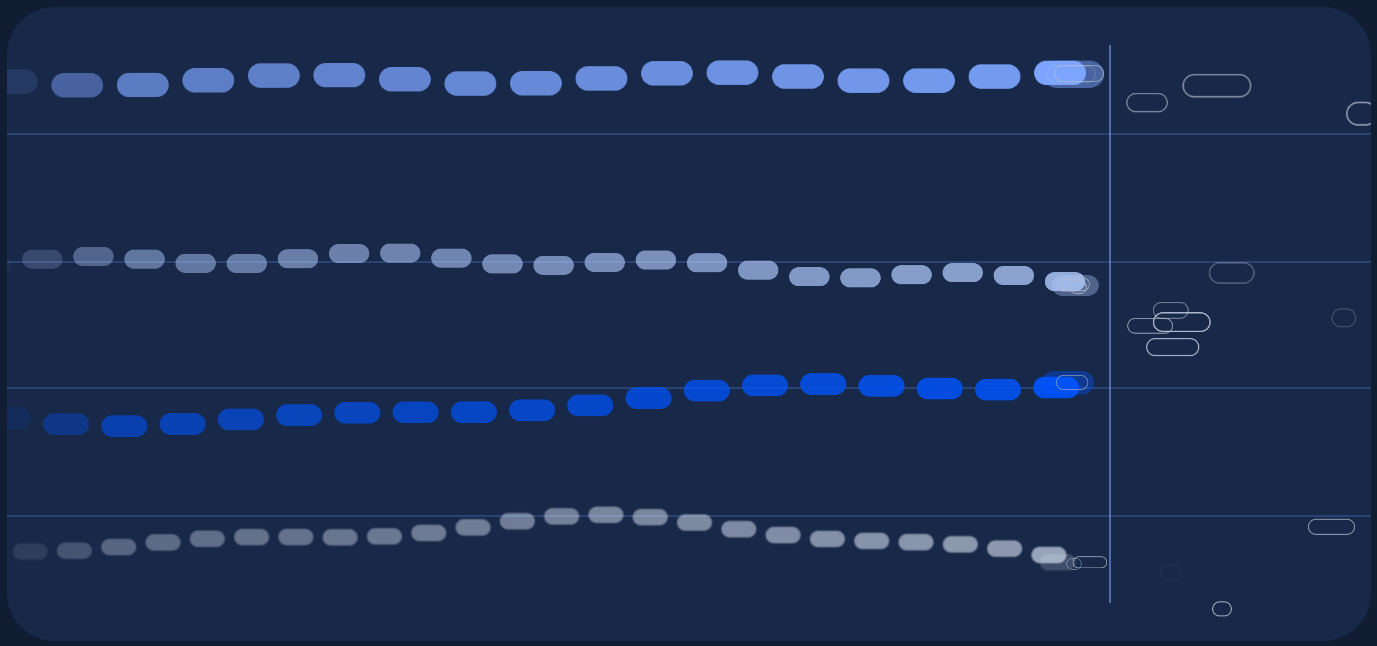
Rillor

RESEARCH REPORT

GPU depreciation and useful life in filings and market evidence

Filings put GPU and server useful lives at up to six years. This report compares those book lives with tax rules, failure studies, rents for older GPUs and resale evidence.



Rillor

Contents

01 Summary

02 Background

03 Method

04 Findings

05 Implications

06 Limits of this analysis

07 Sources

Summary

A useful life in an annual report is a management estimate, and how long a data-center GPU keeps working and earning shows up in other records: tax rules, failure studies, rents, contracts and resale disclosures.

