


The quiet layer under the compute boom

Building out the financial infrastructure for compute: the credit, the data, and the markets forming around it, and where the value will accrue



Everyone is watching the visible part of the compute boom. The attention goes to the chips, the data centers, and the first announced exchange futures. All of that is happening on the surface. Underneath it, largely unremarked, a full financial system is being built to fund the whole thing, and that is where the real value will accrue.

In our view, the paradox that makes this urgent is simple. Capital is underwriting AI infrastructure at hundreds of billions of dollars a year¹, before the market has a stable collateral price, a depreciation curve, or a way to hedge. Finance is running ahead of the thing it is financing. Whoever builds the missing price and data layer sits beneath it all, powering the entire system.

At Motive, we are looking at financial infrastructure and emerging opportunities every day. We haven't seen a market of this kind assemble this quickly in sometime. Here's how we see the forces lining up.

WHY CAPITAL IS STARTING TO ORGANIZE AROUND COMPUTE

The four largest US hyperscalers are set to spend roughly \$725 billion on capex this year, up more than 75 percent from last year², and appears to be outstripping the cash flow they generate. The wave of operators building alongside them has no such cushion and, in our view, runs almost entirely on debt-funded capacity. Underneath the spending, something subtler is happening. Compute has started to behave like a financeable asset. Prices fluctuate, capacity gets rationed, buyers reserve years of supply before they need it, and whoever pays most gets served first. Compute capacity is now the constraint the rest of the AI economy is waiting on. The scarce input sets the terms, and the software built on top takes what it can get. When a scarce input gates everything downstream, capital organizes around it: pricing, funding, and hedging.

Today, the money flows on the strength of a few names. A lender funds a data center because a frontier lab or a hyperscaler signed the offtake, a counterparty too big to fail. That is lending against the logo, and the machine barely enters the math. The hyperscalers and large enterprises use the financial layer to sharpen a deal they could already do, while everyone else needs it to do the deal at all. The only way to fund them is to underwrite future compute cash flows as collateral. That will include defining what it is worth, what it will be worth, and whether the risk can be laid off. And that takes a price, a curve, and a market.

¹ Source: JPMorgan, Pitchbook

² Source: Financial Times

READ IT AS A STACK, WITH FINANCE ACROSS ALL OF IT

It helps to see the physical build as a stack of assets, each financed differently, with a financial layer forming around all of them.

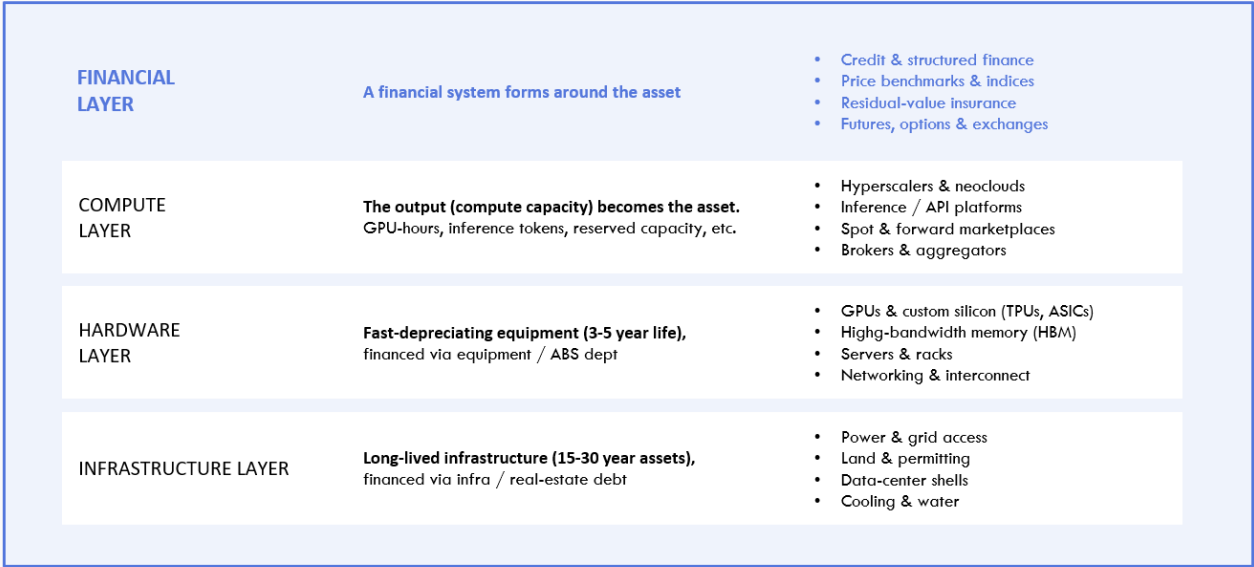
At the base lies infrastructure: power, land, and the building itself. These are long-lived assets. A data center typically runs for fifteen to thirty years, and what increasingly decides whether it gets built is access to power, not access to chips. This layer is generally financed like traditional real estate: long, patient, and secured against a stable structure.

