

Back to Research

How Headlines Move GPU Prices

February 17, 2026

H200: Speculation to Stability

Over the past few months, NVIDIA's H200 has shown us how world events and global discourse around AI and data centers can affect GPU pricing. In early December 2025, H200 rental pricing spiked when the Trump administration announced it would allow NVIDIA to export H200s to China *(1)*. As political discourse around this decision ramped up, spot market uncertainty over the next two weeks increased, and volatility rose by nearly .25, resulting in a 90-day-high of $2.73/ GPU hour on December 20th. On December 22nd, NVIDIA confirmed a mid February delivery timeline, addressing the fear/uncertainty in the market, bringing pricing down to about $2.30. In mid January, AWS raised H200 pricing by around 15%, signaling a supply and demand imbalance *(2)*. This

announcement helped remove speculative volatility for H200 pricing as the market began to accept this increased pricing as the new normal. This resulted in the steepest single-week decline in pricing volatility between January 12th and 15th. Later in January, the DeepSeek “efficiency” conversation came back into the mainstream, pushing prices down to a low of $2.21/hour before rebounding that same week, while volatility kept falling instead of spiking, which is what happens when headlines don’t move actual supply and demand. In early February, Chinese GPU approvals for large orders helped ease expectations, with price tightening around $2.25 to $2.27/ hour *(3)*.

B200: Volatility is Floored

In early to mid December, B200 rental pricing stayed in a pretty tight range around $3.02 to $3.08/ hour, even as headlines pushed the “next chip” narrative for B300s *(4)*. By December 21st, a mix of concern regarding Blackwell supply and end of year buying led to a spike in pricing at $3.69, and volatility stayed high around 1.3 to 1.4 as nobody knew how long that market crunch would last. In early January, the conversation around B200s changed once the market started accepting the high GPU pricing, and the B200 spot price began going back down toward roughly $3.50 to $3.60 while volatility dropped quickly. That change in thinking continued as AWS raised GPU pricing by

around 15% *(2)*, which suggested the whole market had a higher baseline price. The most important break came in late January when China approved around 400,000 plus H200 chips, which reduced the need for buyers to chase the top-line GPUs through back channels, and B200 pricing dropped rapidly to about $2.65 on January 21st. After that drop, the market cooled off, with B200 pricing settling into a more normal range around $3.00 into early February as the big questions got answered and trading became less panic-driven. By February 10th, price movement looked more settled, and volatility sat near its floor around 89 points, which tends to be where things land once the market stops trading off every headline and starts pricing off fundamentals instead.

A100: Older Chip, Tighter Range

In early December, A100 rental pricing dropped from from $1.02/ hour to $0.88 as the market reacted to two headlines, with Trump moving toward allowing H200 sales to China *(1)* and Supermicro starting to ship B300 systems *(4)*, pushing A100 more into the “older but still useful” consideration. Volatility stayed jumpy around 55 to 65 points because people were still guessing how quickly buyers would shift to newer chips and what that would do to supply. By late December, the DeepSeek V3 reporting confirmed the idea that strong

performance could still be accomplished with older hardware, designating the A100 as useful for inference and cost-sensitive workloads, but not enough to push demand for training *(5)*. In early January, AWS increased its H200 price and signaled that even previous-generation hardware is constrained by supply. This move did not raise A100 prices, but it gave the A100 a firmer footing with pricing holding around $0.82 to $0.83 and volatility falling toward 0.35. In late January, Reuters reported that China approved more than 400,000 H200 chips for firms like ByteDance, Alibaba, and Tencent *(3)*, which put a stop to the conversation that “maybe China still needs older chips”. After that, A100 pricing flattened into a tighter $0.80 to $0.82 range, while volatility went down toward 32 percent into early February as the last remaining questions around older chip demand got answered.

