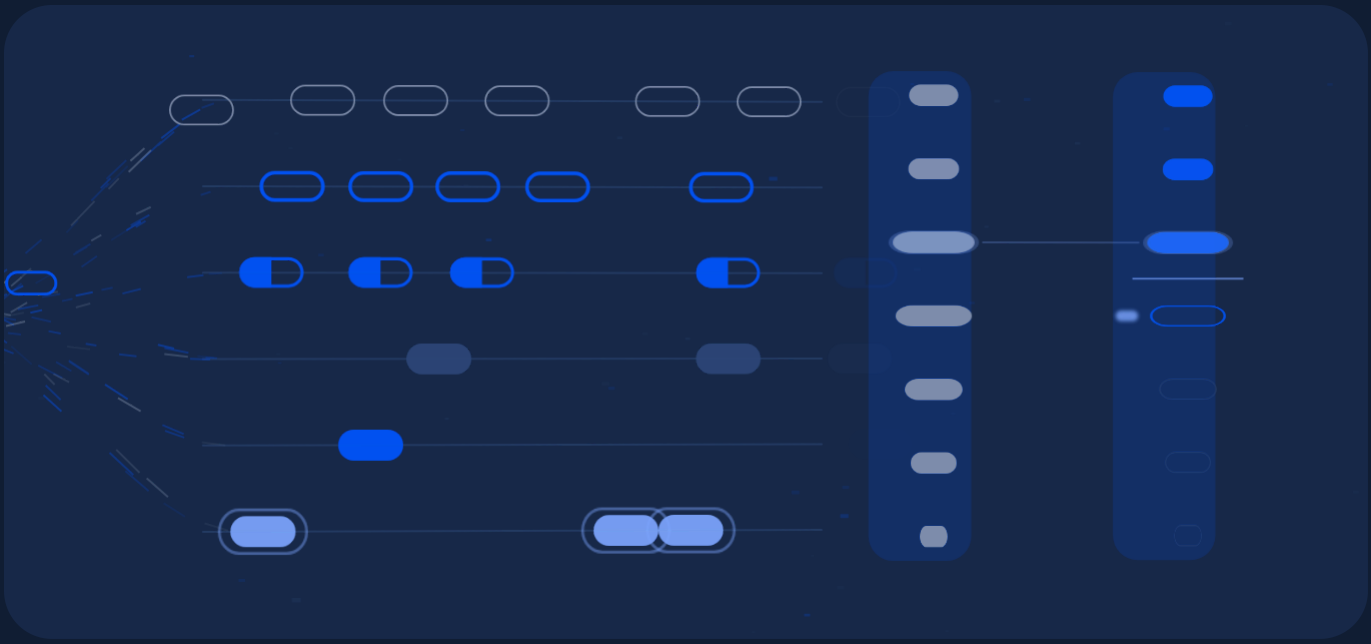
Rillor

EXPLAINER

# What sets the price of a used H100 server

A used H100 server price depends on what the figure covers and who states it, then on condition, warranty, diagnostic records, software licenses and export screening.



Rillor

# Contents


01 Summary

02 Background

03 Method

04 Findings

05 Implications

06 Limits of this analysis

07 Sources

# Summary

A used H100 server price depends on what the figure covers, who states it and what the unit's own records show, and the public figures reviewed here mix all three.

- Published figures for a used eight-GPU H100 server measure different things. A Hashrate Index article of 8 April 2026 gives $150,000 to $180,000 with no source, date or configuration [1]. American Compute's October 2026 value of $195,000 per node is built from listings adjusted to a reference configuration [2]. Two Dell XE9680 listings with eight H100 GPUs, viewed on 8 November 2026 and neither stating the unit's condition, asked $199,999 and $275,000 [13, 14]. None of these figures records a sale.
- Card prices do not add up to a server price. Publishers' 2026 estimated ranges for a used or secondary-market H100 SXM card run from Standard Mobile's $15,000 to $21,000 secondary-market range for the second quarter to BuySellRam's $21,000 to $28,000 for the first, with Mercatus and Compute Exchange inside that span [7, 9, 11, 12]. Servers with the same eight GPUs differ in host CPUs, memory, network cards and drives, and one listing, "Supermicro SYS-821GE-TNHR Barebone - 8U - 8x NVIDIA HGX H100 SXM5 - 2x Intel Xeon 4th Gen", names eight GPUs but lists no GPU baseboard in its contents [13, 15, 16, 18].
- Asking prices run above sales. In American Compute's October 2026 Rack Report, a used H100 SXM card was listed at a median of $16,576 and sold for $14,250 typical, a ratio of 1.16 [3, 4]. In October 2026 the same publisher put one HGX H100 node at $216,000 installed, $195,000 sold whole, $173,000 in an orderly sale and $126,000 in a forced sale, all estimates [3].
- Condition labels have no shared definition. eBay separates manufacturer-approved refurbishment, refurbished grades backed by a one-year warranty, seller refurbishment and plain used [21]. Compute Exchange and BuySellRam price refurbished H100 SXM cards above used ones, while Mercatus prices them below [7, 8, 9, 11].
- Warranty depends on who issued it and when. Supermicro's standard warranty for GPU systems runs three years for labor and one year for parts from Supermicro's invoice date, and applies only to the original purchaser [26]. Compute Exchange states that refurbisher warranties typically run 90 to 365 days, scoped to the refurbisher [8].
- A used H100's memory history stays with the card. The nvidia-smi tool keeps aggregate memory error counts as a lifetime counter that cannot be cleared on Ampere and later GPUs, and reports whether a row remapping, which the GPU applies to memory rows with uncorrectable errors, has failed in the past [22, 23]. NVIDIA's Field Diagnostic decides whether a GPU qualifies for return, and NVIDIA states that its DCGM diagnostic suite is not designed to replace field diagnosis tools [23, 24].
- Licenses and export rules follow the hardware. Included NVIDIA AI Enterprise subscriptions come with H100 PCIe and NVL, H200 NVL and A800 40GB Active cards, tied to the card and its server, and NVIDIA lists the software as an add-on for H100 SXM GPUs [17, 27]. Guidance from the U.S. Bureau of Industry and Security (BIS) counts servers among the items that contain advanced computing integrated circuits, and when red flags are raised an exporter has a duty to check them out [28, 29].