Above it sits the hardware. The GPUs, the memory, the networking that turns a building into a usable cluster. This behaves quite differently from the layer below, as a GPU has three to five good years before the next generation makes it a liability. It gets financed more like traditional equipment.

That mismatch is easy to miss and expensive to ignore. A three-year asset sits inside a thirty-year one, financed by two debt markets that behave nothing alike. The building borrows like real estate, while the chips borrow like equipment: short and increasingly packaged into asset-backed securities that the market is still learning to price. Most coverage treats data-center financing as one thing. It is at least two, and the newer components carry most of the risk.

After comes compute itself, this is where things get interesting. GPUs are the infrastructure, while compute capacity is the economic output they produce. A GPU in a rack just depreciates, but contract an hour of its time, and you have something you can finance, benchmark, and eventually trade. The market is converging on a pricing unit, the GPU-hour, priced by chip: most emerging indices and announced futures settle on the rental rate for an hour of an H100, H200, or B200, and reserved, on-demand, and spot are just ways of buying that same hour. Tokens sit a layer up, model-specific and collapsing in price, so no one finances against them yet.

The catch is that a GPU-hour is not yet fungible. An hour on a tightly networked, highly available cluster is worth more than an hour on a smaller, poorly connected one; location, uptime, latency, software, and contract terms all create basis risk. The closer analogies are electricity, a standardized unit that still trades by location, time, and reliability, and aircraft finance, which underwrites against residual value, redeployment, and an asset's maintenance and configuration records. Oil supplies only the first idea, a reference grade with everything else priced around it. So, the first financeable unit is likely a specified capacity contract, not a naked GPU-hour, with generation, configuration, availability, location, and delivery all defined. Until those contracts throw off repeated arm's-length prices, lenders are underwriting contracts and counterparties, not compute as a standalone asset.



THE FINANCIAL LAYER IS WHERE IT GETS REAL

Emerging around the stack is the financial layer, where compute ceases to be a cost line item and becomes an asset class. Like every financial market before it, it develops in layers. Credit creates demand, data at the base creates transparency, benchmarks establish trust and standardization, insurance broadens participation, with liquid tradable markets emerging last.

- Credit is already here, and it is significant. GPU-backed lending is moving into investment-grade territory, evidenced by deals like CoreWeave, which closed an \$8.5 billion facility backed by its chips and its contracts, rated investment grade, priced near 6 percent³. One of the largest private-credit deals ever assembled to date, about \$35 billion led by Apollo and Blackstone, went to finance chips leased to a single AI lab⁴. Morgan Stanley's analysts put the data-center funding gap above a trillion dollars⁵, much of it likely to come from alternative capital, and see tens of billions a year in securitization, a growing share of it GPU-collateralized. Hundreds of billions are moving already, at investment-grade scale. And yes, most of it still prices the collateral on guesswork, underwriting three-year assets with no reference price, no depreciation curve, no forward curve. It gets away with it only because a well-established customer is on the contract.
- The whole stack rests on data that barely exists yet. Someone has to build the standardized record of what compute costs, what it is worth, and how fast it decays. Much of it already sits with the lenders, buried in the contracts and collateral marks they hold. A handful of young companies are assembling it now and fighting over the right method: whether the true price is what gets advertised, what actually printed, or what settled privately. Quiet and unglamorous today. It is the ground everything above it stands on.
- The benchmark comes next: the index, the curve, the number everyone else references. It does two jobs at once, as it lets lenders underwrite the asset instead of the logo, and it lets insurers stand behind the residual value. Whoever the market trusts to set it becomes the standard everyone else licenses and settles against. Standardization is what turns a contract into a market. Compute will need agreed specs (hardware generation, networking, availability, location, delivery) before a GPU-hour is interchangeable across providers, and only then do benchmarks get trusted and liquidity deepens.
- Risk transfer is the quiet backstop: The largest operators get financed because a blue-chip buyer signed the offtake. Everyone else needs another form of protection, such as insurance on the residual value. A first wave of firms is writing exactly that cover, with reinsurers behind them. The least visible part of the stack, and one of the most load-bearing. Without it, credit only reaches the borrowers who could already raise it.

