ocpi_methodology_public_web

Ornn Compute Price Indices (OCPI): Methodology

Ornn ResearchJuly 2026

> *The Ornn Compute Price Indices (OCPI) measure the real-time market cost of GPU compute using on-demand trades executed across a global network of verified compute providers. For a given GPU type, the index is calculated as the volume-weighted winsorized mean of executed transaction prices over a rolling one-hour window. OCPI is designed as a transparent, settlement-grade measure of the value of compute, expressed in USD per GPU-hour.*

---

1. Introduction

The cost of high-performance compute has become a critical input to the economics of artificial intelligence. GPU rental markets link hardware supply to AI deployment capacity, yet the market has historically lacked a standardized measure of compute pricing.

Ornn's index work began with the Ornn US H100 Index, a representation of rental offers for NVIDIA H100 GPUs across the United States. The Ornn Compute Price Indices (OCPI) generalize that work into a methodology built on live trade data, the prices at which buyers and sellers explicitly converge on the value of compute. Rather than measuring quoted or listed prices, OCPI tracks the traded price of compute where real-time supply and demand intersect.

The indices are constructed to serve as settlement-grade references, enabling market participants to hedge and gain exposure to GPU prices. In doing so, they represent a necessary step in the commodification of compute.

2. Universe Creation

Ornn sources anonymized pricing data from a global network of compute providers operating with dynamic pricing. Ornn records the price of every GPU

under active on-demand rental across its contributor sources, denominated in USD-equivalent per GPU-hour. Each observation captures the price actually being paid for compute at that moment.

OCPI measures on-demand rentals only. Reserved-capacity, long-term, and forward contracts are outside the scope of the indices: such agreements are bespoke and bilaterally negotiated, with custom durations, prepayments, and discount structures, and do not produce comparable arm's-length spot observations. On-demand transactions constitute the most disaggregated, frequently repriced segment of the compute market and therefore provide the cleanest basis for a spot benchmark.

Each observation is represented as an ordered tuple (p_i, n_i, r_i, g_i), denoting the price, the number of GPUs, the region, and the GPU type, respectively (regionality is broken down into North America, South America, Europe, Asia, Africa, and Oceania). The universe U_t is defined as the totality of observed active rentals at time t.

Indices are localized per GPU type, as the value of compute differs substantially across hardware specifications. This is consistent with established practice in commodities benchmarking: it captures the heterogeneity of the underlying market while maintaining a manipulation-resistant measure.

New GPU-specific indices are introduced as market conditions warrant. The criteria for launching a new index include sustained transaction volume across the contributor network, sufficient hardware availability to support a representative sample, and demonstrated demand from market participants. Decisions to add, modify, or retire an index are made through Ornn's index governance process and announced in advance.

Provider identities are not disclosed publicly. This is intentional: disclosure would create incentives for contributors to modify their behavior in anticipation of inclusion reviews, undermining the integrity of the benchmark. This approach is consistent with established financial benchmarking practice.

Composite indices spanning multiple GPU types may be introduced, each with its own specification published at launch. The single-GPU methodology described here governs the constituent inputs of any composite index unless the composite's specification states otherwise.

2.1 Data Fidelity and Filtering

Within the universe U_t, only trades satisfying all of the following criteria are eligible for index calculation:

1. **Verified Counterparties.** Compute must have been successfully transferred between two different counterparties, preventing wash trades from entering the calculation universe. GPU specifications, including CUDA version and VRAM, are verified at onboarding and verified again at the point of clearing.
2. **Periodic Re-verification.** All contributing providers are re-verified at periodic intervals against minimum verification thresholds, per the provider verification standards described in Section 7.2. Providers that no longer satisfy eligibility requirements are suspended from contribution until compliance is restored and confirmed. Re-verification results are not disclosed at the provider level, consistent with source confidentiality.
3. **Trade Finality.** Only realized rental prices, meaning prices actually being paid under active rentals, are included. Indicative or offered prices are excluded.
4. **Trade-Size Treatment.** No minimum trade-size threshold is applied; rentals of all sizes are included. The influence of anomalous prints is limited through winsorization, described in Section 3.
5. **Provider Reliability.** Providers exhibiting incomplete reporting or excessive latency may be temporarily excluded until data quality is restored. Provider-level reliability metrics are monitored continuously.

These criteria are cumulatively designed to guarantee the integrity of the index.

