

October 20, 2026

VIA ELECTRONIC SUBMISSION

Mr. Christopher Kirkpatrick
Secretary of the Commission
Commodity Futures Trading Commission
Three Lafayette Centre
1155 21st Street, N.W.
Washington, DC 20581

Re: Request for Comment on the Listing of Compute Derivatives Contracts (RIN 3038-AF77)

Dear Mr. Kirkpatrick:

Hyperliquid Policy Center (“HPC”) appreciates the opportunity to respond to the Request for Comment on the Listing of Compute Derivatives Contracts¹ issued by the Commodity Futures Trading Commission (“CFTC” or the “Commission”). HPC commends the Commission for undertaking this effort to bring needed risk-management instruments to a critical and fast-growing sector of the global economy, and for the contribution that effort makes to American leadership in financial innovation and in artificial intelligence.

HPC is an independent research and advocacy organization dedicated to advancing a clear, regulated path for Americans to access onchain markets, including those available on Hyperliquid.² Hyperliquid is best known for the trading of perpetual futures, and HPC works to advance policy that onshores and expands the market for perpetual futures, including to new commodity types.

Artificial intelligence has rapidly become a strategic sector of the U.S. economy, and the computing capacity that trains and powers AI models has become one of its scarcest inputs. Private investment in AI data centers, compute hardware, and networking equipment reached roughly 0.8 percent of U.S. GDP in the first quarter of 2026.³ We share the Administration’s view that “[w]hoever has the largest AI ecosystem will set global AI standards and reap broad economic and military benefits,” and “it is imperative that the United States and its allies win this race.”⁴ The Administration’s AI Action Plan ties that outcome explicitly to the financial market for compute. It observes that a company seeking large-scale compute today “must often sign long-term contracts with hyperscalers,” that “America has solved this problem before with other goods through financial markets, such as spot and forward markets for commodities,” and that the Federal government “can accelerate the maturation of a healthy financial

¹ Request for Comment on the Listing of Compute Derivatives Contracts, 91 Fed. Reg. 54,259 (Aug. 21, 2026) (RIN 3038-AF77) (the “RFC”), <https://www.federalregister.gov/documents/2026/08/21/2026-17163/request-for-comment-on-the-listing-of-compute-derivatives-contracts>.

² Hyperliquid Policy Center, *About*, <https://hyperliquidpolicy.org/about> (last visited Oct. 19, 2026).

³ Isabel Juniewicz, *The AI Boom Has Doubled Computing Infrastructure’s Share of US GDP*, Epoch AI (June 5, 2026), <https://epoch.ai/data-insights/ai-datacenter-share-gdp>.

⁴ The White House, *Winning the Race: America’s AI Action Plan*1 (July 2025), <https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf> (“AI Action Plan”).

market for compute.”⁵ Establishing markets that allow participants to transfer compute price risk will be important to advancing those objectives, and we support the Commission’s efforts to do so “in a manner that promotes transparency, liquidity, and responsible innovation.”⁶ As Chairman Selig has stated, “America cannot win the AI race without a robust derivatives market for compute.”⁷

One feature of the compute market is particularly relevant to our RFC responses: compute is a commodity that is produced and consumed continuously. A GPU-hour exists only in the hour it is used. It cannot be stored, inventoried, or delivered later, and the data centers that produce it run every hour of every day. Compute consumers face the same conditions on the demand side. Developers train and query models around the clock, and software agents may increasingly procure and pay for compute without a human in the loop.⁸ A commodity with those features can benefit from instruments that are themselves continuous, such as perpetual contracts. This letter accordingly addresses two questions: whether compute exhibits the features of a commodity that can support a derivatives market, and which contract designs, including perpetual contracts, would best serve participants who bear compute price risk.

I. The Compute Market

The market for compute has developed rapidly over the past few years, and it has several structural characteristics that bear on the questions the RFC poses and on how derivatives referencing compute should be introduced.⁹

A. The supply chain and its participants

Compute, the processing capacity used to train and run AI models, reaches its end users through a stack of businesses that supply energy, chips, infrastructure, models, and applications.¹⁰ At the base of this stack are the chip designers, principally NVIDIA, whose GPUs account for the large majority of high-end AI compute, and the foundries that fabricate them. Above them are the data-center developers that house the chips and the power supply that has become their binding constraint. Large AI campuses require substantial power capacity, and securing electricity now influences construction as much as securing chips does.¹¹ Two kinds of providers own and operate most of the hardware. The hyperscalers (e.g., Amazon, Microsoft, and Alphabet) build capacity for their own models and products and rent the remainder, and a newer class of independent “neocloud” operators rents GPU capacity to third parties. Inference platforms sit above the providers. They take raw rented capacity and package it as on-demand access to a model, so that a developer can send a request and receive a result without operating any hardware (“inference” is the act of running a trained model to answer a query, as distinct from training the model in the first place). The application layer, the software products that consumers and businesses use,

⁵ AI Action Plan, *supra* note 4, at 4–5.

⁶ RFC, 91 Fed. Reg. at 54,260.

⁷ Press Release, CFTC, CFTC Requests Comment on the Listing of Compute Derivatives Contracts (Aug. 19, 2026), <https://www.cftc.gov/PressRoom/PressReleases/9286-26>.

⁸ BlackRock Digital Assets Research, *The Machine-Native Economy: How Digital Assets Connect Intelligence, Commerce, and Compute* 5, 9 (Sept. 2026), <https://www.blackrock.com/us/individual/literature/whitepaper/the-machine-native-economy.pdf> (“BlackRock Report”).

⁹ RFC, 91 Fed. Reg. at 54,263 (Questions 1(a)–1(e)).

¹⁰ RFC, 91 Fed. Reg. at 54,262 n.42.

¹¹ See Russell Brandom, *The Billion-Dollar Infrastructure Deals Powering the AI Boom*, TechCrunch (Feb. 28, 2026), <https://techcrunch.com/2026/02/28/billion-dollar-infrastructure-deals-ai-boom-data-centers-openai-oracle-nvidia-microsoft-google-meta/> (describing multi-gigawatt data-center campuses); BlackRock Report, *supra* note 8, at 7–8.

buys that access and pays by the token, the unit in which a model measures the text it reads and produces.¹² Between the layers, brokers, marketplaces, and over-the-counter desks match demand for a specific configuration with the supply that can serve it, much of it still through bilateral negotiation.

