AI DIGEST

Nvidia reinstates Huang, Claude claims a VP ≠ VNP proof

OpenAI and Cisco release Shirudo, Z.ai ships GLM-5.4, CrowdStrike buys Zscaler for $44 billion.

Tuesday, 24 November 2026

IN THIS ISSUE

1. Nvidia returned Huang to the CEO post, the stock rose 31%
2. An internal Claude model claimed a proof of VP ≠ VNP, and the proof is in Lean
3. OpenAI and Cisco released Shirudo: agents put temporary protection on Linux without a restart
4. Z.ai released GLM-5.4: the first flagship on Loop Transformer
5. The dispute about open models: LeCun and Sweeney against Anthropic
6. CrowdStrike buys Zscaler for $44 billion
7. Virginia and Oregon sued SpaceX over sexual images from Grok
8. A former OpenAI engineer recorded dozens of internal meetings
9. Ladybird is now in alpha

TOPIC 1

Nvidia returned Huang to the CEO post, the stock rose 31%

SOURCES NVIDIA, 23.11 Nvidia · CNBC, 23.11 Cnbc · Semafor, 23.11 Semafor · Foreign Policy, 23.11 Foreignpolicy · @TheRundownAI, 23.11 @TheRundownAI · HN Hacker News

Yesterday the digest wrote that Bessent does not exclude a Huang return. On Monday, the Nvidia board voted for it unanimously. Huang is again president, CEO and board member, with immediate effect. In the press release, the board writes that it got "input from the highest levels of the US government". On this basis, the board is sure that the Huang return "poses no national security risk". Colette Kress, the interim CEO, stays as chief financial officer.

According to CNBC, officials told the board at the end of last week that nobody broke national security laws. The first talk between the board and Huang took place on Saturday, and the board made the decision on Sunday. Trump told CNBC that he spoke with the board "directly and indirectly". In his words, Huang "did nothing wrong". The Nvidia stock closed at $165.05, which is 31.48% higher. Before this, it lost almost 33% since 08.11, when Huang admitted an extramarital affair.

The Rundown AI calculates that the company got back about $1 trillion of market value in one day.

The story does not end here. The internal investigation of the committee of independent directors continues, and the board promises to complete it. Nvidia created a committee for national security and export control. The company also looks for a chief compliance officer. On the same day, the Commerce Department put Megaspeed and six more companies from Singapore, Malaysia and Thailand on a blacklist. Alice Huang was the chief executive of Megaspeed. Now no US company can sell chips to them without a license. Foreign Policy called the Huang return "the biggest national security scandal in US history". On X, the reactions are polar. @levelsio and @balajis celebrate, and the journalists who investigated the case speak about pressure from the White House.

Why it matters: the board of a public company changed its decision about the CEO after a signal from the government, and it wrote this openly. This is a new precedent of governance: the White House now also assesses the risk of a CEO. For the chip market, another thing is more important. The same intermediaries that moved chips to China are now on a blacklist. Thus export control became stronger exactly when the CEO came back.

TOPIC 2

An internal Claude model claimed a proof of VP ≠ VNP, and the proof is in Lean

SOURCES Anthropic formal-math, 23.11 GitHub · HN Hacker News · @blader, 24.11 @blader · Levent Alpöge, 24.11 @__alpoge__ · HN Hacker News

Anthropic put the folder `vp-vnp` into the `formal-math` repository. It contains the paper "VP ≠ VNP over every field"

and its proof in Lean 4. VP and VNP are the algebraic analogs of the classes P and NP. Leslie Valiant stated the conjecture in 1979. The paper claims that for every field, circuits for the Hamiltonian cycle polynomial need more than m^(c·(log m)^(1/3)) gates. The same applies to the permanent outside characteristic 2. From this, VP is strictly smaller than VNP. The README states separately that this does not solve the P vs NP problem.

The main caveat is also in the README. An "internal research model Claude" found the result. AI models wrote both the paper and the formalization, and no mathematician checked the proof yet. Lean checks the proof. But no machine checks that the file `Statement.lean` says exactly what the authors claim. Thus the README gives a correspondence table in appendix Q of the paper. The repository CI checks the proof again with a second, independently written kernel, nanoda.

The HN thread got 287 points.

On the same day, Levent Alpöge commented on his proof of the Riemann hypothesis. The digest reported on it on 21.11 and 22.11. In his words, the number "250,000+ pages" was a

joke. He admits that the Lean code is "full of garbage", although it compiles, and he accepts all criticism. The post got 2.1 million views.