2.2 Provider Eligibility and Enforcement

Hardware must satisfy all verification criteria at onboarding; machines that do not qualify are not admitted to the contributor network. At re-verification, machines that no longer satisfy the criteria are suspended pending restoration of compliance.

Suspension is not retroactive: previously published index values are not restated solely because hardware degradation is discovered at re-verification, unless independent evidence indicates that the degradation constituted intentional misrepresentation.

Intentional submission of false, misleading, or manipulated data is treated as a separate and more serious category of non-compliance. Where reasonable grounds exist, consequences include immediate suspension pending investigation, review of affected historical index values covering the period of the manipulation, restatement where warranted and practicable, and permanent exclusion from the contributor network upon confirmation. Ornn reserves the right to refer confirmed cases to regulatory and law enforcement authorities.

3. Index Calculation

3.1 Overview

Each index is computed as the volume-weighted winsorized mean of executed transaction prices for a given GPU type over a rolling one-hour window. Winsorization retains all observed trades and adjusts only those prices falling outside a defined percentile range, replacing outliers with the corresponding bound values. Because trade volume is retained at the boundaries, the estimator is resistant to small-sample distortion while preserving the full informational content of every executed transaction.

OCPI applies no smoothing, averaging across windows, or other filtering beyond the winsorization described in this section; each published value reflects only the executed transactions in its own calculation window.

3.2 Trade Set and Notation

For a GPU type g, let M_g denote the set of executed transactions in the current one-hour window. Each transaction i is described by the pair (p_i, n_i), where p_i is the executed price in USD per GPU-hour and n_i is the number of GPU units traded:

$$M_g = \{(p_i, n_i) : i = 1, \ldots, N\}$$

The aggregate traded volume is

$$V_g = \sum_{i=1}^{N} n_i$$

3.3 Volume-Weighted Empirical Distribution

Percentile bounds are derived from the volume-weighted empirical cumulative distribution function of transaction prices, rather than the unweighted trade-count distribution. Each transaction contributes n_i observations, one per GPU unit at price p_i, ensuring that a single large-volume trade has proportionally greater influence on the bounds than a small-volume trade at the same price:

$$\hat{F}_{M_g}(x) = \frac{1}{V_g} \sum_{i=1}^{N} n_i \, \mathbf{1}[p_i \leq x]$$

where $\mathbf{1}[\cdot]$ is the indicator function.

3.4 Winsorization Bounds

The lower and upper winsorization bounds are symmetric percentile bounds of the volume-weighted empirical distribution, taken at a lower percentile level α and its upper complement $1 - \alpha$, and defined via the generalized inverse:

$$q_\alpha = \inf\{x : \hat{F}_{M_g}(x) \geq \alpha\}, \qquad q_{1-\alpha} = \inf\{x : \hat{F}_{M_g}(x) \geq 1 - \alpha\}$$

Because the bounds are drawn from the volume-weighted distribution, a participant seeking to manipulate the index by shifting the bounds must transact genuine volume at off-market prices; a single nominal print at an extreme price level is insufficient, as its weight is proportional to the GPU units actually traded.

The percentile level α is set by Ornn and is not published; it is available to licensed counterparties, listing venues, and reviewers under the confidentiality arrangements described in Section 7.2. Ornn reserves the right to review and adjust the winsorization level and bounds as market conditions warrant, subject to the methodology change procedures described in Section 5.6.

3.5 Winsorization Rule

The eligible price range is the closed interval $[q_\alpha, q_{1-\alpha}]$. Each transaction's winsorized price is defined as the lower bound if the executed price falls below it, the executed price itself if within the range, and the upper bound if above it:

$$\tilde{p}_i = \begin{cases} q_\alpha & p_i < q_\alpha \\ p_i & q_\alpha \le p_i \le q_{1-\alpha} \\ q_{1-\alpha} & p_i > q_{1-\alpha} \end{cases}$$

A price equal to either bound lies within the eligible range and is not adjusted. Every transaction in M_g contributes to the final index value; no transaction is excluded.

3.6 Index Computation

The index value for GPU type g is the volume-weighted mean of winsorized prices:

$$I(g) = \frac{\sum_{i=1}^{N} n_i \tilde{p}_i}{\sum_{i=1}^{N} n_i} = \frac{1}{V_g} \sum_{i=1}^{N} n_i \tilde{p}_i$$

The index is expressed in USD per GPU-hour and serves as the settlement reference for derivatives built on OCPI.