Compute prices reflect both what it costs providers to supply compute and what developers will pay to use it. On one side are the costs of producing compute, chiefly chips, electricity, and data-center space. On the other is the value of what compute produces, the training runs and inference workloads that generate revenue for the businesses above it in the chain.¹³

B. How compute is bought and where price risk sits

Buyers acquire compute in four principal ways, and the method of purchase determines whether the price is readily observable:¹⁴

- **On-demand.** Capacity rented by the hour at a published rate or a marketplace price, without a long-term commitment.
- **Spot.** On-demand capacity that the provider may reclaim on short notice, sold at a discount.
- **Reserved.** Term or spending commitments, typically one to three years, at a negotiated rate; guaranteed capacity depends on the agreement.
- **Committed purchase.** Multi-year take-or-pay agreements under which a buyer commits to a volume of capacity, often before the data center that will supply it is built.

Most compute changes hands at prices that are not publicly observable. On-demand and spot prices are at least partly observable on marketplaces and in providers' published price lists. Providers and buyers negotiate reserved and committed prices privately and do not publish them, and most of the economic value in the market transacts on those terms.¹⁵ Unlike certain wholesale electricity markets, compute has no comparable market-wide U.S. transaction-reporting regime.¹⁶ As a result, equivalent capacity trades at widely different prices across providers, regions, and contract terms; one published price series shows the largest cloud providers charging roughly three to four times the marketplace rate for the same chip over the same months.¹⁷

Compute prices have moved sharply over the market's short history. Marketplace rental rates for NVIDIA's H100 fell from approximately \$2.50–\$3.19 per GPU-hour in the second half of 2024 to \$1.92–\$2.00 in the second half of 2025. The publisher of that price series attributes the decline to new chip supply, the arrival of marketplaces that made prices visible, price cuts by the largest providers, and the shift of demand from training toward inference.¹⁸ Hardware generations also depreciate against each other, with older chips renting at a fraction of the rate the newest generation commands. An older generation does not lose its value all at once, however. A chip that is no longer competitive for training the

¹² BlackRock Report, *supra* note 8, at 8–9.

¹³ RFC, 91 Fed. Reg. at 54,262 n.42.

¹⁴ RFC, 91 Fed. Reg. at 54,263 (Question 1(a)).

¹⁵ RFC, 91 Fed. Reg. at 54,263 (Questions 1(b) and 1(c)).

¹⁶ *See, e.g.*, 18 C.F.R. § 35.10b (electric quarterly reporting); RFC, 91 Fed. Reg. at 54,263 (Questions 1(c) and 1(e)(ii)).

¹⁷ Carmen Li, *H100 Rental Price Over Time (2023–2025): A Complete Market Analysis*, Silicon Data (Dec. 21, 2025), <https://www.silicondata.com/blog/h100-rental-price-over-time>.

¹⁸ *Id.* §§ 1, 3, 6 (provider-category price series and the drivers of the 2024–2025 decline).

largest models can remain in demand for inference and other less demanding workloads, and one index provider finds that older chips can remain economical for particular workloads.¹⁹

Price risk sits at several points along the compute supply chain, which creates natural hedging interest on both the supply and demand sides of the market. Infrastructure owners and their lenders bear the risk of declines in rental rates and utilization. A provider that financed hardware against projected rental revenue, and the lender holding that hardware as collateral, both lose when the rate falls. Developers, inference platforms, and application businesses are exposed to increases in the cost of compute, which the RFC describes as “one of the most material costs of AI that affects the entire AI production stack.”²⁰ Hyperscalers operate at several layers of the stack and can net part of this exposure internally, but they remain exposed on the capacity they rent to third parties and on multi-year hardware commitments. Brokers and marketplaces face basis risk: they sell customers a specific configuration (a particular chip, in a particular region, with particular interconnect and service terms) while sourcing supply priced against a generic benchmark, creating exposure when the spread between the two prices moves against them. Power markets have a similar structure—electricity is priced at many individual delivery points, and dealers there hedge the generic hub price on exchange while pricing the location-specific difference to their customers.²¹

C. *Early financialization*

A financial layer has begun to form around compute, and it generally resembles the layers that formed around other commodities. Spot marketplaces and brokers arrived first, and they remain the venues on which observable transaction prices print. Long-term reservation and take-or-pay agreements function as the market’s forward layer, and the academic literature has used term rental prices to construct indicative forward curves, subject to differences between term-rental and futures contracts.²² Price indices are the newest layer. Several providers now publish daily reference prices for rented compute, and they differ in construction. Some draw on providers’ published rates, some only on completed transactions, and aggregated feeds of the kind Pyth publishes in other markets offer another approach, combining independent price contributions into a reference price with a published confidence interval.²³

The choice among pricing benchmarks is an important consideration for contract design. Because a derivatives contract must settle to a single reference price, the index that produces that price, in effect, determines which transactions are treated as interchangeable (i.e., which chips, in which regions, on what rental terms, and with what service levels). A participant whose exposure matches those parameters can hedge with the contract directly. One whose exposure differs bears the difference as basis

¹⁹ Ornn Data, *The Economics of Open-Weight Inference* 2, 17–19 (Sept. 7, 2026), <https://data.ornn.com/the-economics-of-open-weight-inference.pdf>.

²⁰ RFC, 91 Fed. Reg. at 54,262.

²¹ See, e.g., U.S. Energy Info. Admin., *Wholesale Power Price Maps Reflect Real-Time Constraints on Transmission of Electricity* (Sept. 20, 2011), <https://www.eia.gov/todayinenergy/detail.php?id=3150>.

²² Federico M. Bandi & Yinan Su, *(Early) AI Compute Asset Pricing*, abstract (v3, Sept. 10, 2026), <https://arxiv.org/abs/2607.12156> (treating synthetic futures prices inferred from term rental contracts as likely upper bounds on true futures prices).

