Notes from our seat at the table, inside the financial plumbing of compute.

The Great Rerouting

Nov 23, 2026 · @Alaz Tanyeri

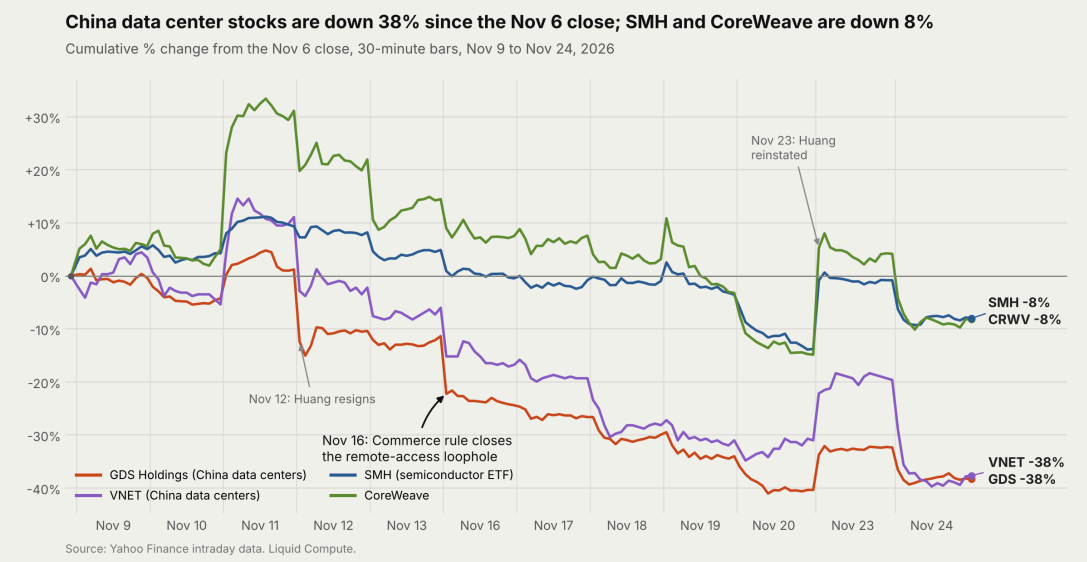
The offshore loophole closes and the compute comes home.

The Great Rerouting

The US has been limiting China's access to cutting-edge chips via export controls on the hardware for a while, but there was a loophole: a Chinese AI lab could rent GPUs in a data center outside China and log in remotely with an SSH key. On Nov 16, a Commerce/BIS interim final rule moved to close this loophole by requiring a license to ship advanced chips to any data center, anywhere in the world, that sells meaningful compute to China-headquartered companies. It also now takes a license to export those chips to five countries Commerce calls "higher-than-normal risk of diversion": Indonesia, Malaysia, Singapore, Thailand, and Brazil. I think the tariff analogy will apply: compute will get pricier inside the restricted zone (China and those five nations) and cheaper outside of it, just as US tariffs mechanically pushed up the price of goods in the US relative to the rest of the world.

The first shock is uncertainty. Nobody knows who will be able to serve whom under which license, and that uncertainty itself freezes financing for Chinese-affiliated offtake. A lot of capacity has no choice but to find new offtake as a result. This supply is real and already showing up. Our desk has already seen a wave of Philippines, Malaysia and Indonesia capacity come to market. Operators are telling us they can't get bridge financing because their offtake is Chinese-affiliated. Some lenders are declining offtake outside the US and EU. That's tough on lenders already holding loans to China-adjacent offtake: if those offtakers get banned, their recovery depends on the re-let market staying hot. Since operators need 3-year offtake to match 3-year loans, this flood of supply is showing up in the term market first. These seem like early signs, and the stress caused by export policy could get worse over the next 6–12 months. For example, some of the Asia-linked capacity we've been shown doesn't deliver until February, which means we haven't even seen the full first-round effects of this rise in supply. On top of that, there is a 90-day compliance window for overseas data center operators to acquire licenses. This not only means the first-round effects of this supply can grow over the next 3 months, but it also means the uncertainty overhang of this decision isn't going away anytime soon. As we learned from the tariff episode last year, the uncertainty around a policy can hit harder than the policy itself.

The public markets woke up to this reality, though US semis and neoclouds only caught up on a lag with the narrative for the selloff applied retrospectively, whereas Chinese datacenters reacted first.



