

Comment on the Request for Comment on the Listing of Compute Derivatives Contracts

To: Commodity Futures Trading Commission, Office of the Secretariat

Re: Request for Comment on the Listing of Compute Derivatives Contracts, RIN 3038-AF77 (91 FR 54259; docket CFTC-2026-1850)

From: Ola Ologunde, independent product designer, London. Submitted in a personal capacity.

Date: 12 October 2026

Summary

A perfect rent hedge would narrow a GPU lender's value range by about 4%; even on loans already operating, by less than a fifth. That is the central result of a working valuation and risk prototype I built for GPU-backed lending, alongside public delivery data for eight AI-infrastructure companies. Compute can support standardised derivatives, but only if the market agrees on what is being priced, settlement sources disclose how they adjust for it, and participants understand what a rent contract leaves uncovered. Four findings follow.

1. **Standardisation is achievable but must be disclosed.** Two GPU-hours of the same chip are not the same product. Region, contract length, cluster size and interconnect change the price. A benchmark is only comparable if it publishes the adjustments it applies.
2. **Settlement data needs a minimum disclosure standard.** Public capacity data today uses at least six different units across eight companies, and two disclose no delivered figure at all. Prices and delivery figures that cannot be traced to a source cannot be surveilled.
3. **Rent contracts leave most of a lender's risk uncovered.** In my prototype, forward rents explain about 30% of value variance on operating loans, and about 4% on a loan still being built. Residual value, delivery timing, renewal, lien ranking and control over where chips can be used account for the rest. Market participants should understand that gap.
4. **Perpetual contracts face a generation problem.** Each new chip generation reprices the one before it. A perpetual contract on a single chip drifts from a constant product to a declining one.

Evidence base

The evidence is a working prototype, Halo, built in 2026 as an independent design study of how lenders value and monitor GPU-backed loans. It is not a commercial product and I have no position in any compute contract or index.

- **Deals and resale prices are synthetic and labelled as such.** Nine illustrative loans and their resale prices are generated, not taken from any firm.

- **Real figures are cited.** Company delivery figures come from results releases and SEC filings for the latest reported quarter, ended May or June 2026. Rent levels start at the median of public September 2026 GPU rental indices (29 cited observations, April to September), so the synthetic loans are valued against observed market rents. The power price is the 2025 German day-ahead average published by Bundesnetzagentur.
- **Methods are disclosed.** Each valuation input carries its source, age and a confidence score, or is shown as absent. Three residual-value methods run side by side and are scored out of sample.

The findings below are about structure: what the market must define and disclose. They are not estimates of real market values.

The prototype is public at <https://halo.olaologunde.com/cftc-2026/>, a fixed copy matching the figures in this comment, with a guide from each recommendation to the screen that supports it.

Cash market: what is being priced

The physical compute market can be standardised, but today it lacks a common definition of both the product and its delivery.

The product. A GPU-hour is not fungible across contracts. In the prototype, a deal's rent is compared with an index only after explicit adjustments for region, contract length, cluster size and interconnect. Without those adjustments, a lender cannot tell whether a contract is priced above or below market. A benchmark that publishes its specification and adjustment rules makes contracts comparable; one that does not leaves each user to rebuild them privately.

Delivery. Contracted capacity is not delivered capacity, and companies report the two differently. The table shows how eight companies reported it for their latest quarter, ended May or June 2026.

Company	Unit reported	Delivered / contracted	Source
CoreWeave	MW active vs contracted power	1,500 / 3,700 (41%)	Q2 2026 results
Galaxy (Helios)	MW critical IT	133 / 526 (25%)	Q2 2026 release, SEC
Applied Digital	MW critical IT	175 / 1,410 (12%)	FY2026 results
TeraWulf	MW critical IT, contracted derived	102 / 839 (12%)	Q2 2026 results, SEC
IREN	\$bn annualised revenue, not MW	1 / 4 (25%)	FY26 results
Crusoe	MW operational vs gross contracted; private, unaudited	1,000 / 6,000 (17%)	Series F release
Nebius	MW contracted; connected not disclosed	Not disclosed / 3,500	Q2 2026 shareholder letter

Company	Unit reported	Delivered / contracted	Source
Cipher Digital	Not disclosed in MW	Not disclosed	Q2 2026 business update

Six different bases across eight companies, and two with no delivered figure. Lessor and tenant figures also overlap (some of these companies lease capacity to others on the list), so the figures cannot be summed into a market total. Even "a megawatt delivered" has no shared definition.

Oversight and manipulation: a disclosure standard for settlement sources

A settlement price is only as surveillable as the data behind it. Compute benchmarks are already used to settle contracts: for example, a contract self-certified on 20 August 2026 asks whether a GPU rental price index for the H100 will be at least \$2.50 per GPU-hour on 15 March 2027, and an over-the-counter block trade in June 2026 settled against a transaction-based H100 rental index. Most of the underlying transactions are bilateral and unpublished. Two listed contracts are now in train: NYMEX has filed Silicon Data H100 and B200 rental-index futures, under extended review to 9 November, and ICE and Ornn have announced cash-settled futures on a transaction-only index. Two settlement sources built differently is exactly where a disclosure standard matters; neither lists the A100, so lenders to older fleets keep basis risk that a published roll rule would address.

In the prototype, every figure must show its source, its age and a confidence score, or appear as absent. Applying that discipline to settlement sources suggests a minimum disclosure for any index that settles a listed compute contract:

1. **Specification.** The chip, configuration, region, term and cluster size the index represents.
2. **Adjustments.** The rules that convert non-standard transactions into that specification, with their current values.
3. **Observations.** The number of transactions or quotes per fixing, and how many came from each of the largest contributors.
4. **Contributor concentration.** Whether any contributor is also a large provider or user of the compute being priced.
5. **Gaps.** What the index does when observations are too few, and how often that has happened.
6. **Track record.** How the index has compared with later observed transaction prices.

