Nvidia's board shuts the door on Huang's return as $5B Isomorphic and $700M Fractile rounds lead the funding tape

VC Tech Radar

2026-11-18

Nvidia's board shuts the door on Huang's return as $5B Isomorphic and $700M Fractile rounds lead the funding tape

By VC Tech Radar • November 18, 2026

The fight over Nvidia's leadership hardened, OpenAI shipped GPT-6.2 to mixed reviews, Perplexity disclosed unit economics behind $1B ARR, and large rounds landed in AI drug design, inference chips and finance agents.

Nvidia: the board rules out Huang, and his backers turn to shareholders

Jensen Huang said publicly that he made "serious mistakes" in his personal life, created problems for the board, and is cooperating with its independent investigation. He added that he would return as CEO if the board decided that served shareholders, employees, the company "and America" [1]. Hours later the chair said the board "does not intend to reinstate" him, according to Polymarket [2]. A campaign backing Huang says the board told the earnings call that he resigned voluntarily and declined to answer questions about him or the investigation [3]. A Semafor scoop says the board is split, with a three-director faction "vehemently" against his return [4].

Garry Tan answered "15% shareholder vote coming" [5]. According to a Grok summary of the bylaws, holders of 15% of voting power for at least a continuous year can call a special meeting. Removing directors there takes a majority with cause, or 66 % without [6]. That is a high bar, not an imminent outcome. Separately, the board announced a "$150B increase" to its buyback. Jason Calacanis argued the money belongs in data-center buildout [7].

Sam Altman said he didn't know enough to judge whether Huang should return, and that it is "hard to imagine the AI revolution without Jensen." Asked about a reported $105B Nvidia guarantee for OpenAI's Ohio data center, he said people "still trust their guarantees very much" [8]. For anyone exposed to Nvidia-financed infrastructure, the governance fight is now open-ended.

OpenAI ships GPT-6.2, and early evals split

OpenAI released GPT-6.2 Astra and Sol, describing them as "super efficient" [9]. Altman pointed to token efficiency and "Sol ultrafast" [10], but conceded the models still need work to be "nice to work with" [11]. Third-party results so far:

- **Perplexity (WANDR benchmark):** Astra scored 0.783 at about $12.54 per task, up 0.095. Sol scored 0.746 at about $2.65, up 0.065 [12].
- **Bindu Reddy:** Astra excels at browser use but is "nerfed on agentic tasks," and Sol trails Sol 6.1. Her read is that OpenAI is trading performance for cost [13].

Perplexity puts numbers behind $1B ARR

Aravind Srinivas said Perplexity crossed $1B in annual run rate. He claimed it burned 5x less to get there than "the company that's going to IPO this week" [14]. Its own fast-search model handles 80–90% of traffic on its own compute, which he estimates saves $100–200M a year [14]. Gross margin is 56% on paid users. Counting free users it is 17% in Q3 2026, up from −219% in Q4 2025 [14].

The products follow the same orchestration thesis:

- **Orchestrator API:** 45+ open and proprietary models behind one API, with failover and unified billing. Routing is designed to keep the KV cache intact. In one example, sending Fable 5.2 refusals to GPT-6.1 Sol raised Terminal-Bench 4.1 results from 38 to 41 of 63 tasks [15].
- **Multiagent:** runs parallel mixed-model agent teams. Perplexity says diverse teams beat homogeneous ones on both performance and cost [16, 17].

Deals worth noting

- **Isomorphic Labs, $5B.** Demis Hassabis announced the funding for AI-first drug design [18]. RAAIS describes it as a Series C [19].
- **Fractile, over $700M Series C** (the transcript spells it "Fractal") [20]. Its thesis is that LLM inference is memory-bound. It stacks DRAM tightly on logic and claims 20x the bandwidth per chip of an HBM-based die [20]. It expects customer workloads running in its racks before the end of next year, and says it will never compete with customers at the model layer [20].

- **Model ML (YC W24), $250M Series B at $1.5B.** It runs AI agents inside banks' and consultancies' own systems. Q3 ARR grew 12x year over year and daily active users 26x [21].
- **River AI, $1B.** Igor Babuschkin's company sells software that lets enterprises post-train and serve models on their own hardware. Backers are General Catalyst, AMP PBC, Nvidia and AMD [22].
- **First Person Labs, $10M seed led by a16z.** Its product, Second, keeps a continuously updated model of the user on the device and serves that context to agents the user approves [23].
- **Alta Ares** won 6 of 16 lots in NATO's counter-drone framework, more than any other supplier. Allied countries can now buy without a new tender [24].
- **Auctor (YC P25)** raised a Series A from Sequoia for AI-led enterprise software implementation. Usage rose 5–10x after it rebuilt the product from scratch [25].

