NYT, the limits stayed in place for more than 10 days. Some drones slowed to 40 miles per hour and less. Some of them lost the connection with the pilot and fell. One commander told the NYT that on the first night his unit lost drones for more than \$1 million.

On the same day, Musk made his first comment about the situation. He wrote that the Ukrainian military used civilian Starlink "for offensive drone strikes", and SpaceX closed this "loophole" last week. According to him, the US government controls Starshield. SpaceX also sent a letter to Congress.

The company writes that it changed nothing in Starshield. It says that it changed Starlink because of "offensive military air missions".

Serhii Sternenko, a former adviser to Defense Minister Mykhailo Fedorov, disagrees. Since February 2026, Ukraine uses whitelists. A terminal that works in motion must be on a special list and on Starshield. Civilian Starlink stops at 60-70 km/h already, and drones fly from 120 km/h.

Thus, according to Sternenko, it is technically impossible to use civilian Starlink for drones.

The NYT adds political context. Senators from both parties warned the White House. They will block Trump nominees for posts if Ukraine does not get full Starshield back.

Why it matters. The connection returned, but one company still writes the rules. The sides still describe in different ways which system did what. For an army that depends on a foreign connection, this is a risk even on days when everything works.

TOPIC 2

Europe after Starlink: "who holds the kill switch"

SOURCES Steve Blank, 24.10 [Steveblank](#) · HN, 24.10 [Hacker News](#) · Computerworld (Denmark), 20.10 [Computerworld](#) · HN, 24.10 [Hacker News](#)

Steve Blank teaches entrepreneurship at Stanford, and he spent a month in London. He writes that in the US the Starshield story was short news, but in Europe it caused a shock. For years, the fear of an "American kill switch" in weapons and critical software looked like paranoia. Blank admits that he also thought so.

Now, according to him, European politicians ask a different question. What happens if Microsoft, OpenAI, Anthropic, Oracle, Palantir or Google also limit access? The conclusion of Blank: an American founder who sells to European defense or the public sector now needs an answer to the question "who holds the kill switch". European founders got sovereignty as a sales argument. Separately, he writes that AI makes it cheaper to build local equivalents of Office, Salesforce or Palantir.

HN shows a counterexample next to it. On 20.10 Computerworld wrote that the Danish ministry of digitalization quietly closed a pilot with Collabora, an open replacement for Microsoft Office. The pilot started in summer 2025, and almost 40 ministry employees tested it. The open part of the article gives no reasons for the closure.

Why it matters. Statements about sovereignty and real purchases still go in different directions. A move away from Office needs money and patience, and a political decision alone is not enough. The question "who holds the kill switch" now comes at the start of each large deal with an American supplier.

TOPIC 3

Apollo: outside the tech sector, AI gave nothing yet

SOURCES Apollo Daily Spark, 24.10 [Apollo](#) · HN, 24.10 [Hacker News](#)

Torsten Slok, the chief economist of Apollo, writes directly: outside the tech sector, AI gave corporate America zero financial return. He compares the Magnificent 7 with the rest of the index, the S&P 493. Revenue per employee and margin grow in the Magnificent 7 and stay flat in the rest. The data also shows no signs of a growth in productivity.

Slok writes that ordinary companies already cut their AI costs. For them, tokens became a cost line that eats the margin. Companies with venture money need the newest model. The rest of the market does not need it and does not want to pay for it. While agents live in a small part of the economy, the macro effect of AI stays small.

On HN the thread got 508 points and 523 comments. The main objection: a quarterly margin is a bad measure for a technology that changes processes over years.

Why it matters. The S&P 493 will decide the AI story, and Slok says this directly. A company that calculates the return from agents gets a strong argument: measure the margin and revenue per employee. The number of tokens proves nothing.

TOPIC 4

Joshua Saxe: where are the hordes of AI attackers?

SOURCES Joshua Saxe, 23.10 [Substack](#) · Marginal Revolution, 24.10 [Marginal Revolution](#) · @levie, 25.10 [@levie](#)

Yesterday this digest wrote about "Black October": Altman no longer believes that it is possible to patch all bugs. Joshua Saxe worked on LLM security at Meta until 2026. He looks at the same picture through data. According to his chart, the damage from cybercrime grew three times in ten years. But the main growth came before ChatGPT, and since 2022 the curve even slowed.

The explanation of Saxe is economic. According to his estimate, approximately 720 000 people work in online crime, and 90% of them sit in scam call centers. There are approximately 76 000 real hackers, and most of them work for states. There are only a few thousand criminal hackers, mostly in ransomware. Saxe checked the fate of 152 known groups, and most of them got legal consequences.

According to Saxe, a vulnerability is only the first step. Then the attacker must get a foothold, demand a ransom and launder the money. Criminals will need years to build a new business model on agents. Tyler Cowen repeated this idea in Marginal Revolution on 24.10.

The conclusion of Saxe does not calm the reader. He thinks that now is a "window of defender advantage", and that defenders must prepare now. On the other side, Aaron Levie from Box writes that cyber agents will soon use 1000 times more tokens.

Why it matters. The number of CVEs and the size of losses are different metrics. A security policy that counts only found vulnerabilities misses the main limit of the attacker: people,

money and the risk of arrest.