- The market sits on top, and that's where all of this is heading: futures, options, and exchange-traded products that let compute be priced, traded, and hedged like any other asset. Some of it settles in cash, some in delivered compute. Some runs on the established exchanges, some on venues built for this from scratch. Get here and compute sits alongside oil, rates, and equities as something you can actually take a position in. This is what everything below is building toward, and once it works, it feeds back down the stack. Hedging closes the loop. Once price risk can be transferred, lenders extend more credit, insurers underwrite more efficiently, and the cost of capital falls across the entire market.

WHO ACTUALLY CAPTURES THE VALUE

Here is the question the analogy raises and most coverage skips. Every serious infrastructure asset class grows a financial system around it. In our experience, the most durable business in that system tends to be the one selling the price and the data, not the one lending the money or trading the contracts. Commercial real estate, oil and gas, aviation, shipping, power. The reference businesses that grew out of them and the exchanges' data arms, have historically stayed small relative to the assets they track while sustaining premium margins, and compounded through every cycle. Digital infrastructure is early on the same road.

The honest complication is that those same incumbents might be thinking about moving into compute pricing right now, and they own the distribution, so the index itself may well end up incumbent-branded. But the raw material underneath it, the standardized transaction data, is still unbuilt and largely ungoverned. A new entrant can still win there, on proprietary executed-transaction data, integration into lender and insurer workflows, independent benchmark governance, and distribution partnerships with the incumbents, even if an established name ends up stamping the final index. It is a very different bet from directional exposure to GPU prices, and one we believe is more attractive.

³ Source: Bloomberg

⁴ Source: Bloomberg

⁵ Source: Morgan Stanley

EARLY IS NOT THE SAME AS IMAGINARY

Some of the financing looks circular, with chipmakers funding their own customers, and that is a real risk rather than a footnote. The market is still early. There is no liquid spot market yet, and no settled unit. But the underlying forces are familiar.

In every market like this, the operators and the prices are remembered for a cycle. The plumbing underneath is the most financially lucrative: the reference everyone trusts, the credit that funds the build, the cover that makes it bankable. It is being built right now, quietly, while the attention stays on the chips.

And what is forming is, in our view, bigger than the futures market everyone is waiting for. It is a financial system where compute can be financed, collateralized, benchmarked, insured, securitized, traded, and hedged. The businesses we find most interesting inside it are not the ones lending the money or trading the contracts. They are the ones that will own the reference price and the data beneath it, the layer we expect

every lender, insurer, and exchange will eventually license.

“ And what is forming is bigger than the futures market everyone is waiting for.”

DISCLAIMER

The views expressed herein are those of Motive Partners employees and do not necessarily represent the views of Motive Capital Management LLC or any Motive managed investment vehicle. Market size and growth figures referenced herein are provided for illustrative purposes only; where attributed to Motive's analysis, they reflect Motive's internal estimates and assumptions, which may differ from third-party analyses.

Case study data has not been independently verified by Motive. This communication contains forward-looking statements; actual outcomes may differ materially from those described. Nothing herein should be construed as investment advice or as a solicitation to invest in any Motive managed vehicle.