²³ NYMEX Submission No. 26-370 to the CFTC 1-3, 5-6 (Aug. 11, 2026), <https://www.cftc.gov/filings/ptc/ptc08112615405.pdf>; Press Release, Intercontinental Exchange, *ICE and Ornn to Launch GPU Compute Futures Contracts* (May 19, 2026), <https://ir.theice.com/press/news-details/2026/ICE-and-Ornn-to-Launch-GPU-Compute-Futures-Contracts/default.aspx>; ICE Futures U.S., Exchange Notice, *New Compute Index Reference Data to Be Made Available on Monday, October 5, 2026* (Sept. 29, 2026), https://www.ice.com/publicdocs/futures_us/exchange_notices/ICE_Futures_US_Exchange_Note_ORNN_Compute_4Q2026_20260929.pdf; Pyth Developer Hub, *Understanding Price Data*, <https://docs.pyth.network/price-feeds/pro/understanding-price-data> (last visited Oct. 19, 2026).

risk, and two indices that define compute differently may diverge enough that a contract settling to one would be of little use in hedging exposure to the other.²⁴

Registered U.S. futures exchanges have moved to list contracts that settle to compute price indices. NYMEX has submitted two cash-settled monthly futures for voluntary Commission approval under Regulation 40.3, with the effective listing date to be announced, each representing 730 GPU-hours and settling to a Silicon Data index.²⁵ ICE Futures U.S. has announced cash-settled futures based on Ornn's indices. Though none of these contracts has begun trading, the two announcements demonstrate interest in compute derivatives from registered exchanges and the market participants they serve.

Beyond conventional futures exchanges, newer venues have begun to offer perpetual and event contracts on compute prices. Architect, for example, announced margined perpetual contracts on compute indices for institutions in eligible jurisdictions through its Bermuda-licensed venue, acquired a U.S. designated contract market ("DCM") with the stated purpose of listing compute derivatives, and announced an exchange-for-physical network intended to connect listed contracts to physical capacity delivery. Other venues, such as Kalshi, have listed event contracts on GPU rental prices.²⁶

Compute benchmarks and listed contracts appear to be developing before the cash market has standardized, which inverts the usual order. In most commodity markets a standard grade and a reference price emerged in the cash market first, and futures contracts followed. In compute, the choice of settlement benchmark is itself doing part of the work of standardization.²⁷ That is unusual but not unprecedented. Henry Hub became the benchmark for American natural gas in large part because NYMEX chose it as the delivery point for its 1990 futures contract.²⁸ The listed contract helped give the cash market a standard convention to converge on.

Part of compute's emerging financial layer also lives onchain. On October 6, 2026, MarketVector announced that Paragon, a deployer of perpetual futures on the Hyperliquid blockchain, had licensed a continuous version of MarketVector's flagship U.S. semiconductor benchmark to underlie perpetual contracts that reference the index directly, with Pyth supplying the real-time pricing that extends the index into overnight and pre- and post-market sessions.²⁹ Stablecoin-funded lending against GPU collateral, tokenized inference credits, and inference networks that settle onchain are operating today, and one industry report describes them as converging with GPU indices and futures into an integrated

²⁴ Rubén Fernández-Fuertes & Vincent Grégoire, *Compute Futures: Who Is Left Exposed?* (Aug. 3, 2026), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=7224223.

²⁵ NYMEX Submission No. 26-370, *supra* note 23, at 1-3, 5-6; CME Group, *Revised New Product Summary: Initial Listing of Two (2) Compute Futures Contracts, Effective Date to Be Announced* (Sept. 25, 2026), <https://www.cmegroup.com/notices/clearing/2026/08/26-274.html>; ICE and Ornn to Launch GPU Compute Futures Contracts, *supra* note 23.

²⁶ Architect, Press (announcements of Jan. 21, May 28, and July 8, 2026), <https://architect.co/insights/press/>; Kalshi, *Compute Has a Market-Implied Forward Curve Now* (July 14, 2026), <https://news.kalshi.com/p/gpu-compute-forward-curve> (both last visited Oct. 8, 2026).

²⁷ Fuli Yang, *Futures Before Fungibility: Compute Derivatives and the Financialization of an Unstandardized Market*, abstract (Aug. 20, 2026), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=7324139.

²⁸ Adila McHich, *Henry Hub Natural Gas Futures: Global Benchmark*, CME Group (July 7, 2023), <https://www.cmegroup.com/articles/whitepapers/henry-hub-natural-gas-futures-global-benchmark.html>.

²⁹ Press Release, MarketVector Indexes, *MarketVector Brings Flagship Semiconductor Benchmark Onchain with Paragon* (Oct. 6, 2026), <https://www.businesswire.com/news/home/20261006879151/en/MarketVector-Brings-Flagship-Semiconductor-Benchmark-Onchain-with-Paragon>.

“inference capital market.”⁵⁰ BlackRock’s digital assets research describes compute as “a distinct, large, and increasingly investable economic resource that could support a new class of digital assets.” BlackRock further observes that onchain markets “may enable more granular regional and hardware-specific contracts within shared settlement infrastructure,” with claims on compute capacity that can be “represented, transferred, pledged as collateral, and settled through programmable infrastructure.” Onchain settlement, on that view, is what allows a benchmark contract and the basis contracts around it to share one margin and settlement system.⁵¹

D. Features that distinguish compute, and analogies to other commodity markets

Compute has attributes that set it apart from the commodity markets the Commission has customarily regulated, and the Commission will need to account for each of them in evaluating contract design and oversight.⁵² These include that:

- **Compute cannot be stored.** A GPU-hour that goes unused is lost, so there is no inventory to carry and no cost-of-carry relationship between spot and forward prices, although a contract could physically deliver future access to specified capacity.
- **Hardware depreciates on a manufacturer-set schedule.** Chip makers now release new generations on roughly an annual cadence, so a contract on a named chip references hardware that will be superseded during a multi-year listing cycle.⁵³
- **Demand moves in jumps while supply moves slowly.** Model releases and new workload types reprice compute in a matter of weeks, while new supply arrives after the time needed to build data centers, secure power, and expand chip output.
- **Hardware performance is variable.** Configuration, interconnect, power, software, and service terms change what a given chip is capable of delivering, so a single derivatives contract cannot capture every variation a participant might want to hedge. Grade and location differentials are a familiar problem in commodity markets, and they are the reason basis contracts have formed alongside benchmark contracts.⁵⁴
- **Supply is concentrated.** A single producer dominates the chip layer and a handful of providers dominate the cloud layer, which the RFC identifies as a source of potential pricing power.⁵⁵
- **The compute supply chain is geopolitically sensitive.** Export controls govern where advanced AI chips may be sold, and energy policy governs where they can be powered.⁵⁶

⁵⁰ Lucas Tcheyan, *Inference Capital Markets: How Crypto Is Building the Financial Layer for AI Compute*, Galaxy Research (July 13, 2026), <https://www.galaxy.com/insights/research/inference-capital-markets-ai-compute-gpu-futures-onchain-crypto>.