# Background

## Who needs a used H100 price

Buyers weighing secondary HGX H100 systems, owners pricing retired systems and lenders checking collateral each need a figure they can interpret, and the checks that change what one unit is worth.

## How H100 systems reached the secondary market

Hashrate Index reports that retail H100 prices held at $25,000 to $40,000 per GPU from mid-2024 into early 2026, and that used and refurbished units traded as high as $50,000 per GPU during the mid-2024 scarcity before dropping sharply [1]. Makers have since ended sales of some H100 systems. Microway relays that NVIDIA's DGX H100 end-of-sale program took final purchases by 25 January 2025 [19]. Cisco announced the end of life of its UCS C885A M8 servers with HGX H100 GPUs on 24 April 2025 and set their end of sale for 19 July 2025 [20].

American Compute's GPU Resale Index tracks what a fixed basket of used data center GPUs aged 2.5 to 8 years, from the T4 to the H100, sells for as a share of its price new, with the first half of 2025 set to 100 [5]. In the October 2026 Rack Report it read 75 for August 2026, down 12% on July, and the publisher writes that "Cards older than Hopper led the fall, while used HGX H100 servers held their price" [5]. The November 2026 issue, released 3 November, puts it at 83 for September, down 4.3% on a restated August of about 87 (arithmetic on [5, 36]).

## The six evidence classes used here

This report sorts every figure into the six evidence classes of Rillor's GPU systems pricing method. A specification gives system and component attributes as the manufacturer publishes them. A listing is a public offer, seen at a stated time. A quote is a price supplied on request for a stated configuration and quantity. An indicative price is a reported price level without a verifiable counterparty. A reported sale is a sale a source reports, without independent confirmation. A verified transaction is a completed sale with confirming evidence.

# Method

**What was examined.** Price figures for used and refurbished H100 servers and cards from six publishers, dated February to November 2026 [1-12, 36], and four seller pages for eight-GPU H100 servers [13-16]. NVIDIA's H100 page and DGX H100 user guide, and two end-of-sale notices [17-20]. One marketplace's condition definitions [21], NVIDIA's diagnostic and memory error documentation and one field study [22-25]. A server maker's warranty terms and NVIDIA's licensing guide [26, 27], export rules and guidance and an NVIDIA quarterly filing [28-30], a rental price notice and three useful-life disclosures [31-34], and NIST's media sanitization guide [35].

**How.** Each figure is given with its publisher, its date and its evidence class. Listings carry their capture date. No figures are averaged, and figures of different classes sit in separate rows. For each server figure, the report records whether it prices a GPU card, a server or a wider system, and which build it states. Arithmetic done for this report is labelled. All web pages were viewed on 8 November 2026, and [36] on 9 November. Seller pages that refused automated access are not cited.

# Findings

## 1. Why published used H100 figures disagree

Hashrate Index, a publication of Luxor Technology, stated on 8 April 2026 that "a used H100 server (8-GPU) trades at roughly" $150,000 to $180,000 [1]. The article names no source, date range or configuration for the figure. Its chain of custody ends at the article, so the figure is an indicative price whose basis cannot be checked. The same article describes public GPU pricing as "a snapshot that mixes reseller asks, refurbished inventory, and marketplace listings", and adds that private deals can clear materially lower [1].

American Compute's HGX H100 page, updated 17 October 2026 from its October 2026 Rack Report with data through August 2026, gives a fair market value of $195,000 per node for an HGX H100 8-GPU server with eight 80 GB GPUs, or $24,400 per GPU [2]. Its method note says the figure comes from "observed listings for whole nodes, adjusted to a reference configuration, not realized sales", and the page calls it "Indicative, not an appraisal" [2]. The same publisher describes its Rack Report as built from "realized resale transactions contributed by brokers, resellers, and ITADs", with dealer asks tracked separately to measure the gap between asks and what closes [6].