Why it matters: this is the second large mathematical result with Claude in one week, and both have the same format.

The model gives a proof in Lean, and human review falls behind. Lean guarantees that the proof is correct for the statement in the file. People must confirm that this statement matches the Valiant conjecture. This review is now the bottleneck. The search for a proof is already faster.

TOPIC 3

OpenAI and Cisco released Shirudo: agents put temporary protection on Linux without a restart

SOURCES OpenAI, 23.11 OpenAI · @OpenAI, 23.11 @OpenAI · @gdb, 23.11 @gdb · WSJ, 23.11 WSJ

Shirudo ("shield" in Japanese) is an open-source environment for Linux. It closes a vulnerability for the time while the team prepares and deploys a patch. The tool runs on eBPF, a Linux technology that runs small programs at points inside the kernel. An agent writes such a "shield" for a specific vulnerability. The operator first looks at how the shield affects applications, and only then turns it on.

OpenAI gives an example. Codex generated a shield that added an input check to the Base64 decoder of Java. The team connected the shield to a live Elasticsearch, and it blocked an attempt of the Java4Shell exploit. The service did not restart. The shields have a cryptographic signature. A TPM module keeps a record of which shields the system loaded and removed. Brockman explains the logic this way: to find a bug today does not mean to patch each system tomorrow.

Why it matters: AI models already find vulnerabilities faster than teams can install patches. Shirudo closes exactly this gap. The risk is also clear: an agent gets programmable access to the kernel of production systems. Thus the signature of shields and manual approval by the operator are the main parts of the design.

TOPIC 4

Z.ai released GLM-5.4: the first flagship on Loop Transformer

SOURCES Z.ai, 23.11 Z · Simon Willison, 23.11 Simon Willison · HN Hacker News

The Chinese lab Z.ai released GLM-5.4. This is the first large model of the lab on the Loop Transformer architecture. In a usual transformer, each layer runs once per pass. In GLM-5.4, a group of 28 layers of the GLM-5 base model runs twice, and the second pass refines the result of the first. This adds no parameters. The model has 750 billion parameters, and Z.ai compares its level with models of 2-3 trillion.

Z.ai did not train the loop from zero. It continued the training of a finished base model, because loop models from zero fall behind at the same budget.

On a 30B model, the loop gave a gain at only ~10% more compute. Mathematics, scientific reasoning and long context gained the most. Answers on GPQA-Diamond and MMLU-Pro became 10-14% shorter. Part of the reasoning moved from tokens into the depth of the loop. The speed is up to 200 tokens/s off peak and 70-80 at peak. The lab itself names the weak point:

the model gets stuck in too long "reasoning marathons".

The weights are closed for now. Z.ai promises to open them after "safety refinement". Simon Willison notes that this is already the second loop model in one week. On 20.11, Thinking Machines Lab released Inkling 2 Mini, which runs the middle layers three times.

Why it matters: two independent labs moved to the loop in one week. Thus the depth of compute becomes a separate axis of scaling, next to model size. This is important for those who run models themselves: the loop gives more reasoning without more memory for weights.

TOPIC 5

The dispute about open models: LeCun and Sweeney against Anthropic

SOURCES @ylecun, 22.11 @ylecun · HN Hacker News · @TimSweeneyEpic, 23.11 @TimSweeneyEpic · @amasad, 24.11 @amasad

The cause was a post that retells a speech of Dario Amodei in Congress. According to this retelling, Amodei said that open models lead to a very dangerous situation. Yann LeCun answered sharply. He writes that since GPT-2, Amodei exaggerates the dangers so that the government bans open models and the control stays with him. LeCun compared this with the ban on printing in the Ottoman Empire. The post got 1.3 million views, and the HN thread got 1325 points and 736 comments.

Tim Sweeney, the CEO of Epic Games, went further. He writes that Anthropic and partly OpenAI offer the government a deal:

give us a monopoly, close open source, and we will do what you say. Amjad Masad, the CEO of Replit, told on the same day about a meeting with Vice President Vance. In his words, US allies more and more often choose Chinese open models. Masad thinks that America can win the race of frontier models and lose the global ecosystem.

Why it matters: the dispute takes place at the moment when Washington decides how to regulate AI. The safety argument and the competition argument now have equal force. The context is also notable: GLM-5.4 came out on the same day, and its weights are closed for now.

