AI DIGEST

Huang discusses a return to Nvidia

Bessent wants an outside CEO, the NYT reports a Lutnick conflict of interest, and a Vulcan rocket lost 40 Amazon satellites.

Monday, 16 November 2026

IN THIS ISSUE

1. Nvidia: Huang thinks about a return, Bessent wants an outside CEO, Huang will testify to Congress
2. Chip exports: Lutnick writes rules that touch clients of Cantor Fitzgerald
3. A Vulcan rocket failed with 40 Amazon Leo satellites
4. OpenAI: GPT-6 Astra proved a lattice construction of the Standard Model
5. Qwen4-Max: 2.4 trillion parameters, cost per task down 62%
6. Bitly: strangers registered again the short links in old tweets
7. A court did not dismiss the DRAM conspiracy case against Samsung, SK Hynix and Micron
8. A week of reports: Nvidia on 17.11, the Anthropic IPO on 19.11
9. US ground stop: Duffy blames a router, controllers do not believe him
10. South Korea: the top score on the tax accountant exam came from AI

TOPIC 1

Nvidia: Huang thinks about a return, Bessent wants an outside CEO, Huang will testify to Congress

SOURCES Reuters, 15.11 Reuters · Financial Times, 15.11 FT · The Wall Street Journal, 15.11 WSJ · BBC, 16.11 BBC · The Wall Street Journal, 16.11 WSJ · @balajis, 16.11 @balajis · HN, 16.11 Hacker News

Yesterday we wrote about the campaign to bring Huang back. Now, according to the FT, Huang himself discusses a return with people close to him. One ally told the FT that he puts together a "war room" of Silicon Valley executives and will look for support from the White House. His words: "Everyone that matters in tech outside the frontier labs loves Jensen". Earlier the FT wrote that the Nvidia board spoke with Arm CEO Rene Haas as a possible successor.

Washington answered at once. Treasury Secretary Scott Bessent said on Fox News on Sunday that the board must look for a CEO outside the company: "They need a dramatic shake-up". Congressman John Moolenaar told CBS that Huang agreed to testify to three House committees. The hearings will be closed, and the dates are not final. The BBC writes that the chance of a return is low while the internal investigation continues.

Megaspeed also made a statement. The company says that Alice Huang left her executive role in August 2025.

Megaspeed started its own investigation and calls part of the reports "not supported by verifiable evidence". The analysts at C4ADS think that Alice Huang stayed "practically in control" of the company after she left.

Two lines go against the general tone. First, according to the WSJ, the Trump administration will name Nvidia this week among more than 20 AI "national champions". The program must help US companies sell AI abroad. Second, Balaji Srinivasan argues in a long thread that the rental of Nvidia chips through Megaspeed does not break current law. He cites the Atlantic Council: the RASA bill on remote access passed the House, but not the Senate.

Why it matters: on 17.11 Nvidia reports its quarter, and this is the first report without Huang. Three forces now push on the company at the same time: the campaign for his return, the government demand for an outsider, and the "national champion" status from the same government. The board decision will show which force is stronger.

TOPIC 2

Chip exports: Lutnick writes rules that touch clients of Cantor Fitzgerald

SOURCES The New York Times, 15.11 NYT · Financial Times, 15.11 FT · Financial Times, 15.11 FT · BBC, 16.11 BBC · HN, 15.11 Hacker News

Chinese companies cannot buy the best Nvidia chips, but they can rent them in data centers in Malaysia, Indonesia or Thailand. The Commerce Department of Howard Lutnick now writes the rules for this rental. The NYT found a conflict of interest. The Singapore company Bitdeer builds data center A201 in Malaysia, and a NYT source says that Tencent will probably use it. Cantor Fitzgerald, which Lutnick led before and which his sons lead now, was a sales agent for a Bitdeer stock offering of up to $1 billion.

The NYT also describes the draft rule from the department. It covers only 7 countries: Singapore, Malaysia, Thailand, Indonesia, Brazil, South Africa and Norway. The rule applies only to new shipments, and data centers that are already full are exempt. A company must only notify the government when a Chinese client trains a model, and no license is necessary. Chris McGuire from the Council on Foreign Relations says: "China will just rent from the data centers outside those countries". The department says that Lutnick obeys all ethics agreements and does not discuss department matters with his sons.

Meanwhile, Congress moves a stricter chip export law, the FT writes. The US and its allies signed an accord to protect AI supply chains from China. Lutnick will testify to the Senate Banking Committee next month, and now he flies to China with Trump.

Why it matters: the scandal around Huang became a question about who controls AI compute outside China, and how. The line between the sale of a chip and the sale of access

to it is not clear. Projects worth billions in Southeast Asia depend on the version of the rules that the government accepts.

