Anthropic's prospectus shows compute at 58% of revenue as Palo Alto pays $4.5B for Applied Compute

VC Tech Radar

2026-11-19

Anthropic's prospectus shows compute at 58% of revenue as Palo Alto pays $4.5B for Applied Compute

By VC Tech Radar • November 19, 2026

Anthropic's filing puts numbers on compute-heavy AI economics. Enterprises are shifting fast to open-weight and multi-model setups, and Washington is weighing a disclosure rule for Chinese models. Also covered: a $4.5B exit for a model-ownership startup and new funding for agent infrastructure and physical AI.

Anthropic's filing: fast revenue growth, compute costs rising faster

SaaStr's read of Anthropic's S-1/A: about 5,300 employees as of September 30 and an $81B run-rate as of October 31. It added $10B of run-rate in October alone [1]. Q3 revenue was $16.96B. Compute and infrastructure cost $9.86B, or 58% of revenue, up from 48% in Q2. Personnel was about 10% [1]. Anthropic posted a $567M operating loss for the quarter [1]. SaaStr's view is that the small headcount is driving growth, not profit, though cheaper inference could change that [1].

Investing in AI points to the balance sheet. Anthropic disclosed about $900.9B of commitments, substantially all compute and infrastructure, with $451.6B due before the end of 2030. These are contracts and leases running to 2045, not near-term debt [2]. The author's risk case is not a liquidity crisis. It is that customer consumption is flexible while the infrastructure bill is not. The stock could disappoint if revenue or compute economics fall short of the valuation [2].

David Sacks says an "Anthropic–Doomer split" has arrived. He argues Anthropic marketed "closing the loop" on recursive self-improvement on its IPO roadshow as a core advantage [3], and relays Tegmark's claim that Anthropic-linked donors have "co-opt[ed]" the safety ecosystem [3]. These are partisan views, but they are a narrative risk to watch around the listing.

For comparison, Greg Brockman said OpenAI doubled its enterprise revenue in Q3 [4]. He called compute the main limit on growth and the thing that sets AI companies apart, and cited OpenAI's in-house inference chip, Jalapeno [4].

Palo Alto Networks buys Applied Compute for $4.5B

Applied Compute signed a definitive agreement to be acquired by Palo Alto Networks for $4.5B [5]. Its thesis is that every company should train, serve and improve its own models. The founders say security, auditability and policy compliance led them to Palo Alto, and customers keep control of their models [5]. They say a team of a few dozen went from zero to over $100M in annualized revenue run-rate in just over 18 months [5]. That is a clear benchmark for how fast an "own your model" enterprise startup can reach a strategic exit.

Enterprises move to open-weight models, and Washington targets Chinese ones

The Databricks AI Index tracks token traffic from more than 10,000 companies. Open-weight models went from 7.3% of tokens in July to 30.9% in November [6], mostly for high-volume bulk work rather than the hardest tasks [6]. The leading open-weight providers are Chinese labs (Zhipu, DeepSeek, Moonshot), and GLM 5.2 accounts for about 40% of open-weight tokens [6]. Companies relying on a single model provider fell from about 64% to 28% [6]. On the proprietary side, Opus 5.5 went from 7.2% to 37.9% of tokens in six weeks [6]. Databricks sells open-weight hosting, so its sample likely overstates the trend [6].

On the same day, Reuters reported that the administration is weighing a rule requiring companies to disclose on their websites when they use Chinese open-source models. Sources describe it as "name and shame," not a ban [7], and Treasury and OSTP have discussed it [8]. Airbnb was named. It says its AI runs "overwhelmingly on U.S.-origin models," and that its few China-origin models run only through approved US providers [9]. Separately, ByteDance is reportedly in advanced talks to spend about $500M more with Microsoft for OpenAI models on Azure after being cut off by Anthropic. OpenAI has objected over distillation concerns [10].

For portfolio companies whose margins now depend on cheap Chinese open weights, this is reputational and compliance exposure to diligence.

Funding and early-stage signals

- **Linq: $90M led by a16z,** with Ben Horowitz joining the board and over $120M raised to date [11]. Linq is messaging infrastructure for agents. DoorDash, TownAI and tomo run agents on it, and customers have processed more than a billion messages [11]. It is expanding into agent phone calls and computer-use infrastructure [11].
- **AMI (Yann LeCun): $1B seed.** Funding came about 40% from Europe, 33% from the US and 27% from Asia, with "not a cent" from Silicon Valley VCs. LeCun cites their desire for control as one reason [12]. AMI is building world models from video and sensor data for industrial process control and planning. LeCun expects robotics to matter, but not in the short term [12].
- **BrightAI: $97M** for critical-infrastructure inspection, combining its Stateful platform with autonomous robots [13].
- **Xlue joins Conviction's portfolio.** The company claims its health world model beats frontier token-based models in blind third-party evaluations at under 0.1% of the compute [14].
- **Scylake (Nir Zuk, Palo Alto Networks founder),** backed by Lightspeed. It sells on-prem, sovereign AI security to air-gapped and regulated buyers [15]. Zuk says some customers burned their full annual LLM token budget within a month or two using cloud tools [15]. His co-founders were already retired [15].
- **Ramp** is reported to be in talks to raise $1B+ at a $70B valuation, a month after its $60B round [16].
- **Netlify** is growing nearly 130% toward $130M ARR after "re-founding" around non-developers new to AI [17]. Martin Casado calls it the most significant turnaround he has directly witnessed [18].

Agents in production

Resolve AI released SR1, its first model, alongside "Primary On Call." Resolve now takes production pages, mutes false alarms, applies fixes that policy allows, and escalates the rest to humans [19]. It says users in shadow mode saw it make the same call they would have about 99% of the time [19]. Those are company-reported figures.

A vendor-run study of 27,280 trials by Runta and Baseten found that the best-value coding harness changes with the model. Codex CLI had the lowest cost per passed task with GLM-5.3 but ranked 12th with another model [20]. LangChain's open-source deepagents harness passed 50M downloads [21]. Perplexity says its API business grew 4.3x year over year to a nine-figure run rate, after launching Monitor, Orchestrator and Context APIs [22].

Nvidia

The fight over Jensen Huang continues. Travis Kalanick says he will show up at Nvidia HQ with employees ready to walk out, and founders, to "fight" for Huang's return. He concedes "we don't have control of the outcome" [23].

Sources

1. Anthropic Will Get to $100B in Revenue with ~7% of the Headcount Google and Meta Needed
2. Anthropic's IPO Refinances Their Compute Bill
3. X post by @DavidSacks
4. Why Sam Altman and Greg Brockman Say Humans Must Stay in Control of AI
5. X post by @ypatil125
6. 5 Interesting Learnings from the Databricks AI Index: Open-Weight Jumps to 31% of Enterprise Tokens, 72% of Companies Now Use Multiple Model Providers, and a 281x Gap Between Median and Power Users
7. X post by @Jason
8. X post by @Jason
9. X post by @Jason
10. X post by @leomschwartz
11. X post by @elliott___potter
12. AMI Founder & Chair Yann LeCun at Silicon Valley & The World | Semafor
13. X post by @nickbkim
14. X post by @saranormous
15. Why More Security Teams Need AI That Runs On Premises | Cyber60
16. r/SaaS post by u/Hot_Hearing3480
17. X post by @alexrkonrad
18. X post by @martin_casado
19. What Happens When AI Fixes Production Without a Human in the Loop? | Lightwork
20. r/artificial post by u/shoumikgoswami
21. X post by @sydneyrunkle
22. X post by @AravSrinivas
23. X post by @dnapway