TOPIC 6

CrowdStrike buys Zscaler for $44 billion

SOURCES Semafor, 23.11 Semafor · FT, 23.11 FT · WSJ, 23.11 WSJ

CrowdStrike buys Zscaler for almost $44 billion in stock. Semafor estimates that this can be the largest deal in the history of cybersecurity. FT, WSJ and Semafor confirmed the deal. Semafor links it to the wave of outages at airlines, banks and government agencies in the US. Because of these outages, the White House fired its cybersecurity director last week. Experts explain most outages with mass patches of old vulnerabilities by AI teams. The digest reported on this on 18.11 and 23.11.

Why it matters: the two largest "pure" cybersecurity companies merge, because the market expects a growth of demand for protection. For customers, this means fewer independent suppliers. FT and WSJ keep the details of the deal behind a paywall.

TOPIC 7

Virginia and Oregon sued SpaceX over sexual images from Grok

SOURCES Ars Technica, 23.11 Ars Technica · NYT, 23.11 NYT

Virginia and Oregon are the first states to sue SpaceX. SpaceX bought xAI in February, so it is now responsible for Grok. The state attorneys general say that Grok still sexualizes real photos, mostly of women and children. Since the time of "spicy mode", the model generated intimate images without consent and child sexual abuse material. The Musk companies said that they put new safeguards in place. The states call these statements "deceptive". In their opinion, the company kept such content on purpose to take a larger market share.

Why it matters: this is the first state-level lawsuit against image generation itself. If the court accepts the logic of the lawsuit, the responsibility for the output of a model will fall on the company that deployed it.

TOPIC 8

A former OpenAI engineer recorded dozens of internal meetings

SOURCES NYT, 23.11 NYT · FT, 23.11 FT

Abdulrahman Alfozan, 35 years old, is a former OpenAI engineer. Prosecutors say that he put a hidden recording device in the office. In a court filing on Monday, they describe hundreds of audio files on his computer. The recordings include meetings with OpenAI customers, meetings about AI safety and a meeting with representatives of a foreign government. The files are in folders by date.

There is one charge: illegal interception of conversations. Alfozan pleads not guilty. A few days ago, prosecutors told the judge that they have no evidence of espionage or theft of trade secrets. His lawyer writes that most recordings are "noise and silence".

Why it matters: labs discuss weights, plans and talks with governments in closed meetings. One insider with a cheap device collected this over months. The internal security of labs is now as important a topic as the safety of models.

TOPIC 9

Ladybird is now in alpha

SOURCES Ladybird, 23.11 Ladybird · Simon Willison, 23.11 Simon Willison · HN Hacker News

Andreas Kling announced the alpha of the Ladybird browser for macOS and Linux. The browser has its own engine. The team wrote it from zero, without a fork of Chromium, WebKit or Gecko. The alpha works as a channel with automatic updates.

Some sites do not work yet, and the browser sometimes crashes. The team does not guarantee security yet and asks hackers to break the browser. For developers, there is WebDriver and the Firefox DevTools protocol. The HN thread got 1560 points, the most in the day after Huang.

Why it matters: two engines, Blink and WebKit, hold almost the full web, plus Gecko with a small share. A new independent engine in the hands of users is the first real alternative in many years.

---

misc

- **FSNotes:** Apple removed the open-source notes app FSNotes from all European App Stores after 9 years. The author, a Ukrainian sole proprietor, submitted documents under the DSA and from the state register of Ukraine. Since 05.11, Apple did not answer even once. GitHub
- **Arm:** Arm lawyers demand to close the unofficial AArch64 instruction reference arm.jonpalmisc.com. The site will disappear on 04.12. The HN thread got 914 points. Hacker News
- **Nuclear weapons in orbit:** according to the German NDR, WDR and Süddeutsche Zeitung, Russia can launch a satellite with nuclear weapons as early as December. An explosion in orbit would disable thousands of satellites. FT writes the same with reference to Germany.

  Ars Technica
- **AI toys:** more than 150 child protection groups ask Amazon and eBay to stop the sale of AI toys until a safety check. A plush bunny "from 3 years" connects to ChatGPT, although the ChatGPT terms require an age of 13 years.

  Guardian
- **OpenWorks:** OpenAI pays local AI experts up to $8,000 per month to help small businesses in 20 US cities.

  @gdb @gdb

- **Quantum circuits:** @GoogleDeepMind writes that teams of Gemini 4 agents built quantum circuits that make future simulations up to 10 times cheaper. The team checked the optimizations in Lean and opened the code.

  @GoogleDeepMind