- In filings made from 2023 to 2026, the nine companies reviewed here state lives of up to six years for servers, GPUs or technology equipment: six at Alphabet and CoreWeave, five to six at Amazon and Nscale, five to 5.5 at Meta, five at Nebius and IREN, two to six at Microsoft and 1 to 6 at Oracle [2, 4, 5, 7, 8, 9, 11, 12, 13].
- Alphabet, Amazon, Meta, Microsoft and Oracle report servers with network or other equipment as one asset class, with no separate life for GPUs. Nscale, IREN and Nebius tie a stated life to GPUs by name, and CoreWeave gives six years for technology equipment [2, 4, 5, 7, 8, 9, 10, 12, 13].
- Six of the seven useful-life changes in the 2023 to 2026 filings reviewed here lengthened lives. Amazon shortened a subset of servers and networking equipment from six years to five from 1 January 2025, citing an increased pace of technology development [1, 3, 4, 5, 6, 9, 10, 11].
- A change in useful life is a change in accounting estimate, applied from the change onward without restating earlier years. Each filer that changed a life disclosed one year's effect, from $20 million at CoreWeave to $3.9 billion at Alphabet, both in 2023 [1, 9, 16, 17].
- For U.S. federal tax, IRS Publication 946 gives computers and peripheral equipment a 5-year recovery period, and a 100% special allowance applies to certain qualified property acquired and placed in service after 19 January 2025 [18].
- Failure studies count interruptions and replacements, not the end of working life. Two A100 research clusters logged 6.50 and 2.34 failures per thousand node-days, and the Titan supercomputer replaced about 11,000 of its GPU assemblies for a failing board resistor in 2016 and 2017 [19, 22].
- Older GPU classes still rent and re-contract. Lambda lists the V100, announced in 2017, at $0.79 per GPU-hour from 16 November 2026, and CoreWeave said in August 2026 that it had signed an A100 contract extending into 2029 [24, 29, 30].
- No verified resale price turned up in this review. HIVE Digital reported a gain on selling fully depreciated GPUs without a price, and a firm that offers to structure residual value insurance publishes residual levels from proprietary data [33, 34].

Background

Four lives of one asset

Accounting defines useful life as "The period over which an asset is expected to contribute directly or indirectly to future cash flows" [16]. PwC's guide adds that "the useful life could be different than the economic life or actual physical life of the asset", and gives the example of a machine designed to last ten years that a company plans to use for five, whose useful life is five years [16]. This report keeps four lives apart.

- **Book life** is the period over which a company spreads an asset's cost in its financial statements. It is a management estimate, disclosed in the notes to the annual report.
- **Tax life** is the recovery period a tax system assigns. In the United States it comes from the depreciation system described in IRS Publication 946 [18].
- **Physical life** is how long the hardware keeps working, measured in failure and replacement studies.
- **Economic life** is how long customers keep paying for the hardware's capacity. KBRA's June 2026 research on GPU and TPU financings also discusses the distinction between useful life and economic life [31].

Why the estimate carries weight

Alphabet depreciates on a straight-line basis "over the estimated useful lives of the assets" [2], so a longer life spreads the same cost over more years. The amounts are large. Meta's depreciation expense was $18.00 billion in 2025, $13.36 billion of it for servers and network assets [5]. Microsoft's rose from $15.2 billion in fiscal 2024 to $34.3 billion in fiscal 2026 [7]. Nebius reported depreciation and amortization of $259.7 million in the second quarter of 2026, 245% more than a year earlier [11]. IREN's GPU hardware stood at $2,034.8 million at 30 June 2026, against $76.0 million a year earlier [12].

GPU generations now arrive quickly. NVIDIA's annual report for fiscal 2026 describes "a new product and architecture cadence of our Data Center solutions where we seek to complete new computing solutions each year" [23]. Lenders tie terms to the estimate too: Lambda's $926 million term loan of August 2026 is repaid "on a fully amortizing schedule aligned with the contracted cash flows and useful life of the underlying GPU infrastructure" [32].

This report is for lenders, credit analysts and finance teams who need to know what filings assume about GPU lives and what public evidence sits beside those assumptions.

Method

What was examined. The book lives that companies running server or GPU fleets state in filings with the U.S. Securities and Exchange Commission (SEC), the changes they made from 2023 to 2026, and the public evidence on tax, physical and economic life.

Sources. Filings first: annual reports, registration statements and a shareholder letter of nine companies [1-13], a credit agreement filed with Nscale's registration statement [14] and an SEC staff comment letter to CoreWeave [15]. Accounting and tax guidance from PwC, BDO and the IRS [16, 17, 18]. Four studies of GPU fleets [19, 20, 21, 22]. NVIDIA, cloud provider and Lambda documents [23-29], CoreWeave's second-quarter 2026 earnings call [30], KBRA [31] and Lambda's loan announcement [32]. Resale disclosures from HIVE Digital and American Compute [33, 34], and the Rillor Compute Index page [35].