⁵¹ BlackRock Report, *supra* note 8, at 7–8.

⁵² RFC, 91 Fed. Reg. at 54,263 (Question 1(e)).

⁵³ NVIDIA, *GTC 2025 Announcements and Live Updates* (Mar. 25, 2025), <https://blogs.nvidia.com/blog/nvidia-keynote-at-gtc-2025-ai-news-live-updates/>.

⁵⁴ Fernández-Fuertes & Grégoire, *supra* note 24; BlackRock Report, *supra* note 8, at 8.

⁵⁵ RFC, 91 Fed. Reg. at 54,262 (observing that “dominant market participants may wield significant pricing power”); *id.* at 54,263 (Questions 1(e)(iii) and 1(f)).

⁵⁶ RFC, 91 Fed. Reg. at 54,264 (Question 3(b)); Bureau of Indus. & Sec., *Guidance Regarding Enforcement of License Requirements for Advanced Computing Items for Entities Headquartered in Country Group D:5 and Macau* (May 31, 2026), <https://www.bis.gov/media/documents/bis-guidance-may-31-2026.pdf>.

Each of the attributes that distinguishes compute has a counterpart in a commodity market the Commission already oversees. Electricity is the closest analogue. It is a service that cannot be stored and is consumed as it is produced, and its price forms at many physical delivery points that a handful of hub benchmarks aggregate into the reference prices for exchange-traded contracts, an arrangement that dates to the restructuring of the 1990s.³⁷ The Commission has itself observed that electricity “must be used or consumed on an as-needed basis,” and the same is true of a GPU-hour.³⁸ Brent crude is the precedent for a benchmark that outlives the reference beneath it. Brent began as the price of oil from a single North Sea field. As that field’s output declined, the industry widened the deliverable basket to other grades, and the benchmark remained an important pricing reference while the physical commodity underneath it changed almost entirely.³⁹ Compute benchmarks may similarly adapt to hardware changes, although each transition must account for its effect on the exposures being hedged.

The participants with compute exposure to hedge are therefore numerous and varied. Providers and their lenders carry rental revenue and collateral value that fall when the price of compute falls. Developers, inference platforms, and application businesses carry costs that rise when it increases. Brokers and marketplaces carry the spread between the configurations they sell and the benchmark they can hedge. Hyperscalers carry some of each. And a growing number of participants, from startups budgeting for inference to software agents buying it by the token, carry compute exposure indirectly through the products they build and use.

II. Defining the Compute Commodity

A compute derivative must specify the commodity it references, and the market has not yet converged on a single specification. Section 1a(9) of the Commodity Exchange Act (“CEA”) defines “commodity” to include an enumerated list of agricultural products together with “all other goods and articles except onions...and motion picture box office receipts” and “all services, rights, and interests...in which contracts for future delivery are presently or in the future dealt in.”⁴⁰ The final clause defines commodity status by reference to trading rather than by reference to the nature of the asset. That drafting gives the definition an expansive reach, capturing any asset in which a futures market may develop, including assets that did not exist when Congress wrote the statute. Courts and the Commission have read the language accordingly, extending it to assets as far from the enumerated list as virtual currencies.⁴¹ Rented processing capacity, measured in GPU-hours as discussed below, fits comfortably within that definition, and, as Part III explains, the market for it has exhibited many of the features that are historically predictive of successful futures markets.

The RFC treats the commodity underlier as “access to rented compute capacity from hardware the purchaser of such capacity does not own (e.g., the hourly rental price of compute from a B200).”⁴² This letter focuses primarily on that form of compute, the GPU-hour, which the announced exchange contracts also reference. A GPU-hour is the hourly rental of a specified chip, produced and consumed in the same instant, and it cannot be stored. The announced contracts settle in cash against an index,

³⁷ Fed. Energy Regul. Comm’n, *Electric Power Markets*, <https://www.ferc.gov/electric-power-markets> (last visited Oct. 8, 2026).

³⁸ RFC, 91 Fed. Reg. at 54,262 n.45.

³⁹ Energy Intelligence, *Global Crude Benchmarks: Brent Sets the Standard*, Intercontinental Exchange, <https://www.ice.com/evolution-of-brent-its-markets-and-why-its-ecosystem-is-relied-upon-by-commercial-participants> (last visited Oct. 19, 2026).

⁴⁰ CEA § 1a(9), 7 U.S.C. § 1a(9).

⁴¹ See *CFTC v. McDonnell*, 287 F. Supp. 3d 213, 228–29 (E.D.N.Y. 2018) (holding that virtual currencies fall within the CEA’s definition of “commodity”); *In re Coinflip, Inc.*, CFTC No. 15–29, at 3 (Sept. 17, 2015) (finding that bitcoin and other virtual currencies are commodities under the CEA).

⁴² RFC, 91 Fed. Reg. at 54,262.

although non-storability does not itself preclude physical delivery of future access to specified capacity. Because the same chip rents on different terms in different places, an index must standardize several variables, including chip, region, term, interruptibility, cluster size, and service-level terms, with the consequence that index selection becomes an economic term of the contract. Three features of an index bear most directly on that selection. The first is how the index handles hardware migration when the reference chip is superseded. The second is whether it draws on posted rates or on printed transactions. The third is whether the Commission and the listing exchange can observe the transactions behind it.⁴³ Parts III and IV address each.

The RFC also recognizes that “the underlier may be a different type of compute-related commodity, such as access to a stated volume of LLM inference tokens.”⁴⁴ HPC supports that view. Inference is the act of running a trained model to answer a query, and the token is the unit in which a model meters the text it reads and produces. Tokens are the form in which many businesses and consumers actually purchase compute, since an application developer pays an inference platform by the token rather than renting hardware, and a growing share of inference demand comes from software agents that consume tokens without a human intervention. A token’s price reflects hardware costs and model efficiency together, so a token contract would serve application businesses and inference platforms whose costs are denominated in tokens, participants whom a GPU-hour contract would hedge only imperfectly. Token contracts raise standardization questions of their own, chiefly the capability threshold that makes tokens from different models comparable and the fungibility of tokens across providers, but those questions are of the same kind as the index questions above and admit of the same answers.

