Nvidia Reinstates Jensen Huang After Trump Joins Board Call; Claude Model Credited With VP VNP Proof

AI High Signal Digest

2026-11-24

Nvidia Reinstates Jensen Huang After Trump Joins Board Call; Claude Model Credited With VP VNP Proof

By AI High Signal Digest • November 24, 2026

Jensen Huang returned as Nvidia CEO after direct White House involvement, setting off a fight over export controls and corporate governance. Also covered: a second major AI math result, more evidence of users moving to Gemini 4 Neon, and GLM-5.4 joining the looped-transformer wave.

Jensen Huang is back at Nvidia, and the White House had a hand in it

Nvidia's board reinstated Jensen Huang as president, CEO and board director [1]. According to The Rundown, the vote was unanimous and Nvidia gained about $1T in market value at the open [2]. A reporter's source says President Trump dialed into the board meeting to argue for Huang's return. White House communications director Steven Cheung confirmed Trump "was totally involved" and said "We have to protect our geniuses" [3]. In a later interview, Trump said that if Huang hadn't been reinstated, Xi might have asked him to lead a Chinese company [4].

The fight now turns to China policy: - **Congress:** Sen. Kevin Cramer said the national-security risk now falls on the US military. He urged Congress to "immediately pass the strictest export controls possible" [5]. Rep. Jeffries called the reinstatement "egregious" [6]. - **Megaspeed listing:** Chris McGuire points out that Commerce added Megaspeed to the Entity List but did not apply the foreign-produced direct product rule. Without that rule, the listing doesn't block shipments of Taiwan-made Nvidia chips, even though every other AI-chip-

related Chinese listing included it [7]. - **Governance:** OpenAI's Joshua Achiam argued that both the exit and the return were "hasty, ill-conceived" and "a farce of corporate governance." In his view, the alleged wrongdoing needed careful investigation, not an abrupt ouster followed by an abrupt reversal [8].

A second landmark math result: VP VNP

Last week's Riemann proof now has a companion. According to alphaXiv, an internal Claude model is credited as author of a 70-page paper proving Valiant's conjecture, the algebraic analogue of P vs NP, and the Lean proof passed automated checks [9]. The claim: over the real numbers, computing the permanent eventually takes more arithmetic steps than any fixed polynomial in matrix size, so it has no shortcut like the determinant's [9]. The formalization sits in Anthropic's formal-math repository. Stella Biderman noted that two years ago this would have been a candidate for math news of the decade, yet it wasn't even the week's biggest [10]. No independent mathematical review has been reported yet.

OpenAI staff defend the perplexity chart

This continues last brief's debate over OpenAI's chart of pretraining progress measured by perplexity (lower means the model predicts text better). Critics had suggested the evaluation set was skewed by AI-written code. OpenAI researchers rejected that. John Hallman said the team uses "a fixed eval across time" and guards against model-generated code [11]. Will DePue said the eval is "held out, frozen, tokenizer agnostic" [12]. Boaz Barak called two critical analyses "nonsense." He argued that math progress refutes the claim that models are good at software engineering but not research [13].

Attempts to put other labs on the same plot ran into trouble. teortaxes conceded that Inkling's bits-per-byte figures (a tokenizer-independent loss measure) aren't directly comparable to OpenAI's [14]. Lucas Beyer noted that building a clean code eval now means hand-writing code, carefully deduplicating old GitHub commits, or buying an exclusively licensed dataset [15].

Users are switching to Gemini 4 Neon

Aravind Srinivas says Neon's rise on Perplexity is organic, since Perplexity doesn't route traffic to it by default. The growth comes largely from users switching away from Opus 5.5 and Sol 6.2 [16]. A day later, Neon was up 18.8% day-over-day in Ask and 45.4% in Computer, and Computer had already passed the previous day's full volume with 7 hours left [17]. On OpenRouter, users spent more on Google models than on OpenAI's last week, putting Google over 50% [18]. Neon ranks #3 on FrontierSWE at 62.4%, on par with Opus 5.5 [19].

Arena's five-month Agent Arena trend, built on 3.2M sessions, shows Anthropic

extending its lead and Google moving past OpenAI [20]. Google DeepMind also reports that teams of Gemini 4 agents designed quantum circuits that cut future simulation costs by up to 10× [21].

Looped transformers spread

Looped transformers reuse the same layers several times per token to get more depth without adding weights. Z.ai's GLM-5.4 is its first one: it loops existing GLM-5 layers, adds vision, and extends coding to video, games, 3D shaders and CAD. API access waits on a safety evaluation [22]. The Turing Post reports it loops 28 middle layers twice at 750B scale and claims 2T–3T-class intelligence [23]. On BugBunny's VulnPR-100, it found 47 of 100 vulnerabilities, 15 more than GLM-5.3. DeepSeek V4.1 Flash 1017 found 54 at about 14× lower cost [24].