Compute Exchange gives a used H100 SXM5 card an indicative range of $22,000 to $27,000 for the third quarter of 2026, "Based on aggregated supplier quotes and broker market data", and states that it "does not publish active counterparty listings" [7]. It adds that final pricing depends on quantity, region, condition and warranty terms [7].

Asks and sales differ in a measured way. American Compute's liquidation page reports that a used H100 SXM card "was listed at a median $16,576 and sold for $14,250 typical in the October 2026 Rack Report", and its valuation page states the same gap as asks running 1.16 times the realized price [3, 4]. The page draws the consequence: "A fair market value built on listings starts too high, and every liquidation value derived from it does too" [4].

**Table 1. Card-level figures for a used or refurbished H100 SXM 80 GB, by source**

| Publisher | Date or period | Used | Refurbished | Evidence class |
|---|---|---|---|---|
| **Standard Mobile** | 10 April 2026 (second quarter of 2026) | $15,000 to $21,000 (secondary market, used and refurbished not split) | Not given | Indicative price [12] |
| **American Compute** | October 2026 Rack Report (August data) | $16,576 median listing, $14,250 typical sale | Not given | Listing and reported sale [4] |

| Publisher | Date or period | Used | Refurbished | Evidence class |
|---|---|---|---|---|
| **BuySellRam** | 18 February 2026 (first quarter of 2026) | $21,000 to $28,000 (good) | $30,000 to $34,000 | Indicative price [11] |
| **Mercatus** | 27 March and 23 June 2026 | $22,500 to $25,500 (18 to 30 months old) | $18,000 to $22,000 | Indicative price [9, 10] |
| **Compute Exchange** | Third quarter of 2026 | $22,000 to $27,000 | $26,000 to $30,000 | Indicative price [7, 8] |

The rows differ in class, period and scope, so their spread says little about the price of any one card.

## 2. Card, server or system: what a figure covers

NVIDIA's H100 page compares two versions. Its specification table lists the H100 SXM with 80 GB of memory and up to 700 W, for HGX H100 partner and NVIDIA-Certified Systems with four or eight GPUs and for the DGX H100 [17]. It lists the H100 NVL with 94 GB, 350 to 400 W and a dual-slot air-cooled PCIe form factor, for partner and certified systems with one to eight GPUs [17]. Standard Mobile notes that SXM GPUs "require specific baseboard and cooling infrastructure, limiting the buyer pool" [12].

A server figure covers more than eight cards. American Compute's reference node is a whole server in a standard configuration, with CPUs, memory, storage and node network cards, and leaves out fabric, racks, installation and support [2]. Mercatus, in a page of 27 March 2026 updated 9 October 2026, lists $250,000 to $320,000 for a new, fully built 8-GPU HGX H100 server, or $31,000 to $40,000 per GPU "including server, NVLink, and networking", and calls about $285,000 a typical integrated quote [9]. Selling a server whole also differs from selling its parts. American Compute reports an HGX H100 node trading 4% above the sum of its parts in its October 2026 Rack Report, against 40% for an HGX A100 80GB node [2, 4].

NVIDIA's DGX H100 gives a reference build to compare listings against, part by part [18].

**Table 2. Three eight-GPU H100 builds**

| Part | NVIDIA DGX H100 [18] | Dell XE9680, The Server Store [13] | Dell XE9680, Interpro (decommissioned) [16] |
|---|---|---|---|
| **Host CPUs** | 2 x Xeon 8480C, 56 cores each | 2 x Xeon Platinum 8462Y+, 32 cores each | 2 x Xeon Platinum 8558, 48 cores each |

| Part | NVIDIA DGX H100 [18] | Dell XE9680, The Server Store [13] | Dell XE9680, Interpro (decommissioned) [16] |
|---|---|---|---|
| System memory | 2 TB in 32 DIMMs | 2,048 GB (32 x 64 GB DDR5-4800) described | 32 x 64 GB DDR5-4800 |
| Cluster network | 8 x ConnectX-7 single-port InfiniBand | 1 x ConnectX-7, plus a dual-port ConnectX-5 EDR and 100GbE card | 8 x ConnectX-7 400G |
| Drives | 2 x 1.92 TB M.2 (OS), 8 x 3.84 TB U.2 self-encrypting (data) | No hard drives, trays included | 8 x 3.84 TB Samsung PM893, 2 x 480 GB M.2 |
| Power supplies | 6 x 3.3 kW | 6 x 2,800 W | Not stated |
| Price | End of sale, final purchases by 25 January 2025 [19] | $199,999 listing, out of stock | Contact sales |

Listings viewed 8 November 2026.

The Server Store's option list prices memory upgrades separately, for example "2TB (32x 64GB) 5600MHz DDR5 Registered Memory (+$92,000.00)", so the page does not make clear whether its $199,999 total includes the 2,048 GB its description names [13]. Vipera's XE9680 bundle, listed at $275,000 on 8 November 2026, was in stock with a lead time of three to four weeks, and its page stated no condition, CPU, memory, storage or network build [14]. On 8 November 2026 the same page showed an HPE Cray XD670 HGX H100 system at $288,500 [14]. Neither can be compared with Table 2 without a quote.

