



GPU Useful Life: Depreciation, Rental Income and Recovery Value

Plan and build private AI compute with GPU Smith. We connect workload requirements with GPU hardware, networking, deployment and operating decisions so your infrastructure fits the work it must perform.

Published by GPU Smith · 2026-10-15 · gpu useful life · gpu depreciation · gpu economic useful life · gpu residual value · gpu collateral valuation · h100 · b200 · b300 · gpu financing

Original article: <https://gpusmith.com/articles/en/gpu-collateral-useful-life>



In this article

1. Executive Summary
 2. Introduction and Background
 3. Key Changes: Four Clocks Behind GPU Useful Life
 4. Implementation Considerations and Process Changes
 5. Market Evidence and Comparable Collateral
 6. Cash Generation and GPU Economic Useful Life
 7. GPU Collateral Recovery Value: From Gross Price to Cash
 8. Data Analysis and Evidence
 9. Implications and Future Directions
 10. Conclusion
 11. References
-

Executive Summary

GPU useful life is a set of different estimates, not a single expiry date. A graphics processing unit (GPU) has several relevant clocks. A lender needs to distinguish **accounting depreciation**, **physical serviceability**, **profitable workload life**, and **recoverable net sale proceeds**. Public filings answer the first question at an asset-category level. CoreWeave's fiscal-year 2025 policy uses **six years** for technology equipment; Amazon reports **five to six years** for server and networking equipment, having shortened a subset to five years effective January 2025. Neither disclosure supplies an H100, B200, or B300 liquidation appraisal. (www.sec.gov) (www.sec.gov)

The market evidence is equally specific about what it measures. Retrieved **October 15, 2026**, Lambda's eight-GPU table lists **H100 SXM at \$3.99 per GPU-hour** and **B200 SXM6 at \$6.69**, plus applicable tax. Those are public rental tariffs, not the borrower's collected revenue. SSP's page, labeled updated **September 22, 2026**, displays a **\$240,000** refurbished Dell H100 HGX node ask in the United States (US), but marks the relevant rows **unavailable**. That observation is neither an executable bid nor a completed sale. (lambda.ai) (lambda.ai) (sspglobal.ai)

Configuration and operating conditions can matter more than a generation label. NVIDIA documents an eight-GPU DGX B200 system, while Supermicro describes a GB300 NVL72 rack containing **72 B300 GPUs**. Comparable evidence must identify the unit sold, supporting components, cooling and power architecture, condition, location, and transferable rights. Software support also has its own clock: CUDA compatibility depends on driver requirements, while operating-system coverage follows separately documented lifecycles. (docs.nvidia.com) (www.supermicro.com) (docs.nvidia.com) (ubuntu.com)

The defensible underwriting approach is therefore to reconcile several independent worksheets. Establish the accounting policy, verify the exact installed collateral, model collected workload income after operating costs, obtain an appraisal on the lender's accepted basis, and compare scheduled debt with explicitly assumed or evidenced **net recoveries**. General asset-based lending guidance supports independent appraisal and continuing collateral assessment; it does not provide a GPU-specific residual floor. Where executable bids, completed trades, borrower utilization, or support-transfer terms are absent, the answer should remain

unknown. The illustrative arithmetic below demonstrates how a functioning server with positive book value can still leave a debt-recovery shortfall. It is a sensitivity exercise, not an appraisal or a recommended lending standard. (www.occ.gov) (ivsc.org)

Introduction and Background

For an equipment lessor or private-credit investment committee, “How long does a GPU last?” is usually shorthand for several decisions: how much principal can remain outstanding, when amortization should accelerate, and whether a terminal sale assumption can support the proposed tenor. A depreciation schedule supplies context, but the recovery decision requires its own evidence. CoreWeave's carrying-value policy, for example, measures historical cost less accumulated depreciation and amortization. That is a different measurement from a purchaser's offer for equipment at a particular location and date. (www.sec.gov)

This GPU Smith report adopts an **independent engineering perspective** on that credit question. It examines what can be verified from public filings, original commercial pages, specifications, standards, and official guidance, then identifies the asset-level information a financing decision still requires. It does not adopt a reseller depreciation curve as an appraisal. The public glossary of the International Valuation Standards Council (IVSC) distinguishes comparison with comparable assets from an income approach based on future cash flow, reinforcing the need to state which question a valuation is answering. (ivsc.org)

The analysis is anchored to **October 15, 2026, Coordinated Universal Time (UTC)**. Observations describe what the retrieved sources actually disclose. **Assumptions** identify inputs selected for illustration. **Illustrative arithmetic** shows reproducible calculations. **Hypotheses** describe possible future workload uses that require testing. Keeping those labels visible prevents an attractive rental price, a broad accounting category, or a newer product announcement from silently becoming a liquidation forecast.

