Comment of Laela Zorana, ComputeLab.co

Conflicts of Interest and Affiliations. RIN 3038-AF76.

91 FR 50926, published 6 August 2026. Docket CFTC-2026-1686. Comments due 5 October 2026.

1. Who I am

I have measured delivered work across four accelerator vendors, AMD, NVIDIA, Google and AWS, and published 68 runs with their toolchain and hardware so anyone can repeat them. The same nominal hour of that hardware rents between 2.68 and 7.28 dollars depending on who sells it, and the Commission’s own release cites an estimate of 430 billion to 1.3 trillion dollars a year for the compute service flow behind it.

I work with frontier AI labs on the compute behind their models. I tune GPU kernels and inference paths so a model serves more tokens per dollar, I benchmark accelerators and publish the runs, and I author reinforcement learning environments and benchmark problems used to train and grade models. I size the compute behind a dataset, which chips, how many, and what it costs. That sizing runs on the same rental prices these two contracts settle against, so the question of who calculates that number is one I meet in the work.

My kernels, runs and datasets are public under LaelaZorana on GitHub and Hugging Face. The benchmark numbers here are my own runs, on hardware I name. The rental prices and the market size come from published sources.

2. The terms this comment uses

An accelerator is the specialised chip that does the heavy work in artificial intelligence. People often call it a GPU. The H100 and the B200 are two such chips. A firm rents time on one by the hour, and a GPU-hour is one chip rented for one hour.

A kernel is the small program that runs directly on that chip and does the actual arithmetic. Two kernels can do the same job at very different speeds.

A settlement index is a daily number published by an outside firm. A cash settled contract pays against that number rather than against delivered goods.

3. The facts I am commenting against

I answer questions 30, 31, 35, 42 and 69.

On 11 August 2026 NYMEX filed submission 26-370 to list two compute futures contracts. They are the Silicon Data H100 Rental Index Futures and the Silicon Data B200 Rental Index Futures. Trading is set to begin 5 October 2026. Each contract pays out in cash. Each is sized at 730 GPU-hours and quoted in dollars per GPU-hour, so each one covers about a month of one chip.

Both contracts settle against a daily index published by Silicon Data, a private firm. CME Group and Silicon Data announced the launch on 11 August 2026, and that announcement states that Silicon Data is backed by the global trading firm DRW.

The Silicon Data methodology overview dated 27 August 2026 sets out how the number is made. Each day the index takes in more than 10,000 GPU rental records from more than 40 cloud providers. A pricing model owned by the firm then applies a basis adjustment to every record.

That adjustment is worth stating in ordinary terms. Rental deals differ from each other, so a raw average of them would carry no meaning. The adjustment restates each deal as though it were the one standard deal the contract is written on. Crude oil grades work the same way, and so do the par and non par quality rules in an existing physical contract. A barrel of one grade is converted to the reference grade before it counts. The crude adjustment is published, and this one stays inside the firm.

The methodology document names the four things the adjustment considers. They are rental type, country, CPU platform and GPU variant. The model itself, the coefficients inside it, and the size of the adjustment applied to any single record all remain the firm’s own. Adjusted prices are then averaged for each provider, and the daily index is a weighted average across providers. Each provider weight is said to reflect that provider’s market relevance, and those weights also remain the firm’s own. The pricing model is recalibrated once a week. The document is marked as customer material that requires written consent before it is shared.

So the settlement price of a listed futures contract is produced by a private firm. The inputs come from the sellers of the thing being priced. The step in between can be reproduced only by that one firm.

4. Question 30, the scope should reach the index administrator

Question 30 asks whether the scope should be widened to take in other types of entities. It should reach the administrator of a settlement index. The answer matters most where the underlying thing is calculated rather than delivered.

The Proposal draws its lines around intermediaries and around firms that trade. That is the right map for a market where the price comes from trades on the exchange itself. On a cash settled contract written on an outside index, the price comes from the administrator’s model rather than from the order book. A firm can hold no position, hold no membership and place no order, and still set the number that every open contract pays against. The conflicts framework in this Proposal stops short of that firm. On these two contracts that firm is the control point.