Both forms of compute are suitable underliers for derivatives contracts, and they serve different participants. The GPU-hour serves the providers, lenders, and large buyers whose exposure is to hardware rental rates, and the token serves the businesses and users whose exposure is to the price of model output. The Commission should encourage the introduction of derivative products referencing both classes of compute rather than treat one as the sole template for the market.

III. Compute Exhibits the Economic Conditions Associated with Successful Futures Markets

A commodity can satisfy the statutory definition and still fail to support a futures market. The economic literature on why futures contracts succeed or fail helps identify which commodities can, and it informs contract design alongside the statutory listing requirements. Compute meets most of those conditions today, and the ones it does not yet meet can be addressed through contract and benchmark design.

A. Analytical framework

The economic literature on futures markets has long asked why some contracts attract durable trading interest while others fail shortly after listing. The most influential account remains Dennis Carlton’s 1984 study, which the RFC cites, surveying the history of U.S. futures markets to identify the attributes a commodity must have for a contract on it to succeed. Carlton identified five: (i) a commodity must exhibit uncertainty about price, so that there are economic incentives to use hedging instruments; (ii) its prices across slightly different grades and locations must be correlated, so that a standardized contract hedges the physical product with limited basis risk; (iii) it must have a large number of potential participants in a vertically disintegrated industry, so that there are consumers of hedging instruments;

⁴³ RFC, 91 Fed. Reg. at 54,263–64 (Questions 1(c), 2(b), 2(c), and 2(g)).

⁴⁴ *Id.* at 54,262.

(iv) its transactions must be large in value, so that there is an incentive to draw in speculators to assume price risk; and (v) market forces rather than government programs must determine its price.⁴⁵

Subsequent scholarship has largely confirmed Carlton's criteria while condensing them into fewer conditions. Hilary Till's 2014 survey of the later literature distills it into three conditions: (i) a commercial hedging need, (ii) a pool of speculators willing to take the other side of commercial hedgers, and (iii) public policy that is not adverse, which in turn requires a convincing economic rationale and a contract seen as serving the national interest. The Supreme Court's account in *Curran*, which the RFC quotes, adds a legal gloss on the same economics, observing that a futures market requires both a fungible commodity and a standardized contract.⁴⁶ Taken together, this literature provides the criteria against which to measure the nascent compute market.

B. *Factors present*

The compute market already exhibits the conditions that the futures-market literature treats as most indicative of supporting a derivatives market for a commodity:

- **Price uncertainty.** Rental rates have moved sharply in both directions over the market's short history, and participants on both sides of the market have reason to hedge.
- **Large value of transactions.** The RFC cites an estimate of an annual compute service flow of \$430 billion to \$1.3 trillion.⁴⁷
- **Two-sided commercial hedging need.** Providers seek protection against falling rental rates, purchasers seek protection against rising compute costs, and neither side can offset its exposure without a counterparty from the other.
- **A pool of speculators.** Existing event-contract markets on compute prices and the planned listing of compute futures on Architect, ICE Futures U.S., NYMEX, and other venues are potentially indicative of speculator demand.
- **Public policy that is not adverse.** The AI Action Plan directs the federal government to accelerate a financial market for compute. The public debate over the pace of the data-center buildout concerns the level of investment rather than the availability of risk-transfer instruments, which are valuable to participants on either side of that debate.
- **A freely determined price.** No government program sets compute prices, although the concentration of private market participants, discussed below, bears on benchmark design.

On the dimensions the literature treats as most important, then, compute resembles the commodities on which futures markets have succeeded.

⁴⁵ Dennis W. Carlton, *Futures Markets: Their Purpose, Their History, Their Growth, Their Successes and Failures*, 4 J. Futures Mkts. 237 (1984).

⁴⁶ Hilary Till, *Why Some Futures Contracts Succeed and Others Fail: A Survey of Relevant Research* 3, 17 (EDHEC-Risk Inst., Oct. 2014), https://www.premiacap.com/publications/EDHEC_Working_Paper_Why_Some_Futures_Contracts_Succeed_and_Others_Fail.pdf; *Merrill Lynch, Pierce, Fenner & Smith, Inc. v. Curran*, 456 U.S. 353, 357-58 (1982) (quoted in RFC, 91 Fed. Reg. at 54,262 n.44).

⁴⁷ RFC, 91 Fed. Reg. at 54,260 n.5.

C. Unresolved factors

The futures-market literature and the RFC also identify five features of the compute market that remain unsettled. Each of them, however, bears on how a compute contract should be designed, rather than whether compute derivatives should be permitted to list.

- **Standardization and fungibility.** Compute is heterogeneous across chip, configuration, location, and terms, so a contract must define the grade it references with care.
- **Cash-market transparency and the reference price.** Most volume transacts in private bilateral agreements, and some published series rest on provider-administered posted rates rather than completed transactions.⁴⁸
- **Concentration and pricing power.** The parties best positioned to influence a reference price also supply the capacity it measures.⁴⁹
- **Non-storability and deliverable supply.** The Commission will need to adapt its Core Principle 5 analysis for a commodity with no inventory, as it has for electricity.⁵⁰
- **Vintage obsolescence.** A reference chip may be superseded during a contract's life.

None of these five open questions is unique to compute, and the Commission has resolved each of them before in other markets. Electricity contracts required the Commission to adapt its deliverable-supply analysis to a commodity that cannot be stored. The Brent benchmark survived the depletion of the field it originally referenced because the industry widened its deliverable basket. And cash-settled contracts in many markets settle to indices drawn from bilateral transactions that no regulator observes in full. In each instance, contract design and surveillance, rather than a more transparent or standardized cash market, supplied the answer. The conclusions below recommend the same approach for compute.