The evidence hierarchy is practical: a dated executable bid for the exact equipment answers a different question from a public ask; a completed comparable sale needs configuration and transaction adjustments; and a rental index needs methodology and service-term matching. Silicon Data explicitly separates provider segments and standardizes observations for several commercial and technical characteristics. Its rental observations should therefore inform rental analysis, rather than being repurposed as a universal hardware resale curve. (www.silicondata.com) (www.silicondata.com)

Key Changes: Four Clocks Behind GPU Useful Life

Accounting life changes the expense schedule

Accounting useful life is an estimate used to allocate depreciable cost. It is not an assertion that an asset will sell for its remaining carrying amount. The September 2026 update from the International Financial Reporting Standards (IFRS) Foundation quotes the existing International Accounting Standard (IAS) 16 residual-value definition by reference to disposal costs and the age and condition expected at useful-life end. That definition should not be confused with a current bid for a younger machine. The update's discussion is tentative; the definition quoted there is the relevant point here. (www.ifrs.org)

Table 1 summarizes dated policy observations, with fiscal year abbreviated FY. The categories in the second column are deliberately retained because they limit what the numbers can establish.

Company and evidence date	Disclosed asset category	Life and effective date	Reason or comparison limit
CoreWeave, FY2025	Technology equipment	Six years; computing-equipment estimate changed from five to six years January 1, 2023	Filing cites hardware performance, software optimization, and data-center design improvements. No named-GPU resale value. (www.sec.gov)
Amazon, FY2025	Server and networking equipment	Five to six years; subset reduced from six to five years January 1, 2025	Faster technology development, particularly artificial intelligence and machine learning. Subset is not an H100/B200/B300 appraisal. (www.sec.gov)
Alphabet, FY2025	Servers and network equipment	Generally six years	Broad equipment category; planned use and technological obsolescence remain estimation considerations. (d18rn0p25nwr6d.cloudfront.net) (d18rn0p25nwr6d.cloudfront.net)
Meta, FY2025	Most servers and network assets	5.5 years effective January 1, 2025	The policy is a company-specific estimate for a broad asset group. (www.sec.gov) (investor.atmeta.com)

These observations establish **different management estimates**, not a consensus recovery curve. A mean of the disclosed lives would combine different fleets, workloads, operating policies, and asset categories. It would add arithmetic without establishing a market price. Microsoft's separately dated FY2025 report illustrates the category problem by giving computer equipment a range of two to six years; it is not presented here as Microsoft's latest FY2026 policy. (www.microsoft.com)

The interim check also needs careful wording. CoreWeave's June 2026 quarterly filing reports no material changes to significant accounting policies relative to its annual filing. Amazon's June 2026 filing says useful lives remain under ongoing review; that is not the same as a newly disclosed GPU-specific estimate. An underwriting memo should record both the annual policy and what the subsequent filing actually says. (www.sec.gov) (www.sec.gov)

Serviceability, earning life, and recovery have separate endpoints

Physical serviceability asks whether the actual system can operate reliably with available parts and support. A warranty term alone does not answer that question. Lenovo identifies a three-year warranty for its SR780a V3, but the term is neither a demonstrated failure distribution nor an expiry date for every component. Condition testing and operating records remain necessary. (lenovopress.lenovo.com)

Profitable workload life asks whether collected revenue exceeds the avoidable and committed costs of serving a workload. **Recovery life** asks whether an eligible buyer will pay enough after disposal costs and selling time. These can diverge: a supported server may lose a particular customer workload, while an older

system may retain a different use. The latter is a hypothesis to test against workload requirements, not an automatic second-life credit. MLCommons defines inference benchmarks by dataset and quality target, which is why a generic throughput claim cannot settle workload suitability. (mlcommons.org)

Implementation Considerations and Process Changes

Separate generation age from deployment age

NVIDIA announced Hopper and H100 in **March 2022**. That dates the product generation, not the commissioning of a particular server offered as collateral. The original announcement also distinguished SXM and Peripheral Component Interconnect Express (PCIe) form factors. A machine installed substantially later can have a different service history while sharing the same architecture name. (nvidianews.nvidia.com)

The credit file should maintain separate dates for:

- **Generation introduction:** the architecture or product announcement used for market context.
- **Manufacture and shipment:** documented from original-equipment-manufacturer records where available.
- **Delivery and acceptance:** evidence that the specified system reached the required site.
- **Commissioning:** when the system was ready for its intended operational use.
- **Revenue commencement:** when a paying customer's contractual billing began.
- **Inspection and quotation:** when health evidence and market evidence were last refreshed.