A title that names the GPUs does not settle scope. On 8 November 2026 SabrePC listed "Supermicro SYS-821GE-TNHR Barebone - 8U - 8x NVIDIA HGX H100 SXM5 - 2x Intel Xeon 4th Gen" at $23,283.36, out of stock [15]. Its product type is "Server Barebone System", and its contents list a motherboard, chassis, fans, IO, NIC, PCIe switch and midplane boards, a backplane, power distribution boards, cables and 3,000 W power supplies (six installed, per its specifications), with no GPU baseboard [15].

## 3. Condition labels and what each asserts

No condition scale is shared across sellers, and none of the sources defines one for the industry. eBay publishes its own labels for computers and electronics [21]. "Certified - Refurbished" means the item was "professionally inspected, cleaned, and refurbished by the manufacturer or a manufacturer-approved vendor to meet manufacturer specifications." Its Excellent, Very Good and Good refurbished grades are each "backed by a one year warranty". "Seller refurbished" means "restored to working order by the eBay

seller or a third party". "Used" means "previously used" and "fully operational", with possible cosmetic wear [21]. The same word, refurbished, can therefore mean manufacturer-approved work with no warranty stated, refurbishment by qualified sellers backed by a one-year warranty, or seller work with no warranty stated [21].

In its article of 8 April 2026, Hashrate Index draws the line by risk: "Used implies as-is from a prior deployment", with the buyer carrying more risk, while "Refurbished implies the hardware has been tested, certified, and potentially comes with a warranty" [1]. It puts refurbished H100s at 15 to 25 percentage points above used ones [1]. Compute Exchange's third-quarter 2026 page describes a 10 to 15 percent premium for refurbished cards "with documented diagnostic and burn-in testing" [7]. BuySellRam's first-quarter 2026 table also places refurbished above used [11]. Mercatus, in pages of 27 March and 23 June 2026, reverses the order: refurbished H100s from data center decommissioning sell for $18,000 to $22,000 in 2026, below its $22,500 to $25,500 for used cards 18 to 30 months old [9, 10].

Age is a separate variable. In a page of 23 June 2026, against a $30,000 OEM list reference, Mercatus gives 85% to 90% for cards 6 to 12 months old, 75% to 85% at 18 to 30 months, 40% to 75% at 36 months and 25% to 35% at 60 months and more [10].

Compute Exchange lists what a typical H100 SXM5 refurbishment covers: NVIDIA Field Diagnostic runs, nvidia-smi memory error and thermal checks, "full-power stress testing for 24 to 72 hours", connector cleaning, and "Output ships with a per-unit test report" [8]. Hashrate Index adds that memory and storage often need refreshing on used GPU servers [1]. A condition label carries weight only to the extent that such records stand behind it.

## 4. Warranty: who holds it and for how long

Supermicro's standard limited warranty table gives GPU systems three years of labor, one year of parts and one year of advance cross-ship [26]. Products are warranted "for a period of time from the date of Supermicro's invoice", and the table applies "In the absence of special warranty terms set forth in a corresponding invoice" [26]. The warranty "applies only to the original purchaser", is voided by "removal, alteration, or destruction of the original product identification markings", and does not cover "products, components, or software supplied by third parties" [26]. A return needs the "invoice date, invoice number, product model name, serial number, and reason for return" [26]. On the table terms, a GPU system invoiced before 8 November 2023 was past both terms on 8 November 2026 (arithmetic on [26]).

Under such terms, a second owner relies on the seller's warranty. Compute Exchange states that "Most refurbishers offer a 90 to 365 day limited warranty against defects, scoped to the refurbisher rather than NVIDIA" [8]. Compute Exchange's FAQ also states that the original NVIDIA warranty does not transfer with the unit [8]. Mercatus cites a 12-month limited warranty for refurbished cards [9, 10]. The Server Store's XE9680 page reads "All servers are fully tested and comes with a 90 warranty", with no unit stated [13]. Return terms vary too. Vipera's page says "All sales final. No returns or cancellations." [14], and SabrePC's says "Non-cancelable and non-returnable" [15].

Support calendars bound what a used unit can still get from its maker. Cisco's end-of-life notice for its UCS C885A M8 servers with HGX H100 GPUs sets the end of software maintenance releases and of routine failure analysis at 19 July 2026, and the last date of hardware support at 31 July 2030 [20].

## 5. Testing a used system with GPU diagnostics

An H100 keeps records that a seller cannot reset. NVIDIA's nvidia-smi documentation states that "Volatile error counters track the number of errors detected since the last driver load" and that "Aggregate error counts persist indefinitely and thus act as a lifetime counter" [22]. For the option that clears them, it states: "Clearing aggregate counts is not supported on Ampere+" [22]. The H100 is a Hopper GPU, and the A100 an Ampere GPU [25]. Each board's serial number "matches the serial number physically printed on each board" and is "a globally unique immutable alphanumeric value" [22], so records can be tied to a physical card.