D. Conclusions and recommendations

Based on the above analysis, HPC offers the following conclusions and recommendations:

1. **Compute already exhibits the core conditions predictive of successful futures markets,** namely price uncertainty, transaction scale, a two-sided hedging need, and supportive policy. The open questions are matters of contract and benchmark design rather than disqualifying features of the commodity.
2. **Listed markets can themselves improve standardization, price discovery, and disclosure,** as RFC Question 1(f) recognizes in asking how listing might affect provider incentives. The Commission therefore should not condition listing on a level of cash-market transparency that the market can develop through trading, provided the proposed contract has a reliable settlement basis and satisfies applicable listing requirements at launch.⁵¹

⁴⁸ RFC, 91 Fed. Reg. at 54,263–64 (Questions 2(a), 2(b), 2(d), and 2(g)).

⁴⁹ RFC, 91 Fed. Reg. at 54,263 (Question 1(f)).

⁵⁰ RFC, 91 Fed. Reg. at 54,263–64 (Questions 1(g) and 2(h)).

⁵¹ CEA § 5(d)(3)–(5), 7 U.S.C. § 7(d)(3)–(5); 17 C.F.R. §§ 38.200–.201, 38.250–.258, 38.300–.301; 17 C.F.R. pt. 38, app. C(a)(2), (c)(2)–(3).

3. **The Commission should permit contracts to settle to competing reference indices** that meet applicable standards and allow market participants to select the benchmarks that best serve their exposures, rather than designating a single administered series.
4. **Manipulation and concentration concerns are best addressed through heightened surveillance**, information-sharing arrangements, and appropriate position limits or accountability during the market's early phase, rather than through a delay of a product for which there is need and demand.
5. **The Commission should articulate what a reliable compute settlement index looks like and let DCMs list contracts that settle to indices meeting that standard**, rather than condition listing on cash-market maturity. Douro Labs, the core contributor to the Pyth Network, has proposed that standard in its submission (contributor breadth and independence, outlier-resistant aggregation, a published dispersion signal, time-weighted settlement, and attribution of every input), together with two related requests: that a DCM be permitted to satisfy Core Principle 4 through an information-sharing arrangement with the index administrator rather than separate arrangements with every data source, provided the DCM and Commission retain access to the underlying information needed for surveillance, and that index robustness be measured by disclosed metrics rather than fixed transaction-count thresholds. HPC supports those positions and agrees in particular that an aggregated, multi-contributor price feed of the kind Pyth publishes should be an acceptable means of determining a settlement price where it meets that standard.⁵²
6. **Onchain derivatives markets can make orders, trades, funding payments, liquidations, and settlements directly observable**, improving the DCM's and the Commission's ability to reconstruct activity and detect abuse. These records complement customer identification and access to relevant offchain and cash-market information.⁵³
7. **The Commission should permit perpetual compute contracts to proceed through its existing case-by-case review under Regulation 40.3 and the Commission's 2026 policy statement on perpetual contracts**,⁵⁴ while allowing venues to compute funding from a continuously updated, multi-contributor index over a time-weighted window; the existing framework provides a route for reviewing these products.⁵⁵

IV. Dated Futures and Perpetual Contracts Each Serve the Hedging Needs of Compute Market Participants

Dated futures and perpetual contracts serve different hedging needs, and the compute market has participants that will benefit from both. The RFC asks whether perpetual compute futures would offer advantages over fixed-date futures or meet risk-management needs that existing products cannot, and whether they would pose unique risks warranting additional safeguards.⁵⁶ The Commission should accordingly facilitate the development of both markets.

⁵² Douro Labs LLC, Comment Letter on the Request for Comment on the Listing of Compute Derivatives Contracts (Sept. 29, 2026), Docket No. CFTC-2026-1850-0023, <https://www.regulations.gov/comment/CFTC-2026-1850-0023> ("Douro Labs Letter"); *see also* 17 C.F.R. pt. 38, app. C(c)(2)-(3); 17 C.F.R. §§ 38.253-.254; 17 C.F.R. pt. 38, app. C(c)(3)(i), (iv)-(v).

⁵³ RFC, 91 Fed. Reg. at 54,263 (Questions 1(c) and 2(d)).

⁵⁴ Policy Statement Concerning the Listing of Perpetual Contracts, 91 Fed. Reg. 33,160, 33,161-62 (June 3, 2026); 17 C.F.R. § 40.3.

⁵⁵ RFC, 91 Fed. Reg. at 54,264 (Question 4(a)).

⁵⁶ RFC, 91 Fed. Reg. at 54,264 (Questions 4(a) and 4(b)).

A. Perpetual contracts

A perpetual contract is a futures contract with no expiration date. A dated future settles on a fixed day, and its price converges to the price of the underlying commodity as that day approaches. A perpetual contract stays aligned with its reference price through funding payments exchanged periodically between longs and shorts, which impose a cost on whichever side has pushed the contract price away from the reference and so give traders a continuous incentive to trade the deviation away. A holder of a perpetual contract therefore maintains exposure for as long as it chooses, without rolling from one contract month to the next, and pays or receives funding in the interim.

Perpetual contracts are now part of the regulated U.S. market. Perpetual-style contracts referencing digital assets have traded on a registered U.S. exchange since July 2025,⁵⁷ and in May 2026 the Commission approved the first true perpetual, KalshiEX's BTCPERP, as a futures contract and issued a policy statement describing how it will review perpetual contracts under its existing rules.⁵⁸ Trading in perpetual contracts on U.S. venues has grown rapidly since, in digital assets and in precious metals. Perpetual contracts on compute would therefore enter a regulatory framework that already accommodates the product.

Many compute market participants carry price exposure with no fixed end date, and a perpetual contract matches that exposure more closely than a dated contract does. For example, a data center earns rental revenue hour by hour, and an inference platform incurs compute costs every day on each model call. Hedging an exposure of that shape with dated futures requires a strip of contracts and a series of rolls, and each roll carries transaction costs and timing risk.⁵⁹ A perpetual allows a participant to hold one position for as long as the exposure lasts, subject to funding costs and basis risk, without scheduled rolls. The exposure is continuous in time as well. Compute demand runs around the clock, and software agents may increasingly procure and pay for it without a human in the loop.⁶⁰ A market that trades continuously and never expires, as onchain perpetual markets can, matches that profile in a way that a session-based, dated contract does not. HPC described the advantages of continuous, margined markets for continuous exposures in its August 2026 submission on energy perpetuals,⁶¹ and many of those advantages apply similarly to compute derivatives, including:

⁵⁷ Coinbase, *Coming July 21: US Perpetual-Style Futures* (June 26, 2025), <https://www.coinbase.com/blog/coming-july-21-us-perpetual-style-futures>; MK Manoylov, *Coinbase Derivatives' Two US Perpetual-Style Futures Contracts to Begin Trading on July 21*, The Block (June 26, 2025), <https://www.theblock.co/post/359852/coinbase-derivatives-two-us-perpetual-style-futures-contracts-to-begin-trading-on-july-21>.