These dates should be reconciled rather than inferred from one another. Alphabet's investor FAQ says depreciation begins when an asset is ready for intended use. CoreWeave's shareholder letter separately explains that committed contracts can start on a fixed date or when capacity is delivered. Neither statement makes architecture-announcement age equivalent to operating age or revenue age.

(alphabet2025ir.q4web.com) (coreweave.com)

Define the asset being compared

H100 alone is an incomplete collateral description. NVIDIA's specifications distinguish H100 SXM and H100 NVL; Dell's XE9680 documentation identifies an eight-GPU HGX H100/H800 configuration with 80 gigabytes (GB), 700 watts (W) SXM5 GPUs. Those are useful identification facts, not permission to value a bare accelerator, a module, and a complete server interchangeably. (www.nvidia.com) (www.dell.com)

The same discipline applies to **B200** and **B300**. NVIDIA documents eight B200 GPUs and 1,440 GB total GPU memory in DGX B200, together with system components beyond the accelerators. Its DGX B300 documentation distinguishes alternating-current/power-distribution-unit and direct-current/busbar versions, with separate power shelves excluded from the busbar system. Missing supporting equipment can change the cost and feasibility of redeployment. (docs.nvidia.com) (docs.nvidia.com) (docs.nvidia.com)

Record at least:

- **Identity:** serial numbers, exact stock-keeping unit, GPU form factor, count, and original system model.

- **Included hardware:** central processing units (CPUs), random-access memory (RAM), storage, network interface cards (NICs), switches, cables, and power shelves.
- **Installation:** rack dimensions, voltage, power distribution, cooling connections, and required facility interfaces.
- **Condition:** operating history, diagnostic results, replaced parts, errors, and evidence of complete systems. Record original equipment manufacturer (OEM) support separately.
- **Rights:** ownership, security interests, support entitlements, software licenses, and transfer conditions.
- **Evidence class:** ask, executable bid, closed trade, appraisal, or index, with date, source, and confidence.

Cooling details deserve particular attention. Supermicro documents both air-cooled and liquid-cooled HGX B200 offerings. Lenovo's SR780a V3 uses separate water loops and specifies a minimum B200 chassis flow of **12 liters per minute**. A buyer able to operate one configuration may need additional infrastructure for another. The appraisal comparison should identify that dependency rather than assume every B200 system has the same buyer pool. (www.supermicro.com) (lenovopress.lenovo.com) (lenovopress.lenovo.com)

Market Evidence and Comparable Collateral

Preserve the evidence class

Table 2 records the public market observations used here; RAM quantities use gibibytes (GiB). “Unknown” is a missing field, not a zero adjustment.

Source and quote date	Class	Observed amount	Form and GPU count	CPU/RAM/NIC inclusion	Condition and geography	Warranty /support	Confidence and permitted use
Lambda, retrieved October 15, 2026	Rental tariff	H100 SXM \$3.99/GPU-hour; B200 SXM6 \$6.69/GPU-hour, plus tax	Eight-GPU instance tier	Table states 208 virtual CPUs; 1,800/2,900 GiB RAM respectively; network bundle requires term review	Cloud service, location and availability require selection	Service contract, not hardware title	Strong evidence of displayed tariff; no collected-income or resale inference. (lambda.ai)
SSP, page updated September 22, 2026	Indicative ask	\$240,000/node, ex-works; relevant rows unavailable	Dell H100 HGX SXM node; count unverified here	Complete CPU/RAM/NIC bill of materials unknown	Refurbished, air-cooled, US	Transfer terms unknown	Context only; no executable offer or completed trade. (sspglobal.ai) (sspglobal.ai)
ALTA, retrieved October 15, 2026	Reseller channel page	Executed price unavailable	Used H100 PCIe and SXM advertised; quoted lot unknown	Depends on actual quotation	Seller describes refurbished equipment	Seller's ALTA warranty; OEM transfer unverified	Establishes channel description, not immediate liquidity. (altatechnologies.com) (altatechnologies.com)

Source and quote date	Class	Observed amount	Form and GPU count	CPU/RAM/NIC inclusion	Condition and geography	Warranty /support	Confidence and permitted use
Silicon Data, December 21, 2025 history	Rental-history commentary	No single blended value adopted	Provider segments differ	Rental bundles require normalization	Historical provider observations	Service terms differ	Segmentation evidence, not a hardware liquidation index. (www.silicondata.com)

A quotation should be refreshed when its validity expires or its stated conditions change. Record whether a seller confirms availability and whether the buyer can inspect the actual lot. A displayed page date cannot substitute for those confirmations.