The same tool reports row remapping, which the GPU applies to memory rows with uncorrectable errors [23]. It shows whether a row is waiting for remapping, which takes effect only after a reset ("The GPU must be reset for the remapping to go into effect"), whether "a row remapping has failed in the past", and a bank remap availability histogram from Maximum to None [22]. NVIDIA's memory error return (RMA) policy states that "The NVIDIA Field Diagnostic tool determines whether a GPU qualifies for RMA" and that, for row remapping, "the RMA criteria is met when the row-remapping failure flag is set and validated by the field diagnostic" [23]. The flag is set when a remap is attempted on a bank that already has eight uncorrectable error rows remapped, on a row already remapped, or after 512 total remappings for uncorrectable errors [23].

A seller's "tested" can mean very different run times. NVIDIA's DCGM diagnostics come in four run levels [24].

Table 3. DCGM diagnostic run levels, measured on Hopper GPU systems

| Level | Run time on 8-GPU systems | Tests added at this level |
|---|---|---|
| **r1 (Short)** | Under 2.5 seconds | Software |
| **r2 (Medium)** | Under 10.5 minutes | PCIe and NVLink, GPU memory, memory bandwidth |
| **r3 (Long)** | Under 35 minutes | Diagnostic, targeted stress, targeted power, NVBandwidth, NCCL tests |
| **r4 (Extra Long)** | Under 2.25 hours | Memory stress (memtest), input EDPp (pulse) |

Source: [24]. The r1 time is given without a GPU count.

NVIDIA states that DCGM Diagnostics are not designed to "Provide comprehensive hardware diagnostics", "Replace the field diagnosis tools" or "Facilitate any RMA process" [24]. A DCGM pass is therefore a different record from a Field Diagnostic result. Compute Exchange's list of what to verify on a used H100 SXM5 names "NVIDIA Field Diagnostic output, NVSMI ECC error counts, runtime hours, thermal logs from the prior deployment, and physical inspection of the SXM5 mezzanine connector and HBM stack" [7].

In one field study, memory history weighs more on the H100 than on the A100. Cui and others analyzed 2.5 years of operational data, 11.7 million GPU hours, from 1,056 A100 and H100 GPUs on the Delta system [25]. They found H100 memory resilience worse than A100, "with 3.2x lower per-GPU MTBE for memory errors", where MTBE is mean time between errors, and found that H100 GPUs "demonstrate significantly improved GPU hardware resilience over A100 GPUs with respect to critical hardware components" [25].

## 6. Drives left in the server

Drives in a used server come from a prior owner. Interpro's decommissioned XE9680 is listed with eight 3.84 TB drives and two 480 GB boot drives [16], and NVIDIA's DGX H100 ships with eight self-encrypting U.2 data drives [18]. NIST SP 800-88 Rev. 2 defines media sanitization as "a process that renders access to target data on the media infeasible for a given level of effort" [35]. American Compute states its liquidation values as gross, "before de-installation, shipping and data sanitization" [2], so those costs sit outside its figures. A unit sold without drives, as The Server Store's XE9680 is listed [13], leaves storage to the buyer.

## 7. Software licenses tied to the original purchase

NVIDIA AI Enterprise "is licensed on a per-GPU basis" [27]. NVIDIA's licensing guide states that "Each NVIDIA H100 PCIe or NVL Tensor Core GPU includes a five-year NVIDIA AI Enterprise subscription", that each H200 NVL GPU includes a five-year subscription, and that the A800 40GB Active includes three years [27]. The included subscriptions apply to H100 PCIe and NVL, H200 NVL and A800 cards, not to the H100 SXM GPUs on an HGX board, for which NVIDIA's H100 specification table lists NVIDIA AI Enterprise as an "Add-on" [17].

Where a subscription is included, it is bound to the hardware. The software "can be installed on the NVIDIA Certified Systems in which the GPUs are installed, but not on a different server or instance" [27]. Its start date "is based on the ship date of the GPU board to the OEM partner, plus 90 days", and "cannot be modified, as it is tied to the specific card" [27]. A five-year subscription on a used H100 NVL card therefore runs out five years after that start date, however long the card sat before resale (arithmetic on [27]). For DGX systems, "NVIDIA AI Enterprise is included in the NVIDIA DGX software bundle for NVIDIA DGX systems using the NVIDIA Hopper architecture" [27].

## 8. Export screening on resale

U.S. export controls attach to systems as well as chips. NVIDIA's Form 10-Q for the quarter ended 26 July 2026 states that the restrictions announced in August 2022 "specifically impact our A100 and H100 integrated circuits, DGX or any other systems or boards which incorporate A100 or H100 integrated circuits"

[30]. BIS's guidance of 13 May 2025 on diversion of advanced computing integrated circuits states in its first footnote that these chips and the commodities that contain them include Export Control Classification Numbers "3A090.a, 4A090.a, and .z items in Categories 3, 4, and 5, such as servers classified as ECCN 5A992.z" [29].