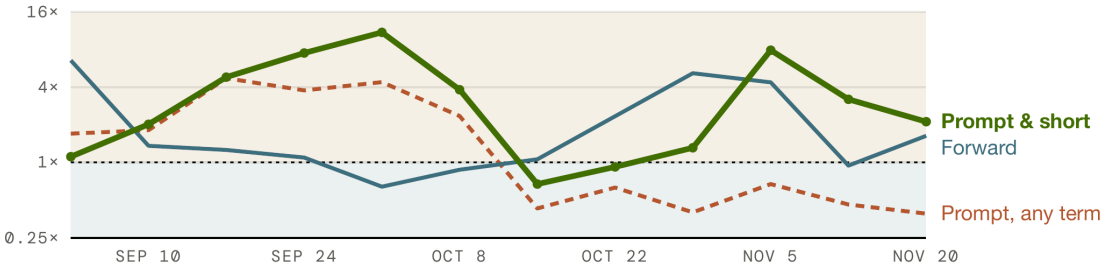
The starkest example so far is Datasection (3905.T), a Tokyo-listed neocloud that was renting most of its ~15,000 Blackwell GPUs to Tencent. On Nov 16, the same day as the rule, it cut its full-year revenue guidance by 89% (from ¥162B to ¥17B) after a ¥80B AI order was cancelled.

You can follow the fingerprints of this story on the compute market in The Compute Channel Check: insights we publish from inside the room where compute is actually traded. We just published the first volume today, you can find here.

Demand is still strong, but it's picky

The bright side is that this rerouted supply is hitting a market in a supply-demand deficit. The catch is that demand is specifically for US-based, near-term, shorter-tenor capacity with low prepay.

Demand-supply imbalance, cut three ways: GPUs requested per GPU offered

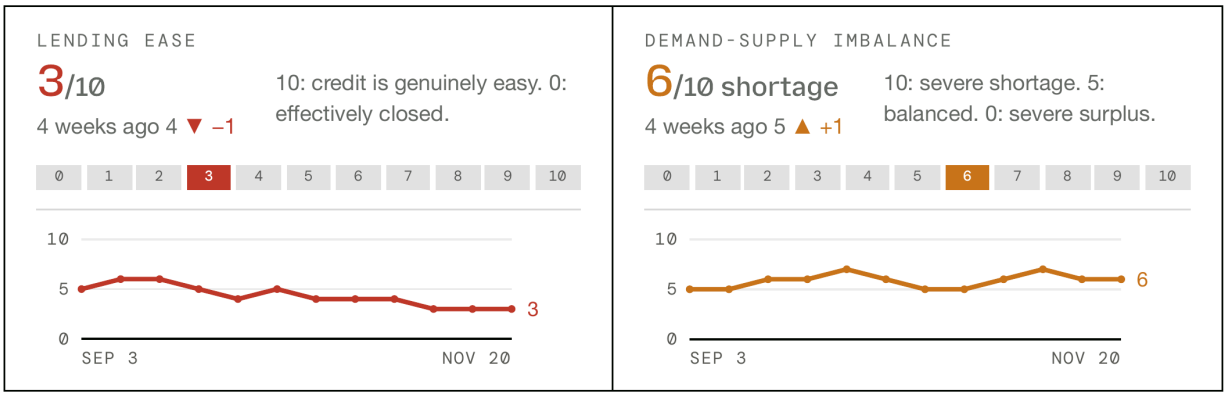


Four-week window, matched by chip family. Above 1× buyers want more than sellers offer; gridlines step ×4, so equal shortages and surpluses sit equally far from 1×. Prompt: delivery within 60 days. Short: a year or less. Weeks under 50 GPUs offered left out.

Source: The Compute Channel Check, Vol. 1 (Nov 20, 2026), Liquid Compute.

Before the export headline even hit the tape last week, we had been seeing a shift since September driven by tightening lending conditions in the compute market. One reason prompt compute pricing rallied so hard to begin with is that tighter credit pushed sellers to ask customers for higher prepays, while the LTVs lenders offer kept coming down (from 70–80% over the summer to 50–60% now). Sellers are asking buyers to put up more capital upfront and simultaneously lenders are being tighter with the money they provide to bridge the gap. So buyers' only way out is to shorten the tenor of what they buy, which lowers both the prepay in absolute dollars and the Total Contract Value (TCV), the total cumulative dollar cost of an entire rental contract. This clears up a common misconception about why GPU rental indices are so high. Of course, a positive demand-supply imbalance is part of it, but another factor flowing into spot demand is lending ease, which has clearly tightened.

