OpenAI proposes shared training limits as Carta shows a market with record prices and fewer deals

VC Tech Radar

2026-11-23

OpenAI proposes shared training limits as Carta shows a market with record prices and fewer deals

By VC Tech Radar • November 23, 2026

OpenAI's chief scientist proposed cross-lab caps on training, tied to a perplexity metric. Carta's Q3 data shows investors paying more for fewer startups. Stripe published early agentic-commerce numbers, and investors argued over where value sits in agent harnesses.

OpenAI asks labs to agree on limits before training further

OpenAI's chief scientist, Jakub Pachocki, published internal log-perplexity figures and proposed that frontier labs agree on common limits they would not cross until they can monitor model behavior better. It is a proposal, not an agreement [1]. Sam Altman quoted the core of it: shared evaluation and red-teaming norms, plus a first limit on development, "measured by some key metric we can track across labs," that stays in place until a concrete milestone is reached, such as a successor to chain-of-thought (CoT) monitors [2]. He called it "the most important thing OpenAI has published in awhile" [3] and linked it to the safety commitments labs made under the White House Accords [4].

What makes the proposal unusual: - **It would cap training, not just releases.** Pachocki would gate training runs themselves on a shared pretraining measure. He expects further transformative capabilities at 0.68–0.7 log-perplexity, and argues pretraining could accelerate sharply as early as 2027 once AI contributes to deep-learning research [5]. - **The metric is weak as a yardstick.** The figures come from a held-out snapshot of OpenAI's own codebase, not a public benchmark. Pachocki himself notes that labs train on different datasets, and that chain-of-thought can make a model more capable without lowering this loss [1]. One commenter pointed out that lower loss hasn't stopped complaints that GPT-6.2 doesn't follow instructions [6]. - **Critics question**

the timing. Bindu Reddy says OpenAI is talking about safety "when their models are losing traction" [7].

Why it matters to investors: if labs adopt the idea, training compute at the frontier would depend on progress in monitoring. That would make interpretability and monitoring tools a gating input rather than a nice-to-have. Ramez Naam offers a sober counterweight. He argues that new OpenAI data show METR-style task-length estimates overstate real capability by 20–100x, and that the data show faster AI usage inside the company but "not a clear acceleration of results" [8].

Carta's Q3 data: record prices, fewer companies funded

Startups on Carta raised $49.3B in Q3 and more than $100B over the last two quarters, the most for any six-month stretch since late 2021. Total deal count fell 21% year over year [9]. At seed, $4.1B was deployed (up 21%) across 426 rounds (down 27%). Median seed valuation is $19.7M, up 84% from 2021, but that median counts only companies that actually closed a round [9].

What matters for an early-stage portfolio: - **Series A is the choke point.** 32% of Series A rounds are bridges, up from 29% in 2025. About 73% of 2022 seed companies had not raised a priced A after 15 quarters [9]. - **Money is rotating at the early stage.** AI's share of seed and Series A dollars fell from 48% to 44%, the first decline since 2022, while hitting a record 57% at Series B and later. Hardware raised nearly $34B through Q3, up 125%, and leads on price at seed ($27.2M median pre-money) [9]. - **Exits aren't keeping pace.** H1 2026 saw 6.7 M&A exits per 1,000 companies, tied for the decade low. Record Series B/C prices create preference stacks that fewer acquirers can pay above [9]. Participating preferred appeared in more than 10% of Q3 rounds, the highest rate in six years [9].

Jason Lemkin adds a buyer-side warning. For two years many AI deals carried an unspoken backstop: Nvidia, or a company close to it, would buy the team or fund the next round. He doesn't think Nvidia's newly promoted CEO will keep doing that, and Salesforce may spend 2027 digesting this year's acquisitions. His advice: "If your exit math assumed Nvidia or Salesforce as the buyer, redo it without them" [10].