BIS's Know Your Customer guidance in the Export Administration Regulations lists red flags [28]. They include a customer or purchasing agent "reluctant to offer information about the end-use of a product", a product whose "capabilities do not fit the buyer's line of business", a customer who declines routine installation, training or maintenance services, vague delivery dates or "out-of-the-way destinations", and "A freight forwarding firm is listed as the product's final destination" [28]. The guidance states that, absent red flags or an express requirement in the EAR, "there is no affirmative duty upon exporters to inquire", and that when red flags are raised "you have a duty to check out the suspicious circumstances" [28]. Its steps include "Do not self-blind" and "Refrain from the transaction or advise BIS and wait" [28].

The 2025 guidance adds due diligence steps for advanced computing items: seek "an end-user certification with detailed information on all proposed transaction parties", and request a written attestation from the data center that the end user is authorized to operate at its location and has the infrastructure to operate the type of server being exported [29]. One seller in this review states on its site that "End-user forms must be completed before we can share quotes" [14]. This section summarizes the regulations' text and is not legal advice.

## 9. Reading a used price against comparable observations

One server has several values, depending on where it sits and how fast it sells. American Compute's valuation page gives each for the same HGX H100 node [3].

Table 4. Values for one HGX H100 8-GPU node, October 2026

| Basis | Value | What it assumes |
|---|---|---|
| **Income value in place (limited sample)** | $216,000 | Installed and earning |
| **Fair market value, removed** | $195,000 | Sold whole between a willing buyer and seller |
| **Parted out** | About $188,000 | Sold as components |
| **Orderly liquidation value** | $173,000 | About 30 days to resell after remarketing, diagnostics and refurbishment |

| Basis | Value | What it assumes |
|---|---|---|
| **Forced liquidation value** | $126,000 | Forced sale through brokers and resellers in parallel |

Source: [3]. All values indicative.

The orderly and forced values are 89% and 65% of fair market value [2]. American Compute defines fair market value in the appraisers' terms and adds that "For GPUs the evidence is completed arm's-length sales, not dealer asks or rental rates" [3]. It also states that "There is no central auction house for data center GPUs", and that parting out a cluster cuts the buyer pool to those who can rebuild it [3].

The installed value rests on what a server earns, and AWS cut on-demand prices for its H100-based P5 instances by 44% from 1 June 2025 [31].

Book lives are accounting estimates and say nothing about what a used server sells for. Amazon changed the useful life of a subset of its servers and networking equipment from six years to five from 1 January 2025, "due to the increased pace of technology development" [32]. Meta raised the estimated useful life of certain servers and network assets to 5.5 years from fiscal 2025 [33]. CoreWeave lists technology equipment at six years [34].

Forward views are forecasts and carry their dates. American Compute's GPU Resale Index page lists a conservative residual band for H100s of 11% to 19% of the price new in 2027, 5% to 7% in 2028 and 2% to 3% in 2029, forecasts its HGX H100 page attributes to its June 2026 residual value report [2, 5]. That HGX H100 page, updated 17 October 2026, puts a server bought used at $195,000 at $14,000 to $22,000 after three years on its band forecasts [2]. NVIDIA's 10-Q states that it "may provide limited residual-value support for a portion of specific projects", and lists a "Form of Residual Value Guaranty" as an exhibit [30].

A used price is therefore read against observations of the same configuration, condition, seller type and evidence class (Findings 1 to 7).

# Implications

**For buyers.** Ask what the figure covers before comparing it. Read a quote for a used HGX H100 server against a part-by-part build (Table 2), the warranty provider and its start date [26], the Field Diagnostic and nvidia-smi records for each serial number [22, 23], the license terms for the GPUs in it [17, 27] and the return terms [14, 15]. A DCGM pass at level 1 and a 2.25-hour level 4 run are different evidence [24].

**For sellers.** Records support a price: per-card serial numbers, aggregate error counts that cannot be cleared, the remap history and a dated diagnostic report [22]. Compute Exchange's refurbished grade rests on such checks and a per-unit test report [8]. Buyer screening is part of the sale, and the red flags are written in the regulation [28, 29].

**For lenders and analysts.** Each figure needs its date, its scope and its evidence class. A listing-based value starts above realized sales [4], and liquidation values derived from it carry the same bias [4]. Installed and forced values differ by $90,000 on one node (arithmetic on [3]). Rillor's GPU systems pricing method records each configuration's condition, such as new, refurbished or used, and its seller type, such as distributor, reseller or broker. Two systems that differ only in condition or seller are different configurations, and each observation keeps its evidence class, so a listing is read as a listing and a quote as a quote.

