SaaStr says agents now pick its software as HubSpot refounds and the fight over Nvidia's ouster continues

VC Tech Radar

2026-11-16

SaaStr says agents now pick its software as HubSpot refounds and the fight over Nvidia's ouster continues

By VC Tech Radar • November 16, 2026

SaaStr's operator report on HubSpot's "Customer Harness" pivot shows how agents can drop a vendor without anyone cancelling. Also covered: the continuing dispute over Jensen Huang's exit, model weights that point to Blackwell training at four Chinese labs, Meta's lessons from launching Muse, and research and security signals investors should weigh.

SaaStr: agents now choose the software, and they can drop vendors without cancelling

Brian Halligan says he and Dharmesh Shah are refounding HubSpot as the "Customer Harness Company" [1]. SaaStr's write-up describes what that looks like at a company with 3 humans and 20+ agents in production [1]:

- **Nobody on the team logs into the CRM.** Agents read and write it through MCP and the API, and it stays the system of record [1].
- **Marketing automation is gone.** In 60 days, marketing to a 450,000+ list moved to agents. They rebuild templates for each send and write emails personalized to each recipient. Consent and deliverability are bought as API services [1].
- **Agents are picked one job at a time, on evals, not bundled from the platform vendor.** As SaaStr puts it, "owning the data can just mean someone else's agent calls your API" [1].
- **Churn without a cancellation.** A core vendor went down for 48 hours, the agent routed around it, and the team never went back. "Usage just

falls to zero." The reverse also holds: a winning agent gets sticky, because every switch means new evals and new context [1].

For SaaS diligence, this hits seat-based pricing and login-based health scores directly [1]. It is one operator's report, but it describes a way to churn that renewal data would not catch.

Jensen Huang's exit: the fight is now about the reporting

Tim Sweeney says the articles have "backtracked" to "Jensen did lawful things that powerful people didn't like." He argues the original claims appear to come from a single competing Big Tech lobbyist, and he wants due process instead of press takedowns [2]. Jason Calacanis guesses that White House China hawks leaked the memo [3], and calls it absurd that Huang would knowingly break export controls [4]. Ed Zitron asserts that Huang "obviously helped a massive illegal smuggling operation with Megaspeed" [5]. All of these are opinions. None is a finding.

On the economics, Bindu Reddy calls Nvidia's 80% margins unsustainable and says financing neoclouds "propped up" Nvidia's own demand [6]. She predicts AI will "fizzle" rather than burst, with private valuations falling 10–100x [7]. @goodwin_ml, quoted by Harry Stebbings, makes the opposite case: GPUs that are powered on are fully utilized, there are "no dark GPU fleets" [8], and frontier model providers are building first-party chips to protect their margins [9].

China's open models: Blackwell fingerprints, and a Huawei-only release

An analysis finds numerical traces in six published models from DeepSeek, Moonshot, MiniMax and Ant Group that match Blackwell chips, where those labs' earlier models matched Hopper. The method was right on 24 of 24 known cases [10]. The report itself says it cannot show where the chips were or whether export controls were broken, and Blackwell can be rented outside China [10]. The newest Qwen, GLM, Tencent, Xiaomi and StepFun models still read as Hopper [10]. With chip controls under scrutiny, investors should treat this as evidence, not proof.

Meituan released LongCat-2.5 under the MIT license. It has 1.6T total parameters with about 48B active. On Terminal Bench v2.1 it scores 88.4 against Claude Opus 5's 89.1, but on v4.0 it trails 38.5 to 53.9. The model card documents deployment only on Huawei NPUs, with no Nvidia path yet [11].

Meta Muse: what launching a consumer agent took

On Lenny's Podcast, Mona Sarantakos said Muse came out of a personal-agent experiment she ran over a weekend [12]. It took about seven months to launch,

while some competing assistants took a couple of months [12]. Most of that time went to reliability, safety and security. Adding a browser opened new attack vectors, and the team's internal rule was "trust is a must" [12].