Agentic commerce: first numbers from Stripe

Stripe reports that daily spend on Link Agent Wallet, which handles payments for personal agents including Muse, Instinct and Grok, was nearly 80x higher in mid-November than at the start of September (seven-day rolling average). Spend doubled in the last month alone [11]. Last week, travel and lodging took more than 40% of volume and retail about a third. More than half of consumers who made one agentic purchase came back for a second [11]. Stripe gives no absolute dollar figures.

Lemkin reports demand on the B2B side too. Three of his portfolio companies sell agents that replace front- and back-office staff at SMBs for $500–$1,000 a month, and close the same day once they get a prospect on the phone. Meanwhile SMBs are putting off HubSpot and Monday purchases [10]. He reads OpenAI's purchase of the travel agent Adessia as a sign that vertical agents will become features inside horizontal agents, and expects "20 more deals like it" [10].

Where value sits in agent harnesses

The "harness" is the code around a model: tools, prompts, memory, sandboxes. Investing in AI argues that harness logic which adds intelligence loses value with each model release, because labs train on harness usage data and absorb the skills. It argues that infrastructure, trust/security and efficiency are durable, because "a model can't train its way into being your audit log, your permissions system, or your cloud bill" [12]. Its screening test for an agent or harness company: does the product get better or need rebuilding when the next model ships [12]? The author discloses that they run an inference-efficiency company [12].

Two data points this period: - **Workflow-specific tuning pays, but doesn't travel.** Hugging Face trained LoRA adapters on Qwen3.8-27B using real agent workflows. Inside the pi coding agent, first-attempt Terminal-Bench 2.1 success rose from 39% to 55% and token use per task roughly halved. There was no clear transfer to business-automation tasks [13]. - **Governance is a layer of its own.** Harrison Chase argues agent platforms need observability/evals and governance, not just a runtime. Governance covers who an agent acts for and how much it can spend, best enforced through LLM and tool gateways [14].

On trust, a follow-up to the OpenRouter provider removals: Together AI separately disclosed that stored prompt samples from June–July were downloaded after an AI coding agent committed a credential to a public GitHub repo [15].

Macro pressure on AI infrastructure

- **Interest rates.** IMF chief Kristalina Georgieva says demand for AI hardware is pushing up consumer prices worldwide and "calls for a tightening of monetary policy in many economies" [16].
- **Permitting.** More than 70 European data-center projects were rejected or restricted in the first four months of 2026. STL Partners estimates opposition has delayed or cancelled about $42B of spending [17].
- **China's chips.** Alibaba is training its next Qwen model on chips from its own subsidiary T-Head, and its Bailian platform already serves inference on Chinese-made chips [18].

For angels: Section 1244

Elizabeth Yin notes that Section 1244 lets investors treat up to $50k of a startup loss as an ordinary loss ($100k for joint filers), rather than the usual $3,000-a-year deduction cap against other income [19, 20]. The shares must have been bought directly from a US operating corporation that had raised $1M or less when it issued them [21]. SAFEs and convertible notes are a gray area, so check with a tax adviser [22].

Sources

1. r/artificial post by u/ZealousidealBend9373
2. X post by @sama
3. X post by @sama
4. X post by @sama
5. r/Futurology comment by u/Gari_305
6. r/artificial comment by u/Antique_Atmosphere97
7. X post by @bindureddy
8. r/artificial post by u/ElroyFlynn
9. 5 Interesting Learnings from Carta's Q3 2026 Data: $100B Raised in Six Months, 21% Fewer Deals, and 1 in 3 Series A Rounds Is a Bridge
10. Jason's Takes on This Week's 20VC: Anthropic's 6,000 Employees, a 2027 Without Nvidia as a Buyer, and Why Muse Is Replacing Marketing Automation
11. X article by @ernietedeschi
12. The Harness-Model Flywheel
13. r/artificial post by u/shoumikgoswami
14. X post by @hwchase17
15. r/artificial post by u/winer666
16. r/Futurology comment by u/Gari_305
17. r/Futurology comment by u/npr
18. r/Futurology comment by u/Gari_305
19. X post by @dunkhippo33
20. X post by @dunkhippo33
21. X post by @dunkhippo33
22. X post by @dunkhippo33