TOPIC 3

A Vulcan rocket failed with 40 Amazon Leo satellites

SOURCES Ars Technica, 15.11 Ars Technica · Space.com, 15.11 Space · HN, 15.11 Hacker News

Yesterday we wrote that SpaceX caught the Starship upper stage with its tower for the first time. One day later a competitor failed. On 15.11 a ULA Vulcan rocket launched from Cape Canaveral with 40 Amazon Leo satellites. At approximately minute 10 of the flight ULA reported an "anomaly", and the satellites did not reach orbit. This is the first total mission loss for ULA since 2006.

Sources told Ars that the problem was in the first stage. The Blue Origin BE-4 engines burned at least 30 seconds longer than the plan, and then the rocket started to spin. This was the first Vulcan launch after a 9-month pause because of failures in the side boosters. This time the boosters worked normally.

The effects are wide. Amazon has 37 more launches booked on Vulcan, and only 396 of 3,292 first-generation satellites are in orbit. The BE-4 engine also powers New Glenn, which exploded on a test stand in May. The US Space Force and the Artemis program also depend on Vulcan.

Why it matters: the US space sector becomes even more dependent on one company. Ars writes that the Space Force will move missions to SpaceX again. At the same time SpaceX retires Falcon 9, and Starship does not yet have certification for military payloads.

TOPIC 4

OpenAI: GPT-6 Astra proved a lattice construction of the Standard Model

SOURCES OpenAI, 16.11 OpenAI · @OpenAI, 16.11 @OpenAI

OpenAI opened a science blog. The first post is by physicist Nathaniel Craig from UC Santa Barbara. He describes how he and GPT-6 Astra proposed a definition of the Standard Model on a lattice, a grid of points in spacetime.

The problem is this. The weak force acts differently on left-handed and right-handed particles, so the theory is difficult to put on a lattice. The direct method gives extra "doubler" particles. Physicists found part of the solution decades ago, but the phase rule for the full Standard Model stayed open for more than 25 years. Craig's idea was to compare the same field configuration on finer and finer grids. He says that Astra turned the idea into a construction and "took the lead" in the detailed proof. Three lattice field theory experts read the proof and asked questions.

Craig himself is careful. He calls the result a "candidate" and writes that it is an answer "in principle", with much work still necessary before realistic simulations. The preprint came out in October, and other experts must still evaluate it.

Why it matters: formal mathematical physics becomes one more field where a model writes most of a proof, and a human sets the task and checks the result. Yesterday we wrote about Google /teamwork and Lean. The difference here is the check: Lean checks a proof by machine, and here people read the proof.

TOPIC 5

Qwen4-Max: 2.4 trillion parameters, cost per task down 62%

SOURCES Qwen, 16.11 Qwen · HN, 16.11 Hacker News

Alibaba released Qwen4-Max on the new Qwen4 architecture. The model has 2.4 trillion parameters, and 71 billion of them are active for each token. The context is 1 million tokens, and vision is native. The architecture combines Gated DeltaNet layers with its own sparse attention to make long context cheaper.

According to Qwen, the model scores 52.69 on its own Work360 benchmark, against 47.74 for Qwen3.8-Max. It uses 24-35% fewer tokens, and approximately 48% fewer on Terminal-Bench 4.0. The average cost per task on Work360 fell from $4.43 to $1.67. Max is available only through the QwenCloud API. Qwen promises open weights for smaller models of the family.

Why it matters: Chinese labs compete on the price of a task, and $1.67 in place of $4.43 matters more for agent workloads than a few benchmark points. But all the numbers come from Qwen itself, and Work360 is its own benchmark.

There are no independent measurements yet.

TOPIC 6

Bitly: strangers registered again the short links in old tweets

SOURCES Ransomnews, 15.11 Ransomnews · HN, 16.11 Hacker News

On 15.11 a Bitly short link in a 2011 Burger King tweet suddenly led to the Wikipedia list of Nazi concentration camps. Nobody hacked the account. Somebody registered the same bit.ly code again that morning, at 05:32 UTC. 18 minutes before that, the same thing happened to a link in a 2011 McDonald's tweet. A post that quoted the Burger King tweet got more than 2 million views.

Ransomnews checked 97 Bitly links in Donald Trump tweets from 2009-2013. 56 of them were created after the tweets, and 7 lead somewhere other than the original target. Bitly confirmed the problem: bad actors were able to "repurpose certain old, inactive links". The

company suspended the accounts and disables these links. According to CNN, links in old posts from Bill Gates, the FBI and the Pentagon also changed.

Why it matters: a short link is safe only while the service keeps the owner behind it. The Bitly help pages promise that a code cannot be used again, but this already broke in February 2025. For archives, documentation and old publications, a full address is more reliable.

TOPIC 7

A court did not dismiss the DRAM conspiracy case against Samsung, SK Hynix and Micron

SOURCES Litnews, 15.11 Litnews · HN, 15.11 Hacker News

On 14.11 federal judge Noël Wise in California denied in full the motion of Samsung, SK Hynix and Micron to dismiss a proposed class action. The plaintiffs say that the three companies agreed to cut the supply of conventional DRAM, for example DDR5, and move capacity to HBM for AI accelerators. HBM needs 3-4 times more wafer area per gigabyte.