3.7 Publication

Index values are published as promptly as practicable following the close of each one-hour window, with the reference timestamp corresponding to the window close. Values are calculated and delivered to the nearest tenth of a cent ($0.001). Current and historical index values are available via the Ornn data platform at data.ornn.com and through authorized data partners, with publication tiers described there.

4. Index Properties

4.1 Robustness

Beyond the pre-processing controls described in Section 2.1, OCPI benchmarks are calculated in a manner that expresses liquidity sensitivity, further enhancing robustness against distortion.

The volume-weighted winsorized mean provides two complementary layers of protection. First, volume weighting ensures that a party seeking to shift the index

must command sufficient traded volume to alter the weighted mean; this raises the cost of manipulation materially, as influence over the published value is proportional to genuine economic participation in the market. Second, winsorization caps the influence of extreme prints without discarding trades: every transaction contributes to the calculation, but prices falling outside the eligible range are adjusted to the nearest bound rather than removed.

The winsorizing bounds are themselves robust statistics, drawn from the volume-weighted empirical distribution of the calculation window. As a result, a coordinated series of off-market prints has limited effect on the published value unless the manipulator commands sufficient traded volume to displace the corresponding fraction of total window volume. This architecture makes concentrated manipulation materially costly.

4.2 Validation and Restatement

Prior to publication, each index value undergoes internal validation designed to identify anomalous or potentially erroneous values. Values that deviate materially from recent historical distributions are flagged for review before release.

If a published value is subsequently determined to contain a material error, Ornn publishes a corrected value and notifies affected stakeholders through appropriate channels. As a general matter, Ornn does not restate historical index values. An exception applies where an error is attributable to data falsification by a contributing provider, or to a calculation error that Ornn determines to be material to benchmark integrity.

Users relying on OCPI values for settlement purposes are advised to maintain independent records of published values.

4.3 Market Disruption Events

A Market Disruption Event (“MDE”) occurs when Ornn determines that the transaction data available for a calculation window is insufficient to support a reliable index determination.

Upon declaration of an MDE, the following fallback hierarchy applies:

1. **Carry-forward.** The most recently published index value is carried forward as the index value for the affected window. Carry-forward may be applied for a

limited number of consecutive calculation windows.

2. **Suspension.** If carry-forward would be inconsistent with the integrity of the index, or if the carry-forward limit is reached without resolution of the disruption, publication is suspended for the affected windows until the disruption is resolved.

If an MDE is active at the scheduled expiry of a listed contract referencing an affected index, the final settlement price is determined from the most recent published index value that was computed from Tier 1 transaction data and was not itself a carry-forward value. Where no sufficiently recent Tier 1 value is available, Ornn will convene the Oversight Committee (Section 5.1) promptly following the scheduled expiry to determine a final settlement price using the best available transaction data, supported by a documented rationale. The detailed settlement-determination procedures, including the applicable recency and timing parameters, are provided to licensed counterparties and listing venues.

MDEs are communicated to relevant stakeholders, including any exchange listing a derivative referencing the affected index, as promptly as practicable.

5. Governance and Administration

5.1 Independent Oversight

Ornn is establishing an oversight function, constituted as a standing committee (the “Oversight Committee”), that will be responsible for benchmark governance, the integrity of the methodology, and alignment with applicable benchmark standards. The Oversight Committee will operate independently of the day-to-day calculation and data collection functions.

The responsibilities of the Oversight Committee will include: periodic review of the methodology; assessment of whether the methodology continues to reflect the underlying market; review of proposed changes to calculation procedures or eligibility criteria; review of complaints; and investigation of suspected manipulation or misconduct.

The Oversight Committee will comprise no fewer than three and no more than five members. A majority of members must at all times satisfy the independence standard: a member is independent if the member is not, and has not been within

the preceding 24 months, an employee, officer, director, or equity holder of Ornn AI Inc. or any affiliate, and has had no material commercial relationship with Ornn or an affiliate within that period. Members who become aware of circumstances affecting their independence must disclose those circumstances promptly and recuse themselves from affected determinations.

A quorum will consist of a majority of seated members, including at least one independent member. Decisions will be taken by majority vote. Material methodology changes, as defined in Section 5.6, will require an affirmative committee vote before taking effect.