Venture bar and market structure

SaaStr argues a company now needs to grow from $1M to $100M ARR within five years to raise venture capital. That is 151% compounded a year, which makes T2D3 roughly the minimum [26]. Citing Carta, the AI Series A bar is about $3.5M ARR, and AI took over 60 cents of every venture dollar in Q1 2026 [26]. As of May there had been zero software IPOs among 32 listings [26].

The counterpoint is Stripe's index of non-AI B2B revenue, which is up more than 30% year over year, with firms under a year old growing fastest [27]. But the sample is mostly self-serve, card-on-file companies, not the enterprise names that were sold off [27].

Newcomer's Cerebral Valley summit added three data points:

- **Open-weight models:** they went from about one basis point of Factory's tokens at the start of 2026 to "way into double-digit percent." At many of its customers they now carry more than half of all tokens [22].
- **Valuations vs. returns:** the AI 100 is worth $3.6T, with OpenAI and Anthropic 55% of that. Bessemer estimates $8T of capex needs close to $6T in annual revenue to pay off [22].
- **Audience survey:** 79% see a bubble and 67% expect it to pop in 2027. The same room expects Anthropic to be worth $1–2T after its first trading day [22].

Agents, autonomy and security

- **Replit Missions.** Replit launched this agent-swarm product. It reports NumPy optimization speedups of up to 4x with one agent, 17x with 16 and 19x with 32 [28]. It also shipped DataCat, a free 1:1 rebuild of an observability platform it says costs $83,700 a year [29]. That is a direct

shot at feature-layer SaaS.

- **Autonomous exploitation.** Google Threat Intelligence Group and Mandiant report the first LLM framework they have seen fully automate vulnerability research and exploitation with no human in the loop. It diffs patched against unpatched web apps, then writes and runs exploits [30].
- **Waymo off freeways.** California's DMV ordered Waymo off state freeways, citing an "imminent hazard to public safety" [31]. A Reddit commenter says the cause was repeat entries into construction zones after a May fix, with no injuries, and that Waymo has paused freeway rides in all 15 markets [32]. Calacanis, who says he is invested in every AV provider, warns that the industry is moving too fast [33].
- **Qwen4-Max is API-only.** It has 2.4T total parameters, and no weights were released [34]. A commenter's summary of draft Chinese export rules says models with 2T or more parameters would need an export license. Newly open-sourcing one would require a "file-and-go" filing, and the rules are not yet final [35]. That would explain the release, but it hasn't been confirmed.

Sources

1. X post by @JensenHuang
2. X post by @Polymarket
3. X article by @TheJensenLetter
4. X post by @ReedAlbergotti
5. X post by @garrytan
6. X post by @grok
7. X post by @Jason
8. Should Jensen Huang Come Back? Sam Altman Answers
9. X post by @OpenAI
10. X post by @sama
11. X post by @sama
12. X post by @perplexity_ai
13. X post by @bindureddy
14. Perplexity CEO and Founder Aravind Srinivas at Silicon Valley & The World | Semafor
15. X post by @denisyarats
16. X post by @AravSrinivas
17. X post by @AravSrinivas
18. X post by @demishassabis
19. X post by @raais
20. Fractile's $7B Bet on Chip-First Frontier AI | Lightwork
21. X post by @ycombinator
22. The Top Takes from Newcomer's Cerebral Valley AI Summit
23. X post by @a16z

24. X post by @nathanbenaich
25. X post by @ycombinator
26. How Fast Do You Really Have to Be Growing to Raise Venture Capital Today? $1M to $100M ARR in 5 Years, Max. But VCs Are Hoping for 2-3 Years
27. Stripe's New Index Shows Non-AI B2B Revenue Up 30%+ Year Over Year Through the SaaSpocalypse
28. X post by @pirroh
29. X post by @Replit
30. r/artificial post by u/jomic01
31. X post by @scottbudman
32. r/artificial comment by u/michaelTM_ai
33. X post by @Jason
34. r/artificial post by u/Delicious_Character6
35. r/artificial comment by u/CrimsonBolt33