# Limits of this analysis

- No public auction result, hammer price or verified transaction for a used H100 server appears among the sources.
- Every price here is an indicative price, a listing seen on 8 November 2026, a typical sale reported by one publisher or a value derived from listings.
- The Hashrate Index server figure has no stated source or configuration [1].
- American Compute states that an issue can restate recent months as late sales come in [5]. Its server values are listing-based, while its Rack Report describes realized sales as its basis [2, 6].
- Several seller pages refused automated access, and none of their prices is cited. eBay's condition definitions are cited, and no eBay listing is.
- Warranty terms are given only for Supermicro, from its own policy [26]. Dell, HPE and NVIDIA warranty terms were not reviewed.
- The field study of H100 memory errors covers one system [25] and gives error rates by GPU type, with no failure rates by age.
- The export section summarizes the regulations' text and names no destination. It is not legal advice.
- This report describes what public sources show. It is not investment, legal, tax or accounting advice.

# Sources


1. Used GPUs for AI: The Secondary Market Reshaping Enterprise Compute (Ian Philpot). Hashrate Index, Luxor Technology, 8 April 2026, viewed 8 November 2026. https://hashrateindex.com/blog/used-gpu-market-pricing-deprecation-secondary-ai/

2. HGX H100 server value: $195,000 per node (Oct 2026). American Compute, updated 17 October 2026, viewed 8 November 2026. https://www.amcompute.com/gpu-server-values/h100

3. In-situ vs ex-situ GPU value: installed vs removed and sold (Oct 2026). American Compute, updated 17 October 2026, viewed 8 November 2026. https://www.amcompute.com/gpu-server-values/in-situ-vs-ex-situ

4. FMV vs OLV vs FLV: liquidation value for GPU collateral. American Compute, viewed 8 November 2026. https://www.amcompute.com/appraisals/liquidation-value

5. GPU Resale Index: used GPU prices by month. American Compute, viewed 8 November 2026. https://www.amcompute.com/gpu-resale-index

6. Rack Report: datacenter hardware value benchmark. American Compute, viewed 8 November 2026. https://www.amcompute.com/rack-report

7. Used H100 SXM5, used GPU. Compute Exchange, viewed 8 November 2026. https://compute.exchange/hardware-market/used-gpus/used-h100-sxm5

8. Refurbished H100 SXM5, refurbished GPU. Compute Exchange, viewed 8 November 2026. https://compute.exchange/hardware-market/refurbished-gpus/refurbished-h100-sxm5

9. NVIDIA H100 price 2026: $25K GPU, $285K HGX server. Mercatus, 27 March 2026, modified 9 October 2026, viewed 8 November 2026. https://www.mercatus-ai.com/blog/h100-gpu-cost

10. NVIDIA H100 resale value 2026: used sells at $22K to $25K. Mercatus, 23 June 2026, modified 8 July 2026, viewed 8 November 2026. https://www.mercatus-ai.com/blog/h100-resale-value

11. NVIDIA GPU cluster liquidation: maximize ROI and asset recovery. BuySellRam, 18 February 2026, viewed 8 November 2026. https://www.buysellram.com/blog/nvidia-a100-h100-h200-cluster-liquidation-maximize-roi-and-asset-recovery/

12. GPU liquidation: what happens when hyperscalers dump their NVIDIA clusters. Standard Mobile Company, 10 April 2026, viewed 8 November 2026. https://standardmobileco.com/articles/gpu-liquidation-hyperscalers-nvidia-clusters

13. Dell PowerEdge XE9680 6U GPU rackmount server (8x NVIDIA Hopper H100 80GB SXM5), listing. The Server Store, viewed 8 November 2026. https://www.theserverstore.com/dell-poweredge-xe9680-6u-gpu-rackmount-server.html

14. Dell PowerEdge XE9680 rack server 210-BFXH-BUNDLE with 8 x H100 GPUs, listing. Viperatech (Vipera LLC), viewed 8 November 2026. https://viperatech.com/product/dell-poweredge-xe9680-rack-server-210-bfxh-bundle-with-8-x-h100-gpus

15. Supermicro SYS-821GE-TNHR barebone, 8U, 8x NVIDIA HGX H100 SXM5, 2x Intel Xeon 4th Gen, listing. SabrePC, viewed 8 November 2026. https://www.sabrepc.com/SYS-821GE-TNHR-Supermicro-S133318654

16. Decommissioned Dell Edge XE9680 6U server with 8x H100 GPUs, listing. Interpro Microsystems, viewed 8 November 2026. https://www.interpromicro.com/product/40091-dell-edge-xe9680-6u-server-with-8x-h100-gpus

17. NVIDIA H100 GPU. NVIDIA, viewed 8 November 2026. https://www.nvidia.com/en-us/data-center/h100/

18. Introduction to NVIDIA DGX H100/H200 systems, NVIDIA DGX H100/H200 User Guide. NVIDIA, viewed 8 November 2026. https://docs.nvidia.com/dgx/dgxh100-user-guide/introduction-to-dgxh100.html

19. EOL, NVIDIA DGX H100. Microway, viewed 8 November 2026. https://www.microway.com/product/nvidia-dgx-h100/

20. End-of-Sale and End-of-Life Announcement for the Cisco UCS C885A M8 Rack Servers with NVIDIA HGX H100 GPUs (EOL15657). Cisco, 24 April 2025, viewed 8 November 2026. https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/nvidia-hgx-h100-gpus-eol.html

