

Silicon Data GPU Rental Index™

Methodology Overview

1 Introduction

Silicon Data GPU Rental Index™ is the world's first daily benchmark for GPU rental prices, bringing pricing transparency to the highly heterogeneous AI compute market. Each day the index processes over 10,000 GPU rental records collected across 40+ independent, market-relevant cloud providers, spanning a diverse range of rental types and hardware configurations. A proprietary pricing model standardizes these heterogeneous configurations into a consistent baseline, ensuring all records are directly comparable before aggregation.

2 Index Methodology

Each index value is produced through four steps: preprocessing, filtering, standardization, and aggregation.

2.1 Step 1: Data Preprocessing

Disparate billing models (e.g., monthly bare-metal leases vs. hourly on-demand instances) are converted into a standardized per-hour, per-GPU rate:

$$\text{unit_price} = \frac{\text{total_rental_price}}{\text{num_gpu} \times \text{duration}}$$

where

- **total_rental_price** is the price to rent an entire server
- **num_gpu** is the count of GPUs on the server
- **duration** is the length of rental in hours

2.2 Step 2: Data Filtering

Preprocessed data undergoes multi-stage filtering to ensure reliability and remove reporting errors — without ignoring genuine structural price movements:

- **Completeness checks:** drops records lacking valid parameters
- **Anomaly checks:** flags significant day-over-day shifts for manual review
- **Outlier detection:** screens records against dynamic historical price bands

2.3 Step 3: Standardization (Basis Adjustment)

Simply aggregating observed prices would cause index fluctuations driven by the daily mix of configurations rather than true price movements. A proprietary pricing model applies a *basis adjustment* that standardizes each idiosyncratic record to a benchmark-equivalent contract, considering factors such as:

This document is being provided for your use as a Silicon Data customer and is subject to the terms and conditions of your applicable agreements with Silicon Data. This document may not be shared or distributed without Silicon Data's express written consent.

- **Rental Type:** On-Demand, Reserved, Spot (Interruptible)
- **Country:** US, AU, FR, ...
- **CPU Platform:** Intel, AMD, Other
- **GPU Variant:** SXM, PCIe, and other form factors

2.4 Step 4: Index Calculation

Basis-adjusted prices are first aggregated at the provider level, preventing any single provider with massive individual records from skewing the market average. The daily index is then calculated as a weighted average across providers:

$$\text{Index}(t) = \sum_{j=1}^N w_j(t) \bar{P}_j(t),$$

where

- $\bar{P}_j(t)$ denotes the average of all adjusted prices from the j-th cloud provider
- $w_j(t)$ denotes the weight reflecting the market relevance of the j-th cloud provider

3 Index Maintenance & Validation

3.1 Weekly Calibration

The pricing model is recalibrated weekly using recent market data, so basis adjustments adapt to fast-moving market conditions. The index calculation uses the latest model as of the index date.

3.2 Multi-tiered Validation

Dynamic reference ranges, derived from historical data at the provider-GPU variant level, trigger automated alerts when newly ingested data deviates significantly in price level or record count. Flagged records are held for manual review prior to inclusion.

3.3 Quarterly Review

Silicon Data announces material changes to an index in advance of implementation, and separately announces any decision to discontinue an index. Announcements are published on the Silicon Data website and distributed through partner news outlets, and set out the nature of the change, its effective date, and the estimated impact on index values or, in the case of discontinuation, any transition or successor arrangements.