

To: Ryan Taylor, Stack Infrastructure
From: Asher Vandevort, 252 Capital
Date: September 29th, 2025
Subject: Proposal for Standardized Unit of Compute and Market Design for Futures

Opportunity

Compute is the most rapidly depreciating core input in today's economy. Unlike fiber (decades) or power plants (multi-decade), top-tier GPUs traverse a 2 to 3 year obsolescence window. The result is that every idle hour is permanently lost value. Meanwhile, pricing is bilateral and opaque, and capacity planning is a mix of art and science.

Proposal: Standardize a financial grade of compute, an Effective Compute Hour (ECH), and launch cash-settled futures and OTC forwards on a small set of well-defined nodes/grades. This creates transparent price discovery and liquidity so:

- Producers can hedge-to-build and run higher utilization.
- Consumers can budget and hedge training/inference exposure.
- The ecosystem gets signals on location, interconnect, and time of day scarcity.

Phase 2 (optional, later): Add combinatorial auctions for complex allocation (continuity, hardware substitution, job splitting), analogous to spectrum and modern ad markets. Futures provide the financeable base layer; auctions optimize real-time matching when needed.

Why a Market Helps

1. Maximizing Utilization
 - a. Idle GPUs are wasted value given rapid obsolescence.
 - b. A market makes every buyer a potential reseller; unused capacity gets re-offered rather than lost.
2. Price Discovery & Budgeting
 - a. Transparent benchmarks replace closed-door negotiations.
 - b. AI labs can hedge training costs; data centers can lock in revenues.
3. Financing New Buildouts
 - a. Futures curves de-risk cash flows.
 - b. Banks and investors are more willing to finance new data centers when revenue is hedged.
4. Broader Access
 - a. Not just trillion-dollar hyperscalers: mid-sized companies, labs, and universities can participate.
 - b. Liquidity and standardized contracts lower barriers to entry.
5. Liquidity as a Feature
 - a. Speculators and financial institutions bring depth to the market.
 - b. As with oil or power, liquidity reduces volatility and spreads risk.

Futures First, Auctions Later

A futures market is the starting point:

- Standardize an "Effective Compute Hour" (ECH) benchmarked by throughput.
- Trade it on monthly contracts, by location/latency and grade (e.g., US-West Tier A clusters).
- Start with two or three clear contract specs to concentrate liquidity.

Over time, auctions can complement futures where preferences are complex: uninterrupted job windows, hardware substitutions, or co-location requirements. Paul Milgrom's auction design (spectrum

markets) and ad auction mechanics are useful analogies. However, the financial foundation should be established first, with futures for liquidity, and later, auctions for allocation.

Why Compute Is Different from Electricity or Fiber

- **Depreciation speed:** GPUs lose value in 2 to 3 years, vs. 30 to 40 years for power plants or decades for fiber.
- **Technology cycles:** Each generation leap (e.g., H100 > B200) massively improves efficiency. That makes squeezing full utilization from today's fleet critical.
- **Capital intensity:** Data centers require billions in upfront capex. A transparent, hedgeable market makes financing easier and spreads participation beyond the largest firms, including infrastructure builders.

Potential Challenges

- **Concentrated supply:** Hyperscalers may prefer bilateral control, but independent operators and brokers have the most to gain.
- **Hardware heterogeneity:** Specs differ. Not all GPUs and interconnects are equal.
- **Liquidity risk:** Too many contract types could fragment the market. Solution: launch small; 2 to 3 nodes, one contract size.
- **Verification:** Ensuring delivery/performance is critical. Solution: use hardware attestation and simple workload checks to confirm the agreed resources were delivered as promised.
- **Regulatory uncertainty:** Contracts need to fit within existing commodity derivatives frameworks. The path forward is to begin OTC, prove governance and transparency, and expand to listed futures once the market stabilizes.

Conclusion

Compute is following the same path as oil, electricity, spectrum, and ads: from bespoke negotiations to standardized, liquid markets. The difference is the urgency in which GPUs depreciate faster than almost any capital input, which makes maximizing utilization and creating liquidity paramount.

The proposal is to standardize a unit of compute that incorporates key parameters such as:

- Latency (interactive vs. batch jobs)
- Location (region-specific prices, like electricity LMPs)
- Cluster size (Train-8, Train-64, etc.)
- Job type (training vs. inference)
- Hardware grade (Tier A / Tier B, based on memory and interconnect)

With these parameters as the foundation:

1. Start with futures to establish transparent pricing, hedging tools, and liquidity.
2. Add auctions later to handle complex preferences and optimize allocations.

This two-phase approach delivers liquidity and financing first, while leaving room for more sophisticated allocation mechanics as the market matures. The result is a market that helps producers, consumers, and society while giving independent operators and infrastructure builders new tools to compete with hyperscalers. The question isn't if compute will be financialized, but how. Starting with standardized futures for liquidity, and layering in auctions later for complex allocation, is the most practical path.

I'd value your perspective on whether this model could fit into data center operations, and whether it's worth exploring in more depth together.