21. Item condition by category. eBay, viewed 8 November 2026. https://www.ebay.com/help/selling/listings/creating-managing-listings/item-conditions-category?id=4765

22. NVIDIA System Management Interface (nvidia-smi) documentation. NVIDIA, viewed 8 November 2026. https://docs.nvidia.com/deploy/nvidia-smi/index.html

23. RMA policy, NVIDIA GPU Memory Error Management. NVIDIA, viewed 8 November 2026. https://docs.nvidia.com/deploy/a100-gpu-mem-error-mgmt/rma-policy-thresholds-for-row-remapping.html

24. DCGM Diagnostics, NVIDIA DCGM User Guide. NVIDIA, viewed 8 November 2026. https://docs.nvidia.com/datacenter/dcgm/latest/user-guide/dcgm-diagnostics.html

25. Story of Two GPUs: Characterizing the Resilience of Hopper H100 and Ampere A100 GPUs. Cui and others, arXiv:2503.11901, version 4, 10 December 2025 (first version 14 March 2025), viewed 8 November 2026. https://arxiv.org/abs/2503.11901

26. Supermicro Warranty (standard limited warranty). Supermicro, viewed 8 November 2026. https://www.supermicro.com/en/support/warranty

27. NVIDIA AI Enterprise Licensing, NVIDIA Enterprise Licensing Guide. NVIDIA, viewed 8 November 2026. https://docs.nvidia.com/ai-enterprise/planning-resource/licensing-guide/latest/licensing.html

28. Supplement No. 3 to Part 732, BIS's "Know Your Customer" Guidance and Red Flags, 15 CFR. Legal Information Institute, Cornell Law School, viewed 8 November 2026. https://www.law.cornell.edu/cfr/text/15/appendix-Supplement_No_3_to_part_732

29. Industry guidance to prevent diversion of advanced computing integrated circuits. U.S. Bureau of Industry and Security, 13 May 2025, viewed 8 November 2026. https://www.bis.gov/media/documents/ai-counter-diversion-industry-guidance-may-13-2025.pdf

30. NVIDIA Corporation Form 10-Q for the quarter ended 26 July 2026. U.S. Securities and Exchange Commission, filed 26 August 2026, viewed 8 November 2026. https://www.sec.gov/Archives/edgar/data/1045810/000104581026000075/nvda-20260726.htm

31. Announcing up to 45% price reduction for Amazon EC2 NVIDIA GPU-accelerated instances (Channy Yun). AWS News Blog, Amazon Web Services, 5 June 2025, viewed 8 November 2026. https://aws.amazon.com/blogs/aws/announcing-up-to-45-price-reduction-for-amazon-ec2-nvidia-gpu-accelerated-instances/

32. Amazon.com, Inc. Form 10-K for 2025. U.S. Securities and Exchange Commission, filed 6 February 2026, viewed 8 November 2026. https://www.sec.gov/Archives/edgar/data/1018724/000101872426000004/amzn-20251231.htm

33. Meta Platforms, Inc. Form 10-K for 2024. U.S. Securities and Exchange Commission, filed 30 January 2025, viewed 8 November 2026. https://www.sec.gov/Archives/edgar/data/1326801/000132680125000017/meta-20241231.htm

34. CoreWeave, Inc. Form 10-K for 2025. U.S. Securities and Exchange Commission, filed 2 March 2026, viewed 8 November 2026. https://www.sec.gov/Archives/edgar/data/1769628/000176962826000104/crwv-20251231.htm

35. Guidelines for Media Sanitization, NIST SP 800-88 Rev. 2. National Institute of Standards and Technology, September 2025, viewed 8 November 2026. https://csrc.nist.gov/pubs/sp/800/88/r2/final

36. Rack Report: November 2026. American Compute, released 3 November 2026, viewed 9 November 2026. https://www.amcompute.com/rack-report/2026-11

## About this report

Published by Rillor on 10 November 2026. Superintelligence (SI) tools assist research and drafting, and every figure cites its source. Rillor builds and secures SI systems, supplies the data they learn from and delivers the infrastructure they run on. We apply SI to financial and compute markets.

Cite as: Rillor. What sets the price of a used H100 server. 10 November 2026. https://rillor.com/insights/used-h100-server-prices

## Notices

Rillor is not a registered investment adviser or commodity trading advisor and does not provide investment or trading advice. Research, forecasts, data and software described on this site are for research and engineering use. Nothing here is an offer or recommendation to buy or sell any security, commodity interest or digital asset. Past or simulated results do not indicate future results.

NVIDIA, AMD, Ceph, Kubernetes and related product names are trademarks of their owners and are used only to identify products. No affiliation or endorsement is implied.

## Work with Rillor

Send the systems, configurations and conditions you need priced, and what the data has to support. Contact us at rillor.com/contact?type=index or contact@rillor.com. More on GPU systems pricing: rillor.com/markets/gpu-pricing#systems