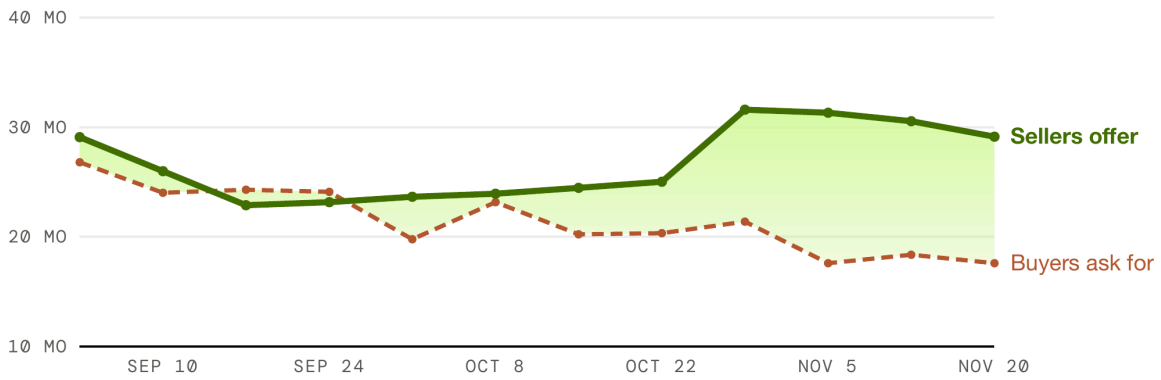
Lending ease vs. demand-supply imbalance, weekly



The two scores, weekly Each point covers the four weeks to that date.

Source: The Compute Channel Check, Vol. 1 (Nov 20, 2026), Liquid Compute.

The tenor gap: what sellers offer vs. what buyers ask for

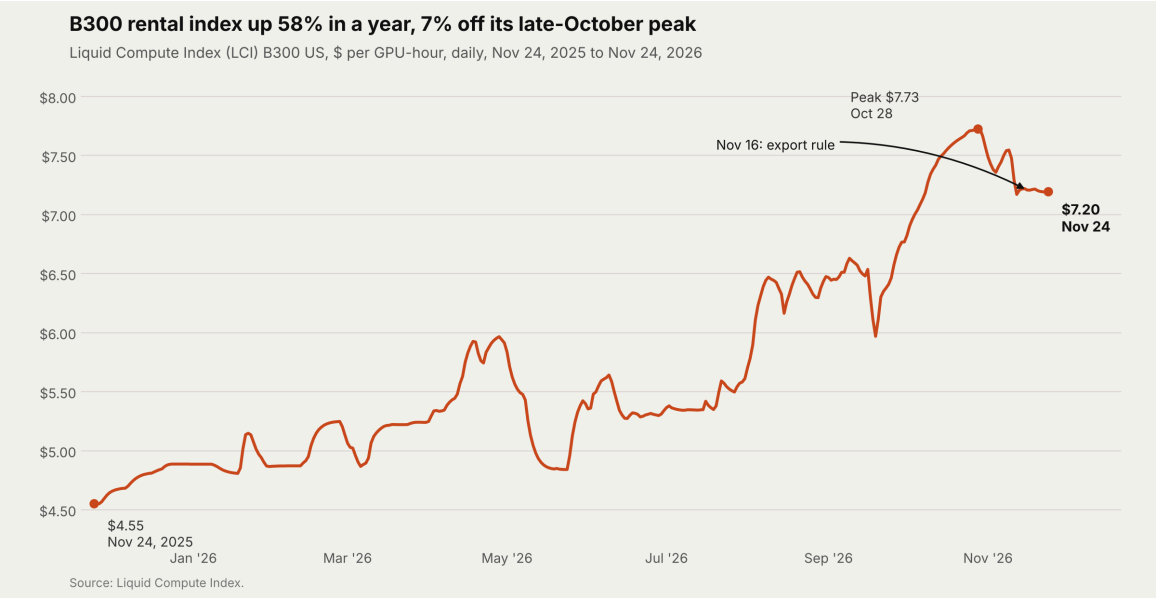


The tenor gap is the average term sellers offer minus the average term buyers ask for: +12 months now, against about zero in September. It opened in October and has held at about a year since early November.

Source: The Compute Channel Check, Vol. 1 (Nov 20, 2026), Liquid Compute.

To make things worse, there is a timing problem with the supply in the market. Most of the GPUs third-party sellers offered us in the past month won't be ready for service (RFS) until sometime between January and June 2027. Sellers market them in advance to line up a customer they can show their lenders before delivery. That's a second wave of supply on top of the export-control one.

In summary, demand is strong and supply is growing, so the B300 price index has softened a bit but has not collapsed. The rerouted supply from Asia is clearing at a discount and on shorter tenors.



Read the full Compute Channel Check · Lending ease, demand-supply imbalance, prepay and tenor data, every four weeks.

Track live GPU rental rates on the Liquid Compute Index. *[link]*