TOPIC 5

Brotli in Python: dependency scanners do not see the vulnerability yet

SOURCES Aikido Security, 23-24.10 Aikido · HN, 24.10 Hacker News · Reuters, 23.10 Reuters · HN, 24.10 Hacker News

Yesterday this digest wrote about HTTP2Shell, the hole CVE-2026-111497 in the Brotli library. Aikido examined what this means for Python. The packages `Brotli` before 1.2.0 and `brotilicffi` 1.2.0.x build the same vulnerable decoder as Node.js. The `Brotli` package has 82,8 million downloads per month.

The problem is that the application code can have no Brotli import at all. `requests`, `httpx` and `aiohttp` use these packages automatically if they are installed. Scrapy installs Brotli as a dependency. The attacker sends a malicious response to a request of the application, or a compressed request body. According to Aikido, on 23.10 the vulnerability databases did not yet connect the CVE with these packages. Thus `pip-audit`, Dependabot and Safety do not show the vulnerable versions.

Fixed versions: `Brotli` 1.2.1 and `brotilicffi` 1.2.1.0. After the update, rebuild the containers. If an update is not possible, Aikido recommends to remove both packages or to reject `br` in `Content-Encoding` before decompression. To remove `br` from `Accept-Encoding` is not enough, because a malicious server can send Brotli all the same.

One more detail of the week came from Reuters. PHP possibly runs on 70% of the world's websites, but the PHP Foundation has one full-time employee and 14 part-time contractors. Executive director Elizabeth Barron says that because of AI, the project gets almost two times more proposed changes than a year ago. The number of contributors did not double.

Why it matters. A green scanner report in the first days after a disclosure proves nothing. Find transitive dependencies manually: `pip show Brotli brotilicffi` in each image. And the maintainers of the infrastructure under the internet get from AI both findings and a double workload.

TOPIC 6

Why AI text is hard to read: researchers measured it

SOURCES arXiv, 21.10 arXiv · HN, 24.10 Hacker News

Thomas Clark, Matias Martin and Kyle Mahowald from UT Austin checked the feeling that model text became harder to read. They collected a corpus of chats, code and essays from 14 models, among them Claude, GPT-5.6, Kimi and GLM. Then they compared it with human texts and with paraphrases through double translation.

All nine frontier models use rare word combinations more often than people. They also put an inanimate subject with a verb for a living actor ("is_file() swallows that"). Most of them also make longer syntactic links, put three or more nouns in a row and make verbs into nouns. Psycholinguistics connected all these signs with hard reading a long time ago.

The authors propose a cause. After the main training, a model gets a reward for a correct answer.

There is no reward for clarity, but there is a penalty for length. In an experiment with such a penalty, the model started to choose shorter constructions, even where nobody trained it to do so.

The effect grows from the base model to the instruct model and then to the reasoning model.

Why it matters. The "AI style" has measured signs, and a script can check them. For a team that publishes text after a model, a simple rule helps: remove the noun stacks and the inanimate "heroes"

of sentences. Measure a text for a reader by how well the reader understood it.

TOPIC 7

The FBI can add an economic espionage charge to the case of an OpenAI engineer

SOURCES The Wall Street Journal, 24.10 [WSJ](#)

[single source] On 20.10 this digest wrote about the case of Abdulrahman Alfozan. Prosecutors accuse the former OpenAI engineer of a hidden recorder in a closed meeting room. The WSJ cites a law enforcement official and writes that the FBI considers an economic espionage charge. The investigation checks if Alfozan stole OpenAI information to offer it to the government of Saudi Arabia. According to the WSJ source, this is possible even without recruitment.

Why it matters. If the charge comes, the case moves from a simple recording of conversations to the protection of trade secrets. For frontier labs, this is one more signal that the state protects them as defense contractors.

TOPIC 8

Spain: the opposition made an ad with a deepfake of the prime minister's wife

SOURCES The Guardian, 23.10 [Guardian](#)

[single source] One month before the elections, the People's Party of Spain released a video with AI copies of people from the circle of Pedro Sánchez's party. A generated Begoña Gómez, the wife of the prime minister, says that she got a university post because of her

husband. By the evening of 24.10 the video had almost 2 million views. The AI label is at the bottom, in small semi-transparent text.

Most of the nine lawyers that El País asked think that the video is legal. Their reason: it has a label and comes as satire. One of the people whose image the video used already filed a complaint in court.

Why it matters. The label "made with AI" protects the author legally. The viewer gets no protection from it. This case will become a precedent for the rules on election ads in the EU.

misc

- **Dream-1:** the startup Dream Labs left stealth mode with \$90 million in seed money and a model for robots. It controls two arms with grippers, even on a robot that it did not see in training.

The price starts at \$3 per hour through the API. The page gives no exact release date, and on HN it became popular on 24.10. [Dreamlabs](#) · [Hacker News](#)

- **GrapheneOS and banks:** the project writes that approximately ten banking apps now explicitly permit GrapheneOS. Approximately 10% of banking apps still block an OS without Google approval.

[Grapheneos](#) · [Hacker News](#)

- **Roku moved from GitHub to GitLab:** all developer repositories moved. Roku separately warns the owners of RAG indexes and agents: an old index gives no error, it silently returns old content.

[Roku](#) · [Hacker News](#)

- **Invented citations in a peer-reviewed paper:** TorrentFreak found that a paper about Z-Library cites a TorrentFreak article that does not exist. Some more sources in the paper also do not exist.

The authors said that they worked with AI. [Torrentfreak](#) · [Hacker News](#)

- **OBS Studio 33:** the team updated the CEF browser engine from version 127 to 150, two years forward. Plugins got a new folder structure. [GitHub](#) · [Hacker News](#)