The composition of the Oversight Committee, including the name and independence status of each member, will be published on Ornn's public documentation site upon its constitution.

5.2 Conflicts of Interest

Ornn administers OCPI and also operates commercial services related to GPU compute capacity. This operational overlap creates a potential conflict between the accurate and credible determination of the benchmark and Ornn's commercial interests as a market participant. Ornn acknowledges this conflict and maintains documented procedures to manage it.

The benchmark determination function is operationally separated from Ornn's commercial activities. Personnel responsible for index calculation and methodology decisions are subject to information barriers restricting the flow of non-public benchmark information to commercial functions, and it is Ornn's policy that such personnel are not compensated on the basis of the commercial trading or marketplace outcomes of Ornn or its affiliates. Once constituted, the Oversight Committee will review identified conflicts periodically and maintain a conflicts register, reviewed and updated as the Oversight Committee determines.

Transactions executed on marketplaces or platforms operated by Ornn or its affiliates receive no preferential treatment in the index. Such transactions are eligible for inclusion only where they satisfy every verification, counterparty, and filtering requirement of Section 2.1, applied identically to all other contributor data, and they receive no distinct treatment in the calculation. The Oversight Committee will monitor the share of index input attributable to Ornn-affiliated venues as part of its periodic conflicts review.

5.3 Contributor Code of Conduct

All entities participating in data contribution to OCPI are subject to the following requirements. Participation in data contribution constitutes acceptance of this code.

Submitted data must reflect bona fide, arm's-length transactions for GPU compute capacity executed on commercially reasonable terms. Only transaction data is eligible for submission: indicative prices, internal transfer prices, notional values, and non-executed quotes are ineligible.

The following conduct is prohibited and constitutes grounds for disciplinary action:

- Submitting fabricated, altered, or coordinated data intended to influence index levels;
- Systematically omitting eligible transactions with the intent to distort the sample;
- Structuring the timing of transactions or submissions to cause a particular index outcome;
- Sharing submission data or methodology details with any party who may use the information to trade ahead of a published index value.

Any party with reason to believe that a contributor has violated this code may report the matter to Ornn's compliance function. Reports are treated confidentially to the extent permitted by law, and retaliation against good-faith reports is prohibited.

Contributors are required to retain all records relating to their data submissions for a minimum of five years.

5.4 Complaints Procedure

Any stakeholder, including index users, contributors, and other interested parties, may submit a complaint regarding an index determination, the methodology, or governance arrangements to Ornn's designated compliance contact. Receipt of a complaint is acknowledged, and a substantive response is provided within a reasonable timeframe.

All complaints will be referred to the Oversight Committee, which will document the complaint, its review, and its outcome. Summary information regarding complaints is published periodically. Individual complaints are not disclosed without the submitter's consent, except as required by law.

5.5 Audit and Record Retention

Ornn retains records of all data inputs received from contributors, all intermediate calculations, and all published values. Records are stored in tamper-evident form that preserves their integrity as at the time of capture and are retained for a minimum period consistent with applicable regulatory standards.

Ornn conducts periodic internal reviews of its calculation processes, data collection procedures, and governance arrangements, with material findings to be escalated to the Oversight Committee. External audit or assurance reviews may be conducted at Ornn's discretion or as required by applicable regulation.

5.6 Methodology Changes and Consultation

The methodology may be updated to reflect changes in market structure, data availability, technological development, or applicable regulatory requirements. Methodology updates are issued as versioned releases with change logs identifying the nature and effective date of each update.

Changes are classified as material or non-material. A change is material where it may meaningfully affect index values, contributor eligibility, or the analytical basis of the index. Material changes are preceded by a consultation period during which stakeholders may submit written comments, and take effect only after reasonable advance notice.

Non-material changes, being corrections and clarifications that do not alter the substantive methodology, may be implemented without prior consultation and are disclosed as soon as practicable. Emergency changes necessary to preserve index integrity or comply with law may be implemented without prior consultation, with notice provided as promptly as circumstances permit.

5.7 Cessation Policy

Discontinuation of an index may have material consequences for users relying on OCPI for settlement, risk management, or commercial purposes. If any index must

be discontinued, Ornn provides advance notice to stakeholders, with a notice period as long as reasonably practicable. In the event of a planned cessation, Ornn publishes a transition plan, where feasible identifying alternative benchmarks or methodologically comparable references.

