METHODOLOGY

Compute Baskets

Construction, NAV & Live Pricing

Combinatoric GPU-hour cost and fleet baskets built from cross-provider compute rental prices — component construction, weighting, historical NAV, and real-time updates.

Produced by Itô Markets
Compute Markets · Quantitative Research · Version 1.0 · Q3 2026

Contents

1 Executive Summary

Itô's compute baskets are transparent, rules-based price indices over the rental market for AI compute. Each basket is a fixed weight vector over *components* — canonical GPU-hour price series defined by GPU SKU, form factor and pricing term (e.g. `H100:SXM:on_demand`) — and its net asset value (NAV) is the weighted cross-provider dollar cost per GPU-hour of that mix. The universe mirrors the combinations institutional counterparties actually request (training fleets, inference fleets, on-demand/spot hybrids, vendor-diversified mixes), packaged as twenty long-only baskets.

Historical NAV is computed daily from the OpenComputePrices archive (2026-03-01 onward, 77 providers, 20 canonical GPU SKUs). Live NAV ticks forward in real time from Itô's one-minute compute quote feed through the same edge NAV engine that prices every other Itô basket, ratio-anchored to the historical series so the published level is continuous. The published level is the basket's actual dollar cost per GPU-hour, not a rebased index.

2 Data Sources

2.1 Historical: OpenComputePrices

The historical layer is built from the OpenComputePrices open dataset — daily snapshots of published rental prices across 77 cloud and marketplace GPU providers (hyperscalers, specialist clouds, and spot marketplaces). Raw observations (roughly 35 million rows over the window) are reduced to daily per-{provider, GPU, form factor, pricing term} medians, retained from 2026-03-01 through the present.

2.2 Live: Itô compute quote feed

The live layer is Itô's own compute price collection: one-minute quotes per provider/SKU/term series (e.g. `lambda:h100:vm:on_demand`), relayed through the Itô edge and persisted to the hot store. These series are the basket legs that tick live NAV between daily historical closes.

3 Component Construction

3.1 GPU normalization

Provider SKU strings are normalized to canonical GPU identifiers (H100, H200, B200, A100 80GB/40GB, MI300X, MI325X, L40S, L4, A10, RTX 4090, RTX A6000, V100, T4, ...). A100 memory variants are split using per-GPU memory (80 GB vs 40 GB).

3.2 Form factor

Instances are classified SXM (NVLink/NVSwitch-class), PCIE, or UNKNOWN from variant and interconnect metadata. The SXM class is the observable carrier of NVLink-fabric value in rental pricing.

3.3 Pricing terms

On-demand and spot/preemptible list prices are kept as separate component series; reserved/commitment prices are excluded from basket NAV.

3.4 Provider aggregation

For each day, a component's value is the *median across providers* of the per-provider daily median price. Median-of-medians is robust to outlier listings and to any single provider's repricing.

3.5 Missing data and forward-fill

A provider series missing on a given day is forward-filled for at most three days; beyond that it drops out of the cross-provider median. A basket publishes historical NAV only on dates where *every* component resolves.

4 Basket Construction

4.1 Weights

Each basket is a fixed vector of non-negative component weights summing to one, chosen to mirror observed procurement mixes (e.g. Training Core = 80% H100 SXM + 20% H200, on-demand). Weights are static; there is no rebalancing or optimization.

4.2 NAV definition

With component prices $p_i(t)$ in \$/GPU-hr and weights w_i,

$$\mathrm{NAV}(t) \;=\; \sum_i w_i\, p_i(t), \qquad \sum_i w_i = 1,\; w_i \geq 0.$$

NAV is therefore itself a dollar cost per GPU-hour of the basket's mix.

4.3 Published level

The published series is the dollar NAV itself: the basket's blended cost in \$/GPU-hour at each date, from the first date the full basket prices (2026-03-01 for most baskets). No rebasing or normalization is applied, so levels are directly comparable to market rental rates.