The table is intentionally unsuitable for averaging. It combines service prices, an unavailable equipment ask, a channel description, and historical rental commentary. A **closed-trade register** should instead record sale date, consideration, exact lot, included components, geography, condition, warranty, marketing period, payment certainty, and disposal expenses. Where these fields are unobserved, the underwriting worksheet should keep them blank.

Cash Generation and GPU Economic Useful Life

Model collected revenue before debt service

GPU rental income over time requires a collection-based model:

Monthly operating cash contribution = available GPU-hours × paid utilization × net realized rate – actual operating costs.

Paid utilization means billable or collected GPU-hours divided by defined available GPU-hours. Processor activity alone does not establish it. FinOps Foundation guidance identifies utilization, queues, and health as management inputs; billing and collection records must connect activity to cash. (www.finops.org)

The revenue bridge should reconcile:

- **Capacity:** installed, healthy, available, reserved, and actually delivered GPU-hours.
- **Billing:** contracted rates, discounts, credits, minimum commitments, and interrupted service.
- **Collections:** invoiced amounts, cash receipts, timing, and counterparty exposure.
- **Energy:** metered load, tariff structure, demand charges, and cooling allocation.
- **Operating costs:** hosting, network connectivity, maintenance, spares, staff, and software.
- **Reinvestment:** replacements and compatibility work required to retain the workload.

Amazon Web Services (AWS) documents a best-effort two-minute notice before stopping or terminating a Spot Instance, with different hibernation treatment. An interruptible rate cannot be substituted for a dedicated committed rate without matching availability and service obligations. (docs.aws.amazon.com)

A **kilowatt (kW)** measures power; a **kilowatt-hour (kWh)** measures energy. The Department of Energy defines power usage effectiveness as annual facility energy divided by information-technology equipment energy. Applying another overhead multiplier to an already allocated total facility load double counts overhead. (www.eia.gov) (www.energy.gov)

Illustrative operating sensitivity

Assumptions, not observed performance: one eight-GPU node, a 730-hour month, no outage hours, 60% paid utilization, and a net realized rate of \$2.50/GPU-hour. Assume allocated total facility load averages 10 kW, energy costs \$0.10/kWh, and other monthly operating costs are \$3,000. These independently selected inputs are not established by tariffs or specifications.

Illustrative arithmetic: available capacity is $8 \times 730 = 5,840$ GPU-hours. Paid usage is 3,504 GPU-hours and revenue is **\$8,760**. Energy is 7,300 kWh costing \$730; subtracting energy and other operating costs leaves **\$5,030** before financing, taxes, and replacement capital expenditure. At 35% paid utilization and \$1.75 realized rate, revenue becomes **\$3,577** and contribution becomes **negative \$153**, holding those costs unchanged.

Calculated break-even paid utilization is **25.5%** at \$2.50, or **36.5%** at \$1.75. This excludes financing, replacement capital expenditure, tax, and omitted costs. Replace simplified inputs with actual idle power, cooling, charges, and service credits before underwriting.

NVIDIA lists **14.3 kW maximum system power** for DGX B200, not this example's assumed consumption. MLCommons warns that benchmark power applies to its accompanying benchmark. Use metering or defined workload measurements; the [GPU power calculator](#) is an input worksheet. (docs.nvidia.com) (mlcommons.org)

Treat software replacement as an economic cost

CUDA minor-version compatibility has minimum driver requirements and documented limitations. Ubuntu long-term-support releases provide **five years** of standard security maintenance, while Red Hat identifies component-specific exceptions within its broader enterprise Linux lifecycle. These policies govern software coverage rather than silicon lifespan. (docs.nvidia.com) (ubuntu.com) (access.redhat.com)

A workload-retention hypothesis should specify the model, precision, memory capacity, latency, throughput, quality target, and migration effort. MLCommons' submission guidance makes scenario requirements dependent on system type and benchmark model. FinOps guidance also includes labor and service disruption in modernization costs. Test whether the workload remains profitable after those costs. (docs.mlcommons.org) (www.finops.org)

GPU Collateral Recovery Value: From Gross Price to Cash

State the valuation basis before choosing a haircut

GPU residual value is ambiguous unless the memo states the valuation date, asset condition, sale conditions, and deductions. A terminal accounting residual, fair value, value in use, orderly liquidation estimate, and rapid-sale recovery are different concepts. IAS 36's recoverable amount uses the higher of fair value less costs to

sell and value in use. That accounting test should not be represented as a secured lender's expected cash recovery. (www.ifrs.org)