⁵⁸ Press Release, CFTC, CFTC Approves BTCPERP Contract Submitted by KalshiEX, LLC (May 29, 2026), <https://www.cftc.gov/PressRoom/PressReleases/9240-26>; Policy Statement Concerning the Listing of Perpetual Contracts, *supra* note 54.

⁵⁹ See Hyperliquid Policy Center, *Perpetual Futures as Complements to Dated Futures* (Aug. 21, 2026), https://hyperliquidpolicy.org/whitepaperfuturesascomplements/hpc_research_-_perpetual_futures_as_complements_to_dated_futures_final_.pdf ("HPC Study") (explaining that perpetual contracts allow hedgers to avoid the transaction costs and timing risk of periodically rolling dated futures positions).

⁶⁰ See BlackRock Report, *supra* note 8, at 5, 9.

⁶¹ Hyperliquid Policy Center & trade[XYZ], Comment Letter on the Request for Comment on the Extension of Standard Futures Contracts to 24/7 Trading and on Perpetual Contracts Referencing Physically Delivered or Storable Energy Commodities 4-8, 16-21 (Aug. 26, 2026), https://hyperliquidpolicy.org/hpctradecftcrfcresponse/2026-08-26_hpc_x_trade-xyz- cftc_rfc_response.pdf ("HPC Energy Letter").

- **Liquidity concentration.** Because a perpetual contract does not expire, trading interest gathers in a single instrument rather than fragmenting across a ladder of contract months, which matters most in a new market where liquidity will be scarce.
- **No roll friction.** A user of dated futures must periodically exit one contract and enter the next to maintain exposure, paying venue fees on both legs and bearing timing risk inside the roll window. A perpetual position never incurs that cost.⁶²
- **Continuous hedging.** A single perpetual position can be held, adjusted, or closed for as long as the exposure lasts, so the hedge follows the exposure rather than the contract calendar.
- **Continuous access.** The events that move compute prices, such as model releases, outages, and export-control announcements, keep no trading calendar, and a continuously traded market lets a hedger act when the information arrives rather than at the next session.⁶³

B. Participants and use cases

The participants for whom a perpetual is the better instrument share a common profile. Their exposure continues indefinitely, or it outlasts the tenors in which a dated market will have liquidity. An infrastructure operator hedging rental revenue against a debt schedule, an inference platform managing a continuing compute budget, a financed operator or its lender managing the rental-price component of a multi-year exposure, and a broker that sells specific chip configurations while hedging against a benchmark contract each fit that description.

Liquidity in a new compute contract is likely to form in the nearest months, because participants cannot forecast which chips will remain commercially relevant at distant tenors. Participants with multi-year exposures (an operator financing hardware over its useful life, or a lender whose collateral value and borrower cash flows depend on rental rates over the loan term) would have to hedge with a rolling stack of short-dated contracts, which introduces both timing risk and expensive roll costs. A perpetual can provide a single instrument for that exposure and concentrate trading interest in one listing rather than across many months and indices, which may also help the nascent compute derivatives market build liquidity.

A perpetual can also mitigate the vintage obsolescence problem identified in Part III. Where a dated contract on a specified chip risks referencing hardware superseded before expiry, a perpetual's reference index can migrate across hardware generations under disclosed, pre-committed transition rules, maintaining the benchmark's relevance through generational turnover. Perpetual markets on Hyperliquid that reference the price of the benchmark front-month crude oil contract provide a related operational example. They roll to the next month over several days by blending the two reference prices in stated proportions, so the transition is gradual, disclosed in advance, and predictable.⁶⁴ A similar approach could be used to transition the referenced hardware type for a perpetual compute contract.

C. Risks specific to perpetual contracts

Perpetual contracts raise a set of concerns that are not specific to compute, including funding costs, convergence without a terminal settlement date, margin and liquidation under continuous trading, retail suitability, and the position-limit treatment of a contract with no spot month. HPC has addressed each of them in prior submissions to the Commission and has explained why, with appropriate disclosure,

⁶² HPC Energy Letter, *supra* note 61, at 5–6.

⁶³ HPC Energy Letter, *supra* note 61, at 14–16.

⁶⁴ trade[XYZ], *Roll Schedules*, <https://docs.trade.xyz/perpetuals/specifications-and-schedules/roll-schedules> (last visited Oct. 19, 2026).

pre-funded margin, automated close-out rules, and position limits, none prevents perpetual contracts from trading in an orderly way on registered venues.⁶⁵

Three risks are specific to compute perpetuals, and each has a corresponding safeguard.⁶⁶

- **Funding-rate volatility around demand shocks.** Model releases and hardware shortages may reprice compute abruptly, and a funding rate computed over a short window could swing with them. Funding-rate caps, or a funding calculation smoothed over a longer time-weighted window, could limit the cost a position can incur in any single interval.
- **Index transition events.** When a reference index migrates from one chip generation to the next, the contract price can move for reasons unrelated to supply and demand. Transition protocols disclosed in advance and phased in over several days, as the crude oil perpetual markets on Hyperliquid do when they roll between reference months, remove the surprise and give hedgers time to adjust.
- **Feedback between a thin cash market and funding.** A reference price built from few transactions can be moved by trading in the derivative itself, and the funding rate would then transmit that move back into the contract. Volatility-calibrated margin, concentration limits, and a multi-contributor index with outlier-resistant aggregation each break that loop.