Points 3 and 4 matter most for manipulation risk. A thin index fed mainly by one provider can be moved by that provider.

Customer protection: what a rent contract does not cover

A lender who hedges rent may believe its compute exposure is covered when most of it is not. Across the prototype's eight operating synthetic loans, forward rents explain about 30% of value variance (15% to

47% per loan). In range terms, even a perfect rent hedge narrows the book's value range by about 4%, because the loan still being built is dominated by delivery risk. The remainder comes from five sources no rent contract reaches.

The gap matters because lenders now hold much of this exposure. On 30 September 2026 the Bank of England's Financial Policy Committee recorded AI-related debt issuance of about \$450 billion by early September, an estimated \$700 billion of 2026–28 data-centre investment to be financed by private credit, and leveraged positions that amplified July's fall in AI stocks. Lenders are already managing this risk through loan structure rather than hedges: a \$5bn GPU-backed loan package launched by JPMorgan and Goldman Sachs on 9 October 2026 is priced at about 11% and amortised to repay in full 5.5 years after funding. A lender that believes a rent hedge covers this exposure is the customer most likely to be surprised.

Uncovered risk	What it is	Prototype example (synthetic)
Residual value	What the chips are worth at the end of the loan	Three accepted methods value the same H100s at 40%, 35% and 23% of cost at maturity
Delivery	Contracted capacity that is late or never energised	On a loan still being built, energisation timing explains 64% of value variance; rent explains about 4%
Renewal	Whether the tenant renews when the contract ends	Value that depends on renewal clusters in a few tenants and dates
Lien ranking	Where a lender stands if a public or senior lender ranks ahead	Loan-to-value of 66% against the covenant, but 107% counting debt that ranks first
Control	Whether chips can be resold, relocated or remotely restricted under export rules or vendor terms	An illustrative 25–30% cut to resale value for current-generation chips whose resale market is restricted

The last two rows are emerging risks. Public lenders are entering this market: the Department of Defense's Office of Strategic Capital is reported to be in talks to lend \$5bn to an AI-infrastructure company. A covenant that measures one loan against all the collateral cannot see debt that ranks ahead of it. Control matters for a similar reason: proposed US legislation would require location verification for advanced chips, which ties a chip's resale value to where it is allowed to run.

I suggest the Commission's customer-protection work asks venues and intermediaries to describe clearly which risks a compute contract does and does not hedge, using categories like these.

Perpetual compute futures: the generation problem

A perpetual contract assumes the underlying stays the same product. A single chip does not. Its rental price steps down each time the next generation reaches volume, then decays as it ages. In the prototype, an H100 reference level loses about 12% a year and a further 10% at the next launch; after 24 months it

sits near 69% of today. On loans with fixed rents, launch timing moves value by 3% or less. On the index a perpetual tracks, the launch step is the largest single move.

A perpetual on one chip therefore tracks a declining asset, and the step-downs are predictable from public launch schedules. Two designs would reduce this:

- **A rolling generation reference.** The contract references "current-generation" capacity and rolls to the next chip on a published rule.
- **Dated contracts per generation.** Expiries are set so that each contract settles before its chip's expected successor reaches volume.

Either way, the roll rule should be disclosed in advance, because the roll date itself is information a large provider may hold before the market.

Recommendations

1. Require any index used to settle a listed compute contract to publish its specification and adjustment rules.
2. Require per-fixing observation counts and contributor concentration, including whether contributors are large compute providers or users.
3. Encourage a common definition of delivered capacity (for example, energised critical-IT megawatts, with power and revenue measures reported alongside, not instead).
4. Ask venues to state plainly which risks a compute contract hedges and which it does not: residual value, delivery, renewal, lien ranking and control (export limits and remote restriction).
5. For perpetual or long-dated contracts, require a published rule for rolling between chip generations.

I would be glad to share the prototype and its data definitions with Commission staff. The public copy already offers its conventions (JSON) and delivery ledger (CSV) for download.

Disclosures and sources

I hold no position in any compute contract, index or company named here, and this comment is not made on behalf of any employer or client. I am also the founder of Human Node, a cognitive-security company; this comment is personal and unrelated to it. Prototype figures are synthetic and illustrative.

- [CFTC press release: Request for Comment on Compute Derivatives](#)
- [Sidley: CFTC Seeks Comment on Listing of Compute Derivatives Contracts](#)
- [DeFi Rate: Polymarket US files AI compute contracts](#)
- [Reuters via Yahoo Finance: Pentagon in talks to lend \\$5bn to Fluidstack](#)
- [Bank of England: Financial Policy Committee summary and record, September 2026 \(published 30 September 2026\)](#)

- [PitchBook: Volta launches \\$5bn leveraged loan package to support data center lease, GPU capital expenditures, 9 October 2026](#)
- [Bundesnetzagentur: 2025 wholesale power prices](#)
- Company delivery sources are linked in the cash-market table.
- [Code on Trial: Pricing the Unobservable](#) (describes the 20 August 2026 self-certification)
- [PR Newswire: Polymarket closes first institutional block trade on a GPU instrument, 2 June 2026](#)
- [CFTC Division of Market Oversight: extension of review, NYMEX Silicon Data H100 and B200 Rental Index Futures, 21 September 2026](#)
- [ICE and Ornn to launch GPU compute futures contracts, 19 May 2026 \(Mondo Visione\)](#)
- [Greenberg Traurig: Navigating GPU Export Controls and AI Use Restrictions in Data Center Operations, 22 September 2026](#)

Ola Ologunde · Independent product designer, London · olaologunde.com · ola@olaologunde.com