CoreWeave's separate US filing describes testing asset or asset-group recoverability with undiscounted future cash flows when relevant indicators arise. Combining that approach with IAS 36 terminology would obscure the applicable framework. Neither test supplies a unit-level bid for the pledged equipment. The article therefore uses **net recovery** only for the explicitly defined disposal calculation below. (www.sec.gov)

General Office of the Comptroller of the Currency (OCC) asset-based lending guidance recommends appraisal during underwriting and updated appraisals through the loan's life. Its net orderly liquidation discussion assumes an orderly sale over a reasonable period and an independent competent valuation party. The public IVSC glossary separately explains how forced-sale conditions constrain marketing time. These sources support documenting sale assumptions; they do not establish a standard GPU haircut or mandatory amortization period. (www.occ.gov) (ivsc.org)

Build an itemized recovery waterfall

Net recovery = evidenced or assumed gross sale proceeds – removal – testing/refurbishment – freight/insurance – selling costs – relevant carrying and transfer costs.

Each deduction should have a scope, source, currency, tax treatment, payer, and timing. SSP's ex-works language explicitly excludes freight, importer-of-record responsibilities, duties, and local charges. Even a valid gross equipment quotation would therefore need a separate bridge to cash available to repay debt. (sspglobal.ai)

The recovery plan should identify:

- **Removal:** de-racking, disconnecting liquid loops, lifting, specialist labor, and site access.
- **Preparation:** diagnostics, refurbishment, packaging, missing parts, and documented configuration.
- **Data handling:** storage sanitization, records, and verification appropriate to the equipment.
- **Logistics:** freight, insurance, customs, destination charges, and exchange-rate exposure.
- **Marketing:** broker commissions, inspection conditions, buyer acceptance, and settlement timing.
- **Delay:** storage, insurance, security, and other carrying costs during the sale period.

The media-sanitization guidance of the National Institute of Standards and Technology (NIST) provides a framework for making access to target data infeasible at a defined effort. It does not mean every GPU requires the same procedure as a storage device. Inventory the node's actual media, identify retained data, and cost the appropriate process and records rather than adding an unexplained "sanitization haircut." (csrc.nist.gov)

Transferability is a separate evidence item. Cornell's reproduction of Uniform Commercial Code (UCC) model text addresses commercially reasonable disposition and transfer of the debtor's rights. It does not establish the governing state's law, the lender's perfected interest, access to leased premises, or transfer of an OEM support contract. Transaction counsel should assess those matters; the engineer and appraiser should identify dependencies that affect operability and value. (www.law.cornell.edu) (www.law.cornell.edu)

Guidance from the Bureau of Industry and Security (BIS) requires analysis of classification, destination, end user, and end use for relevant exports, reexports, and transfers. A foreign buyer's apparent price should therefore remain conditional until the contemplated transaction is reviewed. No H100, B200, or B300 classification, buyer eligibility, or license availability is established by this report. (www.bis.gov)

Data Analysis and Evidence

Book value is a separate calculation

GPU depreciation versus resale value compares an accounting estimate with a dated market outcome, without a fixed conversion percentage.

Illustrative accounting assumption: depreciable cost \$300,000, accounting residual zero, straight-line depreciation, and no impairment or additional capitalization. For elapsed life within the assumed useful life:

Book value = cost – (cost – accounting residual) × elapsed life ÷ estimated useful life.

After two years, a six-year schedule leaves **\$200,000**; a five-year schedule leaves **\$180,000**. The **\$20,000** difference changes book value without creating a buyer. Zero accounting residual is an input in this illustration, not a forecast of zero economic value. CoreWeave's historical-cost policy is evidence of the accounting measurement context, not evidence validating either hypothetical schedule for these assets. (www.sec.gov)

The same distinction prevents a misleading rental calculation. Multiplying Lambda's listed H100 rate by every calendar hour gives a **tariff-times-capacity ceiling under a chosen availability assumption**, not observed income. The listed tier, taxes, customer discounts, downtime, utilization, and service obligations must first be reconciled with actual operating records. A public tariff tests rate plausibility, not collections. (lambda.ai)

Compare scheduled debt with dated downside recovery

Hypothetical Example: debt and recovery worksheet. Assume the same \$300,000 original cost, **\$180,000 initial principal**, no balloon, and equal monthly principal amortization over **36 months**. Exclude accrued interest and enforcement expenses from the debt column. Gross sale assumptions and disposal-cost assumptions below are deliberately selected scenarios, not H100/B200/B300 observations, appraisal conclusions, or forecasts.

Table 3 compares debt with a downside disposal scenario. All amounts are US dollars and are **illustrative assumptions or arithmetic**.