Shopping turned out to be a top use case and home automation underperformed, so the team rebuilt its evals around shopping [12]. Agent swarms produced multi-turn evaluation data that Meta did not have outside coding [12]. The system prompt and behavioral spec must be written by hand [12]. Activation is still "the number one problem": word of mouth works best, and later users need different onboarding than early adopters [12].

Research worth watching

- **Editing weights directly (MIT).** GPT-6 Astra edited about 0.016% of a small Qwen model's weights, with no training, and its math score went from 7.9% to 48.1% [13]. The caveats matter. Most of the gain came from fixing a repetition habit. Answers hitting the length limit fell from 93% to 33.1%. There was no established gain on Gemma, and independent reproduction of the headline result is pending [13].
- **Gödel's notebooks.** Claude produced a 650-page explanation of Gödel's notebooks in a couple of days, though it had to be checked page by page. It found Gödel was "one step away" from a proof of the independence of the axiom of choice [14].
- **LeCun.** He says superintelligence will not arrive within two years and that AI remains "terrible" at understanding the physical world. He is working on this at a new company the transcript renders as "Emmy Labs" [15].
- **Andrew Ng.** He says he will no longer hire engineers who only implement specs, because user feedback, not code, is now the bottleneck [16]. AI projects often stall around 80% until teams switch to golden datasets and rigorous evals [16]. Skill with coding agents does not yet appear in job descriptions [16].
- **Hassabis.** He calls AI that improves itself "an extremely useful technique" that "carries the most risk," and says the race toward it is outpacing safeguards [17].

Security and diligence

According to Cursor, a stolen Fireworks AI key was embedded in a trojanized Claude Code VS Code extension. The failure was in the supply chain, not the model [18]. The post points to static keys, no per-call endpoint verification, and no way to detect a credential swapped mid-session [18]. That points to a gap in tools for securing agent credentials.

Investing in AI argues that vibe coding creates "Recovery Debt": hours saved writing code are borrowed against future time to recover from failures [19]. The author would watch observability, testing, reliability and incident-response com-

panies. In diligence, ask whether mean time to recovery rises as the share of AI-generated code rises [19].

Briefly

- An analysis by Hill estimates that an Anthropic IPO could make about 4,250 of its 5,300 employees millionaires. That assumes a $1.5T valuation, above the roughly $1T target [20]. Clem Delangue answered that concentration of power and wealth is AI's biggest risk without much stronger open source [21].
- Allie K. Miller notes that ChatGPT ads emphasize consumers, while most enterprises she speaks with "still lean Anthropic" [22].
- Apple's health-AI study with OpenEvidence can enroll up to 200,000 US adults. Answers link to clinical sources, and Apple says study data will stay out of AI training [23].

Sources

1. HubSpot Is Refounding as the "Customer Harness Company." Here's What That Means for Us, Right Now, in November 2026
2. X post by @TimSweeneyEpic
3. X post by @Jason
4. X post by @Jason
5. X post by @edzitron
6. X post by @bindureddy
7. X post by @bindureddy
8. X post by @HarryStebbings
9. X post by @HarryStebbings
10. r/artificial comment by u/michaelTM_ai
11. r/artificial post by u/javimbuilds
12. Meta's Muse: How we launched the most successful consumer product this year | Mona Sarantakos
13. r/artificial post by u/winer666
14. r/Futurology comment by u/Gari_305
15. Yann LeCun Unfiltered: The Godfather of AI on Life, Logic, and the Future
16. BayLearn2026 05 AndrewNg HD 1080p
17. r/Futurology comment by u/Gari_305
18. r/deeplearning post by u/No-Conclusion3720
19. Asymmetric Fragility: Why Vibe Coded Software Breaks Harder and Takes Longer To Recover
20. X post by @unusual_whales
21. X post by @ClementDelangue
22. X post by @alliekmiller
23. r/artificial post by u/andrewaltair