Continuous trading also requires safeguards for periods when the underlying cash market is inactive, consistent with the staff's 24/7 trading advisory.⁶⁷ As HPC has explained in prior submissions, mechanisms such as trade[XYZ]'s discovery bounds, which hold the mark price within a dynamic range of the last reference price, support stable trading during reference-market closures.⁶⁸

D. Dated futures and the forward curve

Dated futures and perpetual contracts are complementary risk-management instruments, and the first compute contracts proposed for listing are dated.⁶⁹ A dated contract settles on a fixed day, so it hedges an exposure with a defined horizon without funding cost, and a strip of dated contracts across successive months produces a forward curve, the term-structure information that tells participants how the market values compute at future dates. The contracts NYMEX and ICE Futures U.S. have announced take this form, and event-contract markets have already supplied a market-implied forward curve for compute, which suggests that market participants are building this layer first.

Several compute exposures are better hedged with a dated contract than with a perpetual. A training run planned for a particular quarter, an anticipated purchase with an unfixed price and a known delivery date, a capacity buildout with a known schedule, and a budget cycle with a fixed horizon each end on a date, and a hedger of any of them would rather fix a price for that date than carry a perpetual position and its funding payments past it. A lender or investor marking a multi-year financing against the forward curve likewise needs the term structure that only dated contracts provide.

⁶⁵ See HPC Energy Letter, *supra* note 61, at 8-13.

⁶⁶ RFC, 91 Fed. Reg. at 54,264 (Question 4(b)).

⁶⁷ CFTC Staff Letter No. 26-16, *Staff Advisory for Extending Trading and/or Clearing Operations to a 24 Hours-a-Day, 7 Days-a-Week Basis* 4-5 (May 29, 2026), <https://www.cftc.gov/csl/26-16/download>.

⁶⁸ trade[XYZ], *Discovery Bounds*, <https://docs.trade.xyz/perpetuals/mechanics/discovery-bounds> (last visited Oct. 19, 2026).

⁶⁹ See HPC Study, *supra* note 59 (finding that perpetual markets expanded hedging access without measurable harm to the corresponding dated benchmarks).

Dated and perpetual compute contracts therefore serve different participants with different exposures, and neither substitutes for the other. A market that offers both gives a hedger the choice between a defined horizon and a continuous one, and HPC encourages the Commission to foster the development of both product types so that the instruments available to compute market participants match the range of exposures they carry.

V. Customer Protection and Retail Participation in Compute Derivatives Markets

Question 3 of the RFC asks what customer protections compute derivatives warrant, including for retail participants.⁷⁰ Compute is a commodity that retail participants already buy. Unlike buyers of crude oil or natural gas, many compute buyers are individuals and small businesses, among them independent developers, startups, researchers, and the growing population of users who pay metered rates for model access. The AI Action Plan's stated reason for improving the financial market for compute is to give startups and academics access to large-scale capacity they cannot otherwise afford.⁷¹ A retail participant hedging a compute budget, or taking a view on the input that determines the cost of the AI services it uses, has a commercial relationship to the underlier that few retail participants in energy or metals markets can claim. Many small businesses and startups with genuine compute exposure may not qualify as eligible contract participants under the CEA.⁷²

The protections appropriate for compute perpetuals follow from the features of the product rather than from the identity of the underlier, and HPC has proposed similar safeguards in other submissions on energy and pre-IPO perpetuals,⁷³ including:

- **Product-specific risk disclosure.** A perpetual contract carries costs and mechanics that a dated contract does not. Disclosure should cover contract mechanics, funding rates with illustrative dollar-impact examples over different holding periods, leverage and liquidation thresholds, and the pricing methodology and index definition, including how the index will handle hardware transitions.
- **Leverage limits, position limits, and open interest caps.** Limits calibrated to the asset class bound the loss any one participant can suffer and the influence any one participant can exert on a market whose reference price rests on a thin cash market.
- **Pre-announced oracle and settlement rules.** Venues should publish the rules by which the index is computed and the contract settles, and change them only with notice, so that a compute hedger is never surprised by a change in what the contract references.
- **Disclosure of operator discretion.** Any discretion a venue or index administrator retains, including over hardware transitions, should be disclosed in advance.

These protections matter most for the participants the AI Action Plan singles out, the startups and academics that are least able to negotiate bilateral alternatives, and registered venues can implement each of them under the Commission's existing framework.

⁷⁰ RFC, 91 Fed. Reg. at 54,264 (Questions 3(a)–3(e)).

⁷¹ AI Action Plan, *supra* note 4, at 4.

⁷² CEA § 1a(18), 7 U.S.C. § 1a(18), including the entity eligibility criteria in § 1a(18)(A)(v).

⁷³ HPC Energy Letter, *supra* note 61, at 11–14, 27; Hyperliquid Policy Center & trade[XYZ], Comment Letter on IPO Modernization (File No. CLL-16) 14 (Aug. 18, 2026), <https://hyperliquidpolicy.org/ipocommentletter/20260818hpc-x-tradexyz-ipo-modernization-comment.pdf>.

Existing customer-identification, anti-money-laundering, and suspicious-activity-reporting requirements would apply to registered intermediaries in compute derivatives markets.⁷⁴ Venue surveillance should watch for provider-linked trading and for attempts to influence index inputs, funding, or settlement, and disclosures should explain how export restrictions or other geopolitical developments may affect compute prices and access.

VI. Conclusion

Market participants across the compute supply chain carry exposure to compute prices and would benefit from derivatives markets in which to transfer it. Compute satisfies the CEA's definition of a commodity, exhibits the economic conditions associated with successful futures markets, and presents design questions, chiefly the selection of a reliable settlement index, that contract design and surveillance can resolve. Its distinguishing features, above all that it cannot be stored and that it is produced and consumed continuously, favor contract designs built for continuous exposure, including perpetual contracts, alongside dated futures.

HPC accordingly encourages the Commission to permit compute contracts to proceed through its existing review processes, to accept settlement to any index that meets a disclosed standard of reliability, to foster the development of both dated and perpetual products, and to address the remaining concerns through surveillance, disclosure, and position-level safeguards rather than delay. HPC appreciates the Commission's consideration of these comments and would welcome the opportunity to discuss them with the Commission or its staff.

Respectfully submitted,

/s/ Jake Chervinsky
Jake Chervinsky
Chief Executive Officer
Hyperliquid Policy Center

/s/ Brad Bourque
Brad Bourque
Senior Counsel
Hyperliquid Policy Center

⁷⁴ 31 C.F.R. §§ 1026.210, 1026.220, 1026.320; 17 C.F.R. § 42.2.