Elapsed time	Six-year book value	Scheduled principal	Assumed gross sale	Assumed disposal costs	Calculated net recovery	Net recovery less principal
12 months	\$250,000	\$120,000	\$100,000	\$15,000	\$85,000	-\$35,000
24 months	\$200,000	\$60,000	\$55,000	\$12,000	\$43,000	-\$17,000
36 months	\$150,000	\$0	\$25,000	\$9,000	\$16,000	\$16,000

The assumed schedule leaves a recovery shortfall at the first two dates despite substantial book value. This does not estimate a loss probability. If the first-date gross sale were instead assumed to be **\$150,000**, the same \$15,000 costs would leave \$135,000, or \$15,000 above principal. A committee should see both the inputs and the resulting headroom.

Do not turn the illustration's gross sale values into a labeled "GPU depreciation curve." For actual collateral, leave unsupported recovery cells blank, obtain dated bids or appropriate appraisal evidence, and show separate orderly-sale and constrained-time assumptions. IVSC's comparable-asset principle makes configuration matching central to the market approach; OCC guidance makes continued collateral assessment part of the lending process. (ivsc.org) (www.occ.gov)

Evidence quality matters more than the number of observations

Silicon Data's historical commentary combines internal tracking and public-market pricing, with marketplace observations beginning **June 2024** and neocloud observations beginning **July 2024**. Different starting points and segments prevent treating a blended series as one product. The retrieved commentary also does not provide a complete reproducible sample size and weighting method. This report uses the segmentation insight without deriving a hardware liquidation coefficient. (www.silicondata.com)

Evidence should be ranked by the question it can answer. Public specifications establish what a system requires. Filings establish stated accounting policies. Public tariffs establish displayed service asks. A reseller listing establishes an advertised equipment ask under stated conditions. None replaces the borrower data and buyer evidence missing from this run. The appropriate response to that gap is a conditional underwriting decision, not a universal residual floor.

Implications and Future Directions

What a defensible committee package should contain

A useful approval package connects engineering records, operating cash flow, valuation assumptions, and the debt schedule. Its purpose is to make the unsupported portions visible before approval and provide a way to refresh them afterward. General asset-based lending guidance supports updated collateral appraisals during the loan's life; the frequency and triggers should be agreed for the actual exposure rather than borrowed from an unrelated equipment market. (www.occ.gov)

The package should include:

- **An asset register:** serial-level identity reconciled to invoices, delivery records, inspection, and the security schedule.
- **A policy bridge:** accounting category, life, residual, impairment framework, and filing dates, explicitly separated from appraisal conclusions.
- **An income bridge:** capacity, paid utilization, realized rate, collections, energy, support, and replacement expenditure.
- **A valuation mandate:** accepted basis, valuation date, sale period, configuration, eligible market, and evidence limitations.

- **A recovery bridge:** gross proceeds through every deduction to cash, including time and transfer dependencies.
- **A debt comparison:** outstanding obligations at each review date, downside assumptions, headroom, and escalation triggers.

Public financing announcements reinforce the need to inspect the whole security package. CoreWeave's **March 31, 2026** announcement describes an **\$8.5 billion** facility secured by high-performance-computing infrastructure and an associated customer contract. Blackstone separately describes contract-backed GPU financing. These are neutral examples of contract support alongside equipment; they do not disclose a reusable asset-level recovery methodology for another lender. (investors.coreweave.com) (www.blackstone.com)

Monitor changes in economics and portability

A useful monitoring hypothesis is that workload retention, support availability, buyer eligibility, or redeployment cost may change before hardware stops functioning. Microsoft's January 2026 earnings-call commentary about using software to run newer models on an aging fleet illustrates one operator's approach; it is not a guarantee that an unrelated owner's system will retain the same economics. Borrower-specific evidence remains decisive. (www.microsoft.com)

Monitoring should therefore refresh the variables that can change a credit decision: collected cash per available GPU-hour, required software work, system health, transferable support, configuration-comparable bids, disposal costs, and legal transfer conditions. A later product generation may justify retesting these inputs, but does not by itself establish impairment, physical failure, or a particular resale decline.

No universal GPU obsolescence date follows from the reviewed sources. The practical endpoint is conditional: the equipment ceases to support the modeled credit case when verified operating cash flow, supported workload options, or appropriately evidenced net recovery no longer covers the obligations and assumptions that justified it. That endpoint must be reassessed with the lender, appraiser, accountant, counsel, and relevant technical specialists.

Conclusion

GPU useful life becomes useful for underwriting only when the four clocks remain distinct. Accounting depreciation allocates cost under a stated policy. Physical serviceability depends on the actual installed system. Profitable workload life depends on collected income after the costs required to deliver and retain that workload. Collateral recovery depends on a feasible sale, eligible buyers, transferable rights, disposal costs, and time.