According to the complaint, conventional DRAM prices rose 171.8% in the year to the third quarter of 2025. The judge found the conspiracy plausible because of a set of facts: long HBM contracts, refusals of long DRAM contracts, the closure of the Crucial brand, and no expansion of production at record prices. She also writes that these facts are not proven yet. The court expects answers from the companies by 7.12.

Why it matters: the AI boom raises prices not only for GPUs. More expensive laptops, phones and servers can be partly a result of memory makers who gave wafers to HBM. Now a court will examine this question on the merits.

TOPIC 8

A week of reports: Nvidia on 17.11, the Anthropic IPO on 19.11

SOURCES Semafor, 16.11 Semafor · Financial Times, 15.11 FT · The Wall Street Journal, 16.11 WSJ

Before, we wrote about the Anthropic valuation and orders for the shares that are 6 times larger than the offer. Now the date is known: Semafor writes about a $1.4 trillion IPO on Thursday, 19.11. Much depends on the IPO. The FT calculates that the shares of companies that build data centers and supply equipment for them lost $4.2 trillion since June. Semafor cites the NYT: right after the IPO, Anthropic expects a $10 billion investment from Nvidia.

Nvidia reports on Tuesday, 17.11. The WSJ writes that interim CEO Colette Kress will probably cut the deals where Nvidia financed compute and startups. Startups that counted on this support expect an even larger cut.

Why it matters: in three days the market gets two measurements of the AI economy. The Nvidia report will show if demand holds without money from Nvidia itself. The IPO will

show how much investors will pay for a lab that still operates at a loss.

TOPIC 9

US ground stop: Duffy blames a router, controllers do not believe him

SOURCES Common Dreams, 15.11 Commondreams · Semafor, 16.11 Semafor · The Wall Street Journal, 15.11 WSJ · HN, 15.11 Hacker News

Yesterday we wrote about the L3Harris statement on the FAA move from copper to fiber. Now Transportation Secretary Sean Duffy gave his version. He says that switches and routers "got overloaded with data" when L3Harris replaced a section of copper cable with fiber. The WSJ writes about one router in southern Arkansas. Duffy says that the $875 million AI system SMART "is not even plugged in", and he asks Congress again for $50 billion for modernization.

A Washington Examiner reporter spoke with 60 air traffic controllers, and they do not believe this version. Their argument: crews make such replacements at night in the middle of the week, not on a Friday afternoon. The stop affected approximately 1.5 million travelers. It is only the third nationwide ground stop in history. The Homeland Security Secretary says that investigators found no evidence of a cyberattack.

Why it matters: the official version changed twice in one day, and there is no independent investigation yet.

A GAO report from September warned about exactly this risk: the FAA makes modernization faster at the cost of possible outages.

TOPIC 10

South Korea: the top score on the tax accountant exam came from AI

SOURCES The Guardian, 15.11 Guardian · BBC, 15.11 BBC

A candidate with the surname Han took first place in both rounds of the tax accountant qualification exam, in April and July. In the first round he got one of 160 questions wrong and scored 397.5 of 400 points. On 15.11 he confessed that he photographed the questions with a micro camera, sent them to an AI tool and read the answers on a smartphone. He gave a written confession to the Seoul police.

The scandal broke four days before the national college entrance exam, which is on Thursday. The presidential office called an emergency meeting, and the government promises stricter control.

Why it matters: a camera and a model in a pocket make a closed-question exam in a room useless. A check at the entrance to a profession must now take into account that a model

generates an answer in seconds.

misc

- **Chlorine and Microsoft 365:** on Friday the Microsoft outage stopped email at the Phantasialand park in Germany.

  The hotel pool system sent its alarm by email, and nobody saw that the dosing unit added too much chlorine. 15 guests were hurt, and 4 went to hospital. Wdr · HN Hacker News
- **uBlock Origin:** a researcher found a UXSS in the Firefox version. A website could use the element picker to add a trusted filter and run its own JavaScript on any site. The hole is now fixed. The author says that GPT-5.6 Sol did most of the work to find it. Iter · HN Hacker News
- **"All AI-generated code should be Open Source":** the top HN post of the day, 1,115 points. The author argues that models learned from open code, and that purely generated code has no author under US law.

  Jankiewicz · HN Hacker News
- **Starship in 3D:** Ethan Mollick asked Fable and Opus to find the video of the Starship catch and turn it into a 3D scene in three.js. The result rotates in the browser. @emollick
- **Mollick on stability:** he says that of 270 jobs in the 1950 census, only one is now obsolete, the elevator operator. @emollick
- **Child deepfakes:** a Childlight report says that reports of AI-related child sexual abuse material rose from approximately 4,700 in 2023 to more than 400,000 in 2025.

  Guardian
- **Intel:** shareholders sued the company after a 29% one-day drop of the stock on 30.10. They say that Intel hid a slump in PC sales. BBC
- **Positional embeddings:** Sebastian Raschka explains how modern LLMs encode the position of a token.

  Sebastianraschka · HN Hacker News