Introducing The Compute Channel Check

The Compute Channel Check is our new Market Intelligence offering. We aim to give you a look into the new and growing compute market through our lens. As the financial infrastructure of the compute market, we sit between the three counterparties of compute trading: Buyers, Sellers, and Lenders

Why it matters

GPU clusters are built with private and public credit, sold on multi-year contracts and paid for partly up-front, and the cost and availability of that capital decides what gets built, and at what price. We focus on the middle market for compute, but that market is competing for the same capital as the larger players.

The public signals arrive late. Chipmaker and neocloud earnings report a quarter after the fact. Spot prices show the result of financing conditions, not the conditions themselves. By the time either moves, lenders and buyers have usually already changed their behavior.

What we see

Liquid Compute sits inside these transactions. We trade physical compute, arrange financing, and publish the Liquid Compute Index, so every week our desk sees what sellers ask for, what buyers want, and what lenders will and won't fund.

The Compute Channel Check turns that flow into two scores for investors exposed to AI infrastructure: neocloud equity and credit, chipmakers, private credit and the funds that lend to them.

The two scores

	LENDING EASE	DEMAND-SUPPLY IMBALANCE
What it measures	How easy it is for buyers, sellers and operators to get capital	How many GPUs buyers want for every GPU sellers offer, matched chip for chip and split by how soon and for how long
Scale	0 to 10. 10 is genuinely easy credit; 0 is effectively closed	0 to 10. 5 is balanced; above 5 is a shortage, below 5 a surplus

	LENDING EASE	DEMAND-SUPPLY IMBALANCE
Built from	The prepayment sellers ask for, the gap between the contract length sellers offer and buyers want, the cost of term loans and bridges, how often lenders pass, and how much lending is rescue financing	GPUs requested per GPU offered for near-term, short contracts and for later delivery, and how urgently each side is acting
Why it matters	When it falls, operators slow builds, sellers demand cash up front, and leveraged neoclouds feel it first	A shortage gives owners of available capacity pricing power; a surplus is the early sign of falling prices and idle clusters. It shows which side of a contract has leverage before the index does

Reading the two together

	SHORTAGE (IMBALANCE ABOVE 5)	SURPLUS (IMBALANCE BELOW 5)
Easy credit (lending ease 6 or more)	Build-out accelerates; capacity gets financed and delivered	Overbuilding risk; cheap money meets thin demand
Tight credit (lending ease below 4)	Buyers want capacity, but capital, not demand, is the constraint	Distress: refinancing pressure and forced sales

The combination matters more than either score alone. A shortage with easy credit is a market that is about to add supply; the same shortage with tight credit is one where supply can’t respond, and where contract terms, especially prepayment, start to carry the price.

Where the data comes from

Liquid Compute’s desk and financing records: third-party seller quotes, buyer requests, deal outcomes, lender decisions and rate quotes. Counterparties are never named, and our own offers are excluded. A week is only scored when each input has enough observations behind it, so the history lengthens as the data does.

Key terms

TERM	WHAT IT IS	WHY IT MATTERS
TCV (total contract value)	The full value of a compute contract over its life: price per GPU-hour × number of GPUs × hours in the term. E.g. if an AI lab rents one B300 GPU for 3 years at $4 per GPU-hour, the TCV is $4 × 26,280 hours = $105,120. For a 1,024-GPU cluster on the same terms, that is a TCV of about $107.6 million	It is the base everything else is measured against. Prepay is a share of TCV, and lenders size loans against it, so a long contract with a high TCV is what makes a cluster financeable. It is also why sellers push for longer terms.
Prepay	The share of a contract's value the seller asks the buyer to pay up front. E.g. A 30% prepay means the buyer pays the seller $31,536 up front in cash if the TCV is $105,120.	Prepayment is becoming a form of implicit financing. Buyers prefer putting 0% down on their purchases whereas sellers are asking for more prepayment to secure financing for their hardware purchases. The gap between buyers and sellers measures financing difficulty in the market.
Tenor gap	The average contract length sellers offer minus the average length buyers ask for, in months.	Sellers who own equipment generally need to rent out their compute for longer tenors to be able to secure financing for their equipment, whereas buyers, often financed via VC equity, prefer renting on a shorter tenor because they don't want to tie down capital and because they are uncertain how much compute they will need in the future. The difference is an indicator showing you how difficult it is for both sides to access capital
RFS (ready for service)	The date a cluster is live and the buyer can start using it.	For deals being signed now, RFS is often several months out: half the capacity offered to us in the last four weeks wasn't ready for at least two months. That matters because buyers often have to prepay and commit today for capacity that only starts earning months later, which raises the cost of waiting, and it is why near-term, ready capacity is where the shortage shows up.
Cost of financing	The median rate quoted to us on GPU-backed term loans and on bridge facilities.	It is the price of capital. Higher rates squeeze margins on clusters whose rental income is fixed by contract, and push the weakest operators out first.
Lending appetite (term-loan pass rate)	The share of term-loan decisions where the lender passed (rejected the opportunity), rather than issuing a term sheet or funding.	The most direct read of whether lenders want to commit long-term money. Bridge and rescue requests are left out, so emergency lending isn't mistaken for appetite.

TERM	WHAT IT IS	WHY IT MATTERS
Bridge share	The share of lender activity that is bridge, rescue or prepayment financing.	Bridges are meant to be an emergency: short-term money to cover a gap until permanent funding arrives. A rising share means operators can't get term funding and are patching holes instead, which is a sign of stress rather than of lenders being willing.
Cover ratio	GPUs buyers request divided by GPUs sellers offer, matched chip for chip: B300 requests are compared only with B300 offers, H200 with H200, and so on. Above 1× buyers want more than sellers offer; below 1× there is more on offer than buyers want.	It answers whether buyers can actually find the capacity they are looking for. Matching matters because a buyer who needs B300s can't use an H100: leftover older chips coming off leases would otherwise make a market look loose when the chips people want are scarce.
Prompt cover ratio (prompt and short)	The cover ratio for capacity ready within 60 days on contracts of a year or less.	This is where buyers are today: they want capacity soon and on short terms, while most supply comes on long contracts or months out. It carries half the weight of the imbalance score because it is where shortages show up first and where owners of ready capacity have the most pricing power.
Forward cover ratio	The cover ratio for capacity ready more than 60 days out, at any contract length.	It shows whether supply that is being built will meet demand. If the prompt ratio is short but the forward ratio is balanced, the shortage is about timing and should ease as clusters come online; if both are short, it is a lasting shortage.
Demand vs supply pressure	How urgently each side is acting: the share of buyer requests wanting delivery within 30 days, deals signed and active buyers, against the share of offered GPUs available within 30 days and the volume offered.	The cover ratios count GPUs; pressure reads behavior. When they disagree it is informative: right now short-dated capacity is scarce, but fewer buyers are insisting on immediate delivery, which suggests buyers are being priced out rather than scrambling.

Contacts

- Alaz Tanyeri · alaz@liquidcompute.com
- Stanley Lee · stanley@liquidcompute.com
- Ronit Jain · ronit@liquidcompute.com
- Zach Frederick · zach@liquidcompute.com