# Lean community certifies the Claude-assisted Riemann proof as Palo Alto pays $4.5B for an AI team

VC Tech Radar

2026-11-22

## Lean community certifies the Claude-assisted Riemann proof as Palo Alto pays $4.5B for an AI team

*By VC Tech Radar • November 22, 2026*

The Claude-assisted Riemann proof passes formal checks. Palo Alto pays $4.5B in stock for Applied Compute. Perplexity reports passing $1B in run-rate revenue. Compute stays scarce and inference data leaks surface on OpenRouter. Also: what founders should know before raising over the holidays.

### The Riemann proof has passed its formal checks

Last brief, Levent Alpöge's Claude-assisted proof of the Riemann hypothesis was still waiting on outside checks. Those checks have now come back. The Lean community posted a statement confirming four things [1]: - The Riemann hypothesis statement used is the one already in Mathlib, which experts had vetted. - Dozens of people compiled the library and got the same proof export. - Two independent tools confirmed that the theorem Alpöge proved is exactly Mathlib's Riemann hypothesis. - The exported proof passed Lean's kernel and several independent checkers. Two of those checkers come with formal guarantees against accepting a false proof.

Exponential View notes that Alpöge's proof runs to 17.4M lines of Lean. A separate all-human proof by Alessandro Masullo is being translated into Lean so it can be checked faster [2].

For investors, the takeaway is that machine checking, not human refereeing, is becoming how frontier results get validated, even for human-written proofs. The proof is certified. Whether mathematicians can actually learn anything from a proof this long remains open.

## AI M&A: $4.5B for a team with no disclosed revenue

Palo Alto Networks is buying Applied Compute, which does post-training and inference for open-weight models, for $4.5B in stock. The release gives no revenue or customer figures [3]. CEO Yash Patil becomes Co-Chief AI Officer, and the team will train Palo Alto's own models [3].

Nikesh Arora's reasoning is that enterprises will run a mix of open and closed models, and every one of them will need securing. Palo Alto first saw the platform as a prospective customer [3].

The buyer's own numbers are worth a look. Acquisitions drove most of its 39% Q1 growth, and growth falls to about 18% once the CyberArk deal is in both comparison periods [3]. GAAP operating income was only $96M on $3.44B of revenue, because deal costs and stock compensation are piling up [3].

## Perplexity passes $1B in run-rate revenue

Aravind Srinivas says Perplexity Computer, launched in February, drove 4.7x year-over-year revenue growth [4]. The company crossed $1B ARR after targeting $650M by year-end. He describes a 2028 IPO as the "mental model," and says it could file sooner [4]. He cites 56% gross margins and no large compute commitments [4]. Routing simple answers to its own post-trained model saves about $100M a year [4].

The catch is model dependence. Perplexity pitches itself as model-neutral, yet Opus 5.5 handles 60–70% of agent tokens. Srinivas concedes Perplexity lags on how efficiently agents complete tasks, and says that is why it is building a post-training research lab [4].

## Compute stays scarce, and power is the bottleneck

Kush of Ornn, a market for AI compute, told 20VC: - **Supply:** buying capacity from CoreWeave today can mean a Q4 2027 or Q1 2028 delivery quote. Building your own takes about six months [5]. - **Prices:** renewal options at around $2.50 per GPU-hour are no longer free. He wouldn't be surprised by $20–30 per GPU-hour in three years [5]. The risk to that view is weaker enterprise demand pushing GPU "rents" down [5]. - **Where he'd invest:** he calls "a round done in 2 hours" FOMO-driven [6]. He says the vertical-AI categories now led by companies like Harvey and Rogo are already taken [6]. He would back unglamorous physical infrastructure, such as a startup building better transformers [6].

Exponential View adds that more than 2,000 GW of US generation and storage is waiting to connect to the grid, which is more than all plants running today [2]. With an AI chip renting for about $3 an hour, a month spent waiting for power costs more than the power itself. On-site solar plus batteries is the cheapest data-center power option in half the US [2].

## Data leaks among inference providers

OpenRouter off-boarded three inference providers after security researchers reported compromised data or data policies not being upheld. It also removed 125 endpoints from Zero Data Retention (ZDR) eligibility, some as a precaution. OpenRouter says it found no breach of its own systems and credited a Dragos report [7].

Suhail is pitching Particle on that platform with explicit promises: it respects ZDR and will never route traffic to China [8]. A researcher Suhail credits with uncovering many of the breaches confirmed there were "no leaks through particle from july to november" [9]. Data handling is becoming a selling point inference providers compete on.

## Policy on Chinese models and compute

Clem Delangue praised a new US House report for recognizing that the right answer to Chinese open-source AI is to build more American open source, not to ban it [10].

A WSJ headline says Southeast Asian data centers exploited a sanctions gap to give Chinese firms compute and cover for alleged chip smuggling, and frames Jensen Huang as the builder of that route [11]. One commenter disputes the framing, saying the route was legal until it was banned on November 19 [12].

## Fundraising timing and revenue metrics

Elizabeth Yin's advice for the holidays: - If you are already raising, close quickly and take what you need to reach the new year [13]. - If you are in trouble and haven't started, run a small process aimed at fast-moving emerging managers [14]. - Only companies with standout growth tend to close in this window [15].

Separately, an investor she quoted says many vertical-AI startups exaggerate revenue, and claims to have seen 15 definitions of "CARR" [16]. Ask founders exactly how they define it during diligence.

## Agents: more activity than shipped output

- **NBER study:** a paper covering 500k+ GitHub developers reportedly finds autonomous coding agents raised commits by about 240% but releases by only about 30% [17].
- **Salesforce ROAM:** its release-review agent cut median turnaround from about 24 hours to 12 seconds across 18 products. Code decides whether a patch passes, the model only explains the result, a human gives final approval, and about 8% of requests go to a person [18].
- **LangChain:** it is promoting a "knowledge factory," a curated company knowledge base maintained for its agents [19].

## Elsewhere

- **AMI Labs:** Yann LeCun's Paris company builds world-model AI rather than text predictors. It raised just over $1B in March [20]. He argues extinction warnings help the biggest labs win rules that shut down open research [20].
- **Pharma:** GSK, BMS, Novartis, Lilly, Novo, AstraZeneca and Sanofi say AI has cut by as much as half the time it takes to produce new molecules and move them into the clinic. They now use it to design and rank molecules and to flag safety risks [21].

---

**Sources**


1. r/Futurology comment by u/SgathTriallair
2. Certified human, twin primes, Anthropic++ #611
3. 5 Interesting Learnings from Palo Alto Networks at $14.5 Billion in Revenue: 39% Growth, a $1 Billion GAAP Gap, and a $4.5B Stock Bet on Open-Weight AI
4. Aravind Srinivas: AI and the American Dream | TomorrowX Summit
5. X post by @HarryStebbings
6. Why Neo-Clouds are Great Businesses & The State of Compute with Kush Bavaria, Co-Founder @ Ornn
7. r/artificial post by u/jomic01
8. X post by @Suhail
9. X post by @hdmoore
10. X post by @ClementDelangue
11. r/Futurology post by u/Confident_Salt_8108
12. r/Futurology comment by u/CryMoreT_T
13. X post by @dunkhippo33
14. X post by @dunkhippo33
15. X post by @dunkhippo33
16. X post by @anothercohen
17. r/artificial post by u/Antique_Atmosphere97
18. r/artificial post by u/shoumikgoswami
19. X article by @hwchase17
20. Is Hinton Exaggerating, or Is LeCun Downplaying the Danger of AI?
21. r/Futurology comment by u/Gari_305