Thinking Machines' research model lets users choose 1–5 passes through its middle layers. On the lab's internal destructive-actions eval, the rate fell from 26.6% at one pass to 3.1% at four [25]. LoopRes, which adds looping to existing pretrained LLMs, reports scores up to 37.7% higher than SFT across six LLMs [26]. Nathan Lambert warned that looped architectures are a first step toward chain-of-thought that humans can no longer read [27].

Security and agent tooling

- **OpenAI cyber research:** GPT-6 Cyber helped validate 948 Linux kernel bugs that could grant local root access and generated patches, with fixes sent upstream through maintainers [28]. With Cisco, OpenAI released Shirudo, an open-source runtime that lets agents write eBPF controls to shield systems from exploits until patches are ready [29].
- **Claude Ship (beta):** Reviews PRs, works on fixes, and approves changes that team policy allows [30]. Admins can turn it on for a whole repo [31].
- **Anthropic data retention:** Zero data retention with Enterprise Frontier Safeguards entered public beta. Encrypted copies of API traffic go to the customer's own cloud storage, where automated checks flag misuse [32].

Robotics

Embo came out of stealth with a $110M seed round led by a16z, with NVentures among the investors [33]. Its team is behind Dreamer, VideoGPT and Diffusion Forcing. Its robot brain pretrains on internet video, then practices motor skills in imagined scenarios [33]. RoboCurve reports that GPT-6.2 Astra scored 97% across 120 real-world RoboDojo-RC Tier 1 trials, and it is retiring the benchmark [34].

Benchmarks and measurement

- **Epoch reviews:** Epoch reviewed four benchmarks: one Verified, three Flawed [35]. ARC-AGI-3 was rated Flawed partly because it scores action efficiency without telling models to be efficient [36]. Agents' Last Exam had severe errors in at least 44% of the tasks Epoch sampled [37].
- **Harness comparisons:** Artificial Analysis found that no coding harness is best across models. GPT-6 Luna scores highest in Codex, where GLM-5.3-Flash ranks last [38]. Codex's shell-tool definition alone switched GLM-5.3-Flash into terse reasoning in 12 of 12 samples. A classifier labels 98% of that reasoning as resembling OpenAI's open-weight gpt-oss models [39].
- **RL and hallucinations:** A Cornell paper finds that reasoning-style RL on knowledge tasks makes models memorize hallucinations that got rewarded. Recall of those rose 4.83%, versus 0.17% for hallucinations with no training signal [40].

Also notable

- **ERNIE 5.5 Flash Preview (Baidu):** Weights released under Apache 2.0; 280B total / 14B active parameters, 1M context. Post-training is ongoing [41].
- **Ling-3.2-air (Ant Ling):** Open-weight, 124B total / 5.5B active. AutomationBench-AA rose from 15.7% to 61% [42].
- **GPT-6.2 Sol:** Rolling out in ChatGPT's Chat tab for paid tiers [43].
- **OpenWorks:** OpenAI pays local AI helpers up to $8,000 a month to work with small businesses in 20 US cities [44].
- **River API:** Passed $100M annualized run rate five months after launch [45].
- **Microsoft Rocket:** Microsoft AI detailed its in-house RL infrastructure for runs lasting weeks to months [46].

Sources

1. X post by @nvidianewsroom
2. X post by @TheRundownAI
3. X post by @rogoswami
4. X post by @weijia
5. X post by @SenKevinCramer
6. X post by @RepJeffries
7. X post by @ChrisRMcGuire
8. X post by @jachiam0
9. X post by @askalphaxiv
10. X post by @BlancheMinerva
11. X post by @johnohallman

12. X post by @willdepue
13. X post by @boazbaraktcs
14. X post by @teortaxesTex
15. X post by @giffmana
16. X post by @AravSrinivas
17. X post by @AravSrinivas
18. X post by @OpenRouter
19. X post by @ProximalHQ
20. X post by @arena
21. X post by @GoogleDeepMind
22. X post by @Zai_org
23. X post by @TheTuringPost
24. X post by @Zaddyzaddy
25. X post by @karenzzhou
26. X post by @oahzxl
27. X post by @natolambert
28. X post by @OpenAI
29. X post by @OpenAI
30. X post by @ClaudeDevs
31. X post by @ClaudeDevs
32. X post by @ClaudeDevs
33. X post by @danijar
34. X post by @robocurve
35. X post by @EpochAIResearch
36. X post by @EpochAIResearch
37. X post by @EpochAIResearch
38. X post by @ArtificialAnlys
39. X post by @ArtificialAnlys
40. X post by @dair_ai
41. X post by @Baidu_ERNIE
42. X post by @AntLingAGI
43. X post by @ChatGPT
44. X post by @TheRundownAI
45. X post by @river_ai_inc
46. X post by @MicrosoftAI