The public evidence reviewed here can anchor a disciplined investigation. It cannot establish H100, B200, or B300-specific recovery floors, executable sale proceeds, or a universally defensible loan tenor. The most consequential unknowns are transaction-level: exact serials and configurations, borrower utilization and collections, support transfer, current buyer bids, completed comparable sales, and an appraisal basis the lender accepts.

The committee's decision should therefore show its dependencies explicitly. Keep unsupported recovery values blank or label them as assumptions. Show the debt balance at each date, disclose the recovery waterfall, test adverse operating inputs without assigning invented probabilities, and explain what new evidence would change the decision. A functioning asset and a positive book balance can coexist with insufficient cash recovery.

For independent pre-close technical due diligence, use [GPU Smith's financing due-diligence page](#) to frame the equipment, infrastructure, and operating-evidence review.

References

All sources below were retrieved October 15, 2026 UTC. Filing periods and commercial-page dates are identified in the body; a retrieval date does not establish an executed transaction.

- **Accounting policies:** CoreWeave annual and interim filings (www.sec.gov) (www.sec.gov); Amazon annual and interim filings (www.sec.gov) (www.sec.gov); Alphabet annual filing (d18rn0p25nwr6d.cloudfront.net); Meta policy and announcement (www.sec.gov) (investor.atmeta.com); Microsoft's explicitly historical FY2025 comparison (www.microsoft.com).
- **Deployment and software context:** Alphabet's dated investor FAQ (alphabet2025ir.q4web.com); CoreWeave shareholder letter (coreweave.com); Microsoft's operator commentary (www.microsoft.com); NVIDIA architecture announcement (nvidianews.nvidia.com).
- **Hardware identification:** NVIDIA H100 specifications (www.nvidia.com); Dell XE9680 specifications (www.dell.com); Lenovo SR780a V3 guide (lenovopress.lenovo.com); NVIDIA DGX B200 and B300 guides (docs.nvidia.com) (docs.nvidia.com); Supermicro accelerator configurations (www.supermicro.com).
- **Software and workloads:** NVIDIA CUDA compatibility (docs.nvidia.com); Ubuntu lifecycle (ubuntu.com); Red Hat lifecycle (access.redhat.com); MLCommons benchmark and submission guidance (mlcommons.org) (docs.mlcommons.org).
- **Market observations:** Lambda instance tariffs (lambda.ai); SSP indicative inventory (sspglobal.ai); ALTA used-H100 page (altatechnologies.com); Silicon Data history and index methodology (www.silicondata.com) (www.silicondata.com).
- **Operating economics:** FinOps Foundation AI and modernization guidance (www.finops.org) (www.finops.org); AWS Spot notice terms (docs.aws.amazon.com); EIA electricity units (www.eia.gov); Department of Energy facility-efficiency guidance (www.energy.gov).
- **Valuation and recovery:** OCC asset-based lending handbook (www.occ.gov); dated IVSC glossary (ivsc.org); IFRS impairment summary and residual-value definition (www.ifrs.org) (www.ifrs.org).
- **Disposal and transfer:** Cornell's UCC model text (www.law.cornell.edu) (www.law.cornell.edu); BIS licensing guidance (www.bis.gov); NIST media-sanitization guidance (csrc.nist.gov).
- **Financing context:** CoreWeave's financing announcement (investors.coreweave.com); Blackstone's own financing discussion (www.blackstone.com).