4.4 The basket universe

Ticker	Basket	Mix
ITO-C01	H100 SXM On-Demand	100% H100 SXM, on-demand
ITO-C02	H100 PCIe On-Demand	100% H100 PCIe/generic, on-demand
ITO-C03	H100 Spot	100% H100, spot
ITO-C04	H200 On-Demand	100% H200 SXM, on-demand
ITO-C05	B200 On-Demand	100% B200, on-demand
ITO-C06	A100 80GB On-Demand	100% A100 80GB, on-demand
ITO-C07	MI300X On-Demand	100% AMD MI300X, on-demand
ITO-C08	Training Core	80% H100 SXM + 20% H200
ITO-C09	Frontier Training	50% B200 + 30% H200 + 20% H100 SXM
ITO-C10	Hopper Blend	60% H100 + 40% H200
ITO-C11	Ampere Blend	50% A100 80GB + 30% A100 40GB + 20% A10
ITO-C12	Cost-Optimized Training	60% H100 spot + 40% A100 spot
ITO-C13	Inference Fleet	40% L40S + 30% L4 + 30% A10
ITO-C14	Budget Inference	50% RTX 4090 + 50% RTX A6000
ITO-C15	LLM Serving	60% H100 + 40% L40S
ITO-C16	Legacy Fleet	50% V100 + 50% T4
ITO-C17	Balanced Datacenter	40% H100 + 30% A100 + 20% L40S + 10% L4
ITO-C18	Spot Diversified	40% H100 + 30% A100 + 30% RTX 4090, all spot
ITO-C19	Hybrid H100 Procurement	70% H100 on-demand + 30% H100 spot
ITO-C20	NVIDIA AMD Split	50% H100 + 50% MI300X

5 Historical Series

Daily NAV history runs from 2026-03-01 (over four months, exceeding the two-month design bar) and is served through the same public chart API as every other Itô basket. Provenance is retained per row (`source = opencomputeprices`). Research blends whose legs do not meet the two-month bar (e.g. an AMD MI300X/MI325X training blend — MI325X has only about one month of observed history) remain research-only and are not published.

6 Live Updates

Each published basket's legs map its components onto live one-minute provider series from the Itô compute feed (each component's weight split equally across its constituent series). On every live tick the edge NAV engine recomputes affected baskets as the weighted average of leg prices and publishes over the same SSE/WebSocket channels as all other baskets.

To keep the live series continuous with history, the engine anchors at seed time: with anchor A (the last published dollar NAV) and seed-time raw NAV B, each subsequent tick publishes $A \cdot \mathrm{nav}(t)/B$ — live prices move the series, while the level remains the historical one.

7 Risk Factors & Limitations

7.1 No standalone storage or fabric series

No independent market price series exists today for parallel-filesystem storage (e.g. WEKA) or standalone interconnect fabric. Storage- and fabric-priced legs are excluded rather than proxied; such components can be added when RFQ-derived series exist.

7.2 NVLink exposure is a form-factor proxy

NVLink value is expressed through the SXM form-factor class, not through an independently traded network price. Research spread series (e.g. H100 SXM minus PCIe) quantify that premium but are not published as public baskets: the live NAV engine prices long-only weight vectors, and a signed spread is not economically meaningful as a public price level.

7.3 List prices, not transacted prices

Both data layers observe published/listed rental prices, not negotiated or transacted prices. Large-contract pricing can differ materially.

7.4 Coverage gaps

Provider coverage varies by SKU and term; thin components carry wider uncertainty. Median-of-medians limits, but does not eliminate, composition drift when providers enter or leave the sample.

7.5 Historical/live source transition

History is OpenComputePrices; live ticks are Itô's own feed with a different (smaller, faster) provider mix. Ratio anchoring makes the level continuous, but short-horizon dynamics can differ across the seam.

8 Governance

Basket definitions, component mappings, and this methodology are version controlled. Changes to weights, component definitions, or data sources are released with a changelog entry and a version increment.

Changelog

- v1.0 (Q3 2026) — Initial publication: 20 long-only compute baskets, OpenComputePrices historical layer (2026-03-01 onward), live one-minute feed integration.