Users incorporating OCPI into derivative contracts, structured instruments, or other financial arrangements are encouraged to ensure that those arrangements contain adequate fallback provisions addressing cessation or material modification of the index. Ornn is not obligated to maintain any index in perpetuity.

5.8 IOSCO Alignment

The OCPI methodology is designed with reference to the IOSCO Principles for Financial Benchmarks. As a benchmark for a novel and rapidly evolving asset class, not all IOSCO principles apply directly; Ornn applies them in a manner appropriate to the characteristics of the compute market. Ornn intends to publish a periodic statement describing its alignment with the IOSCO Principles, including areas where a proportionate approach is applied.

Governance structures, oversight mechanisms, and operational controls are reviewed on an ongoing basis and updated as the GPU compute market and applicable standards evolve.

6. Conclusion

By advancing measurement from offered prices to completed trades, OCPI provides a transparent, executable-price benchmark for the on-demand GPU compute market, tracking where demand actually materializes into transactions rather than where sellers set prices. Built on verified transaction data (Sections 2.1 and 7.2), governed by independent oversight (Section 5), and calculated through a manipulation-resistant architecture (Sections 3 and 4), OCPI supports cash-settled derivatives and structured instruments linked to the cost of compute, and advances the commodification of compute through real-time benchmarks of its value.

7. Appendix

7.1 Notation

Symbol	Definition
Trade	A successful transfer of GPU compute between verified counterparties
p_i	Price (USD per GPU-hour) of trade i
n_i	Number of GPUs traded in transaction i
r_i	Region of execution for trade i
g_i	GPU type for trade i
N	Number of executed transactions in the one-hour window
M_g	Set of executed trades for GPU type g in the one-hour calculation window
V_g	Total GPU volume in the window, $\sum_{i=1}^{N} n_i$
$\hat{F}_{M_g}(x)$	Volume-weighted empirical distribution function
q_α	The α-quantile of the volume-weighted distribution
$\tilde{p}_i$	Winsorized price for trade i, equal to p_i clamped to $[q_\alpha, q_{1-\alpha}]$ per Section 3.5
$I(g)$	Index value for GPU type g
U_t	Full universe of trades at time t

7.2 Provider Verification Standards

Eligibility for a given OCPI index is anchored first in hardware conformance: a machine's GPU must be verified as the underlying GPU type of that index. GPU model and specifications, including CUDA version and VRAM, are verified at onboarding and again at the point of clearing (Section 2.1). A machine cannot qualify for an index on the basis of host-level minimums alone; conformance to the index's underlying hardware type is the primary control against non-conforming hardware entering a print.

In addition to hardware anchoring, contributing machines must satisfy minimum verification thresholds across the following dimensions:

- Demonstrated historical reliability;
- Sustained network throughput (download and upload);
- Driver and CUDA currency;

- GPU memory;
- System memory, attached storage, and CPU allocation per GPU;
- Interconnect/PCIe bandwidth

Complete threshold values are available for review by regulators, exchange listing and index-eligibility committees, and independent auditors or assurance reviewers, under confidentiality arrangements. Requests are directed to Ornn's compliance contact.

All contributing providers are re-verified against these standards at periodic intervals; machines failing re-verification are suspended pending remediation. Ornn monitors these metrics continuously and reviews the thresholds as the compute market matures and hardware standards evolve. Changes to verification standards follow the methodology change procedures of Section 5.6.

7.3 Data and Publication

All trade timestamps are recorded via millisecond Unix timestamp and aligned to the close of each one-hour calculation window. Index values are published as promptly as practicable following window close on a continuous basis, calculated and delivered to the nearest tenth of a cent ($0.001). Data aggregation and filtering occur on Ornn's secure analytics infrastructure, with results verified by independent servers. Current and historical index values, including the inception date of each index, are available via the Ornn data platform at data.ornn.com and through authorized data partners.

7.4 Disclaimer

The OCPI suite is designed to provide a transparent, data-driven measure of the value of compute. Ornn maintains strict data validation and review processes, but all published index values are provided "as is." Users should exercise independent judgment when referencing OCPI in contracts or analysis. Ornn accepts no responsibility for any direct or consequential losses resulting from its use.

Questions regarding the methodology, data access, or verification standards: contact Ornn Research at data@ornn.com.

Methodology version 1.0, July 2026. Versioned releases and change logs are available at ornn.com.