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. <https://www.sec.gov/Archives/edgar/data/1769628/000176962826000104/crww-20251231.htm>
2. <https://www.sec.gov/Archives/edgar/data/1018724/000101872426000004/amzn-20251231.htm>
3. <https://lambda.ai/instances>
4. <https://sspglobal.ai/live>
5. <https://docs.nvidia.com/dgx/dgxb200-user-guide/introduction-to-dgxb200.html>
6. <https://www.supermicro.com/en/accelerators/nvidia/?mlg=0>
7. <https://docs.nvidia.com/deploy/cuda-compatibility/minor-version-compatibility.html>
8. <https://ubuntu.com/about/release-cycle>
9. <https://www.occ.gov/publications-and-resources/publications/comptrollers-handbook/files/asset-based-lending/pub-ch-asset-based-lending.pdf>
10. <https://ivsc.org/standards-glossary/>
11. <https://www.silicondata.com/products/silicon-index/h100>
12. <https://www.ifrs.org/news-and-events/updates/ifric/2026/ifric-update-september-2026/>
13. <https://d18rn0p25nwr6d.cloudfront.net/CIK-0001652044/e6a3f2b6-bb98-4f54-900f-2a8794a582cf.html>
14. <https://www.sec.gov/Archives/edgar/data/1326801/000162828026003942/R10.htm>
15. <https://investor.atmeta.com/investor-news/press-release-details/2025/Meta-Reports-Fourth-Quarter-and-Full-Year-2024-Results/>
16. <https://www.microsoft.com/investor/reports/ar25/index.html>
17. <https://www.sec.gov/Archives/edgar/data/1769628/000176962826000366/crww-20260630.htm>
18. <https://www.sec.gov/Archives/edgar/data/1018724/000101872426000026/amzn-20260630.htm>
19. <https://lenovopress.lenovo.com/lp1980-thinksystem-sr780a-v3-server>
20. <https://mlcommons.org/benchmarks/inference-datacenter/>
21. <https://nvidianews.nvidia.com/news/nvidia-announces-hopper-architecture-the-next-generation-of-accelerated-computing>
22. <https://alphabet2025ir.q4web.com/investor/faqs-and-general-information/default.aspx>
23. <https://coreweave.com/blog/coreweave-2025-annual-report-shareholder-letter>
24. <https://www.nvidia.com/en-us/data-center/h100/>
25. https://www.dell.com/support/manuals/en-jm/poweredge-xe9680/xe9680_ism_pub/gpu-specifications?guid=guid-1f91b9fb-b8fd-492e-863f-45febc9355c6&lang=en-us
26. <https://docs.nvidia.com/dgx/dgxb300-user-guide/introduction-to-dgxb300.html>
27. <https://altatechnologies.com/collections/used-nvidia-h100>
28. <https://www.silicondata.com/blog/h100-rental-price-over-time>
29. <https://www.finops.org/wg/finops-for-ai-tools-services-considerations/>
30. <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/spot-instance-termination-notices.html>
31. <https://www.eia.gov/energyexplained/electricity/measuring-electricity.php>
32. <https://www.energy.gov/cmei/femp/cooling-water-efficiency-opportunities-federal-data-centers>
33. <https://gpusmith.com/tools/gpu-power-calculator>
34. <https://access.redhat.com/support/policy/updates/errata>
35. <https://docs.mlcommons.org/inference/submission/>

36. <https://www.finops.org/framework/capabilities/usage-optimization/>
37. <https://www.ifrs.org/issued-standards/list-of-standards/ias-36-impairment-of-assets/>
38. <https://csrc.nist.gov/pubs/sp/800/88/r2/final>
39. <https://www.law.cornell.edu/ucc/9/9-610>
40. <https://www.law.cornell.edu/ucc/9/9-617>
41. <https://www.bis.gov/licensing>
42. <https://investors.coreweave.com/news/news-details/2026/CoreWeave-Closes-Landmark-8-5-Billion-Financing-Facility-Achieving-First-Investment-Grade-Rated-GPU-backed-Financing/default.aspx>
43. <https://www.blackstone.com/insights/article/revisiting-our-megatrends/>
44. <https://www.microsoft.com/en-us/investor/events/fy-2026/earnings-fy-2026-q2>
45. <https://gpusmith.com/gpu-financing-due-diligence>

THE PUBLISHER

About GPU Smith

GPU Smith provides independent engineering and research for private AI infrastructure. We help technical buyers and operators reason about GPU workload sizing, hardware selection, procurement, deployment and inference operations, from the compute system to the networking, power and cooling requirements around it.

Start with the workload

Useful infrastructure decisions begin with the models, data, throughput, latency and operating constraints a team actually has. GPU Smith helps connect those requirements with system design and deployment choices. The goal is a privately controlled AI environment whose capacity and operational demands are understood before a purchasing decision.

Hardware and supplier research

Our public reference library covers GPUs, complete systems and networking components. The server-vendor directory and manufacturing research help teams investigate suppliers, compare options and follow sources behind technical claims. We distinguish manufacturer specifications from measured benchmarks and advertised capabilities from tested configurations. Reference pages are research resources, not a live inventory listing or a binding equipment quote.

Deployment and operations

GPU Smith's engineering scope includes infrastructure integration, networking, power, cooling and the practical operation of inference workloads. Our published articles explain the tradeoffs behind deployment, procurement and ongoing operations so teams can ask better questions and document decisions.

Work with GPU Smith

Explore the [hardware library](#), [GPU references](#), [networking references](#) and [vendor research](#). [Contact GPU Smith](#) with your workloads and deployment constraints to discuss a project and confirm current scope, pricing and availability.

Inclusion of a manufacturer or product in our research does not imply a partnership, certification or endorsement.

Website: <https://gpusmith.com>

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.