The second half of the problem is the input side. The records feeding the index come from the cloud providers who sell the capacity. A provider that posts its rates moves a settlement price without placing a single trade. Averaging within each provider before weighting across providers raises that risk further. It gives a small seller the same starting footing as a large one, and an unpublished weight then decides how much each one counts.

My proposal here is a narrow one. Where a registered entity lists a contract that settles against an outside index, the conflicts rules should follow the settlement input to the firm that produces it. Index vendors in general can stay outside the rule.

5. Questions 31 and 42, control is the wrong threshold here

Questions 31 and 42 ask whether the control formulation captures the right relationships, and whether part ownership below control should be reached. It should, for the class of relationship now in front of the Commission.

Picture a trading firm that holds a minority stake in an index administrator, with a board seat and a commercial agreement alongside it. That firm falls outside control on any ordinary reading, and it also has no need of control. What it wants is influence over an unpublished basis adjustment and an unpublished provider weight, with the model behind them recalibrated every week. A control test written around voting securities passes straight over that relationship and records nothing at all.

My proposal is a second limb keyed to the settlement function rather than to ownership. Suppose a registered entity lists a contract settling on an outside index. Any firm holding equity, a board seat, a revenue share or an exclusive data agreement with that index administrator should be in scope for the conflicts rules, whether or not the holding rises to control.

The reason for the lower threshold is simple. A trading affiliate affects the positions it trades against. A settlement input reaches every open position at once.

6. Question 35, the disclosures that would carry weight

Question 35 asks what further disclosures should be required. Four would do real work where the underlying thing is calculated.

1. The ownership chain of the settlement index administrator, including holdings below control, and any revenue or data agreement between the administrator and a firm that trades the contract.
2. The published rule for the basis adjustment, set out well enough that a person outside the firm can recompute a published index value from the input records.
3. The provider weights in force on each index date, and the date of each weekly recalibration.
4. The number of independent contributors behind each settlement value, so that a thin day is visible on the face of the number.

Question 69 asks a version of the same thing for clearing, and my answer there is the same. Upstream ownership should be disclosed, and the disclosure should reach part holdings that fall short of control. Those are the holdings that carry influence at this layer.

7. Who has to be able to read the evidence

One point sits underneath all of the above. The evidence behind a settlement number on this product is a measurement, and reading a measurement takes someone who has made one.

Here is a measurement I made. I timed a single kernel on an AMD Instinct MI300X accelerator, running under the software stack ROCm 7.2.4. The same job took 86.3898 milliseconds in one version of the code and 0.1015 milliseconds in another. That is a factor of 851. The chip did not change and the run did not change, and only the code changed.

What that number means for the Commission is this. Two ways of writing the same job on one chip can differ by that much. A basis adjustment that standardises across chips and rental types is doing its work in a market where differences of that size are possible, inside a model nobody outside the firm can see. Someone who has timed a kernel reads the risk in that at once. A reader holding only the contract terms sees a chip name and a price, while the evidence sits in a timing trace.

So a firm that lists, designs or settles a compute product should hold that competence at the level where decisions are made, rather than on the payroll below it. An employee can sign a specification and can still be overruled.

A degree or a certificate should stay outside the test, because paper can be bought. What should count is work someone outside the firm can check for themselves. I will provide the criteria to Commission staff on request.

8. Offer to the Commission

I am willing to run measurements for Commission staff on any settlement design under consideration, on the accelerators named in it, and to provide the raw records on request so that staff can check them.

Respectfully submitted,

Laela Zorana, High-Performance Computing Engineer

ComputeLab.co

laela@computelab.co

Patent pending, United States provisional application 64/160,345, filed 23 September 2026. Copyright 2026 Laela Zorana. All rights reserved.

The measurement methods and specifications referenced in this comment are my own work and remain my property. Publication of this comment in the docket grants no licence to use them, and I reserve all rights in them.