Sources

1: CNBC, "Trump greenlights Nvidia H200 AI chip sales to China if U.S. gets 25% cut," December 8, 2025, https://www.cnbc.com/2025/12/08/trump-nvidia-h200-sales-china.html

2: Introl, "AWS GPU Price Increase: H200 January 2026," https://introl.com/blog/aws-gpu-price-increase-h200-january-2026

3: Reuters, "China Clears DeepSeek to Buy Nvidia Chips," January 28, 2026, https://finance.yahoo.com/news/china-clears-deepseek-buy-nvidia-155937026.html

4: Supermicro, "Supermicro Expands NVIDIA Blackwell Portfolio with Liquid-Cooled HGX B300 Solutions Ready for High-Volume Shipment," December 9, 2025, https://www.supermicro.com/en/pressreleases/supermicro-expands-nvidia-blackwell-portfolio-new-4u-and-2-ou-ocp-liquid-cooled

5: SemiAnalysis, "DeepSeek Debates," https://newsletter.semianalysis.com/p/deepseek-debates

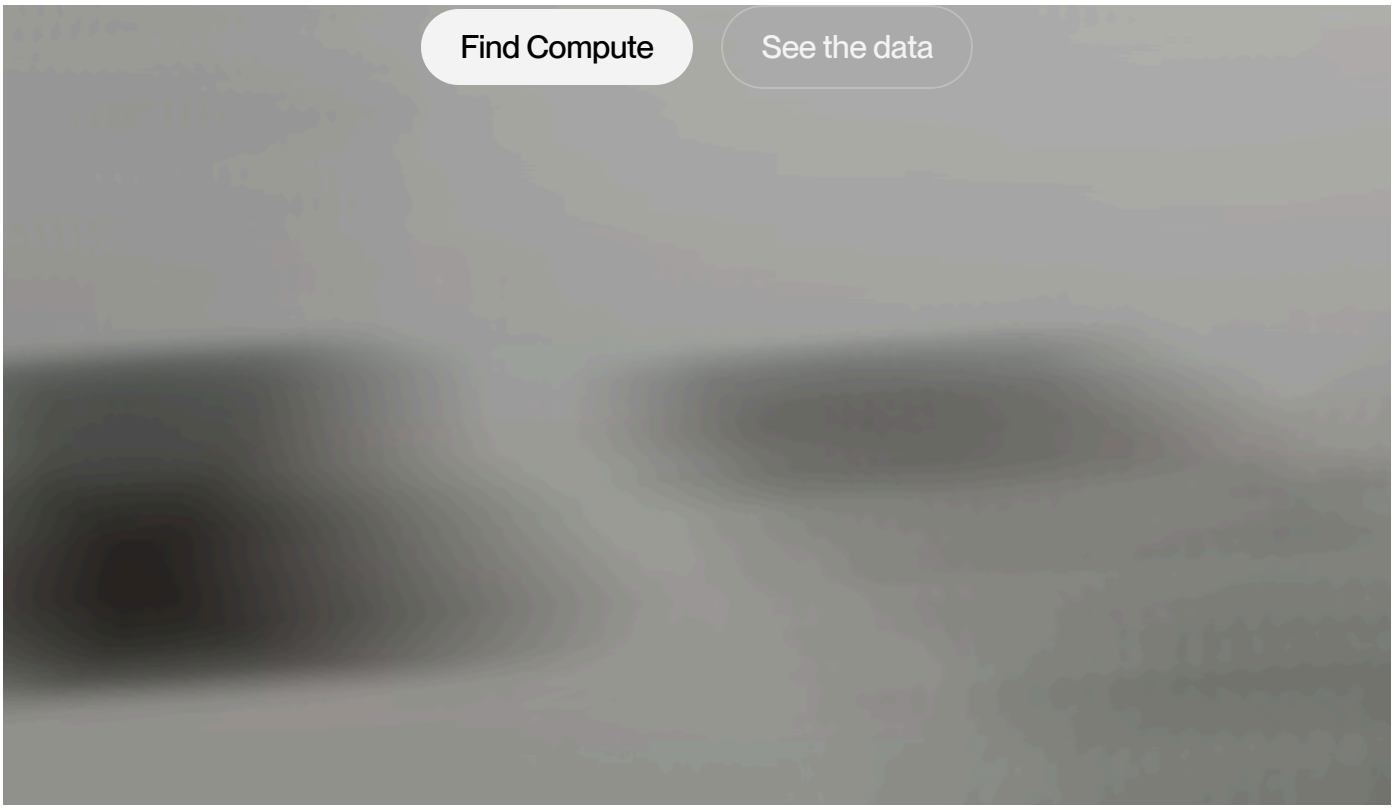

X
LinkedIn
Substack

Careers
About
Research
Docs
Trust Center
Regulatory

Privacy Policy
Terms of Use

Brand + Website by Haptic

©2026 Ornn AI, Inc. All rights reserved.