How. Filings were downloaded from EDGAR and searched in full text, and each policy is quoted or restated with its filing date. Filings are management estimates, not prices. Every GPU price figure and residual value level carries one of the six evidence classes the Rillor Compute Index methodology defines: specification (what a system is as its maker publishes it), listing (a public asking price at a point in time), quote (a price offered for a stated configuration and quantity), indicative price (a price level a source states for reference), reported sale (a sale as a buyer, seller or broker reports it) and verified transaction (a completed sale with confirming documents) [35]. Figures marked "derived" are arithmetic on the cited numbers. All sources were viewed on 8 November 2026, and [27] again on 9 November. This report gives no tax or investment advice and no view on any company's earnings or shares.

Findings

1. Book lives in annual reports, 2023 to 2026

The latest filing of each company reviewed states these lives.

Filer	Filing	Asset line	Life	Source
Alphabet	10-K, filed 5 Feb 2026	Servers and network equipment	Generally six years	[2]
Amazon	10-K, filed 6 Feb 2026	Servers and networking equipment	Five to six years	[4]
Meta	10-K, filed 29 Jan 2026	Servers and network assets	Five to 5.5 years	[5]
Microsoft	10-K, filed 29 Jul 2026	Servers and network equipment	Two to six years	[7]
Oracle	10-K, filed 22 Jun 2026	Computer, network, machinery and equipment	1 to 6 years	[8]
CoreWeave	10-K, filed 2 Mar 2026	Technology equipment	6 years	[9]
Nebius	20-F, filed 30 Apr 2026, and Q2 2026 letter	Servers and network equipment	Four years before 2026, five from 2026	[10, 11]
IREN	10-K, filed 27 Aug 2026	HPC hardware	5 years	[12]
Nscale	S-1, filed 18 Sep 2026	Technology equipment	5 to 6 years	[13]

Seven changes of estimate appear in these filings and the earlier ones reviewed.

Filer	Effective	Change	First-year effect as disclosed	Source
Microsoft	Fiscal 2023	Server and network equipment, four to six years	Operating income up $3.7 billion, net income up $3.0 billion, $0.40 a share	[6]
Alphabet	Fiscal 2023	Servers four to six years, certain network equipment five to six	Depreciation down $3.9 billion, net income up $3.0 billion, $0.24 a share	[1]

Filer	Effective	Change	First-year effect as disclosed	Source
CoreWeave	1 January 2023	Computing equipment in data centers, five to six years	Total expenses down $20 million, $0.10 a share	[9]
Amazon	1 January 2024	Servers, five to six years	Depreciation and amortization down $3.2 billion, net income up $2.5 billion, $0.23 a share	[3]
Meta	1 January 2025	Most servers and network assets, to 5.5 years	Depreciation down $2.92 billion, net income up $2.59 billion, $1.00 a diluted share	[5]
Amazon	1 January 2025	A subset of servers and networking equipment, six to five years	Depreciation and amortization up $1.4 billion, net income down $1.0 billion, $0.10 a share	[4]
Nebius	1 January 2026	Servers and network equipment, four to five years	About $167.6 million less depreciation expected for 2026	[10, 11]

Which assets the life covers. Five filers report servers together with network or other equipment and give no life for GPUs alone [2, 4, 5, 7, 8]. Three tie a life to GPUs by name. Nscale's technology equipment, at five to six years, "primarily relates to our GPU and related infrastructure assets", while its computer equipment carries three [13]. IREN's HPC hardware carries five years, and the GPUs pledged under its GPU financing facility are depreciated straight-line over five [12]. Nebius's 20-F says "Our main assets, GPUs, are currently depreciated over four years" [10], and its Q2 2026 letter reports a five-year life for server and network equipment [11]. CoreWeave gives six years for technology equipment, and its risk factors name GPUs among the infrastructure components whose useful life it estimates [9].

In January 2025 the SEC staff noted that the registration statement CoreWeave first submitted depreciated equipment and software over a range of 3 to 12 years, and asked for a breakdown of asset types and their lives, citing ASC 360-10-50-1(b) [15]. CoreWeave's 10-K for 2025 shows technology equipment at six years on its own line, and data center equipment and leasehold improvements at the shorter of the remaining lease term or up to 12 years [9].

Direction and reasons. Six of the seven changes lengthened lives, and one shortened them. Microsoft cited "investments in software that increased efficiencies in how we operate our server and network equipment, as well as advances in technology" [6]. CoreWeave cited "continuous advancements in hardware performance, software optimization, and data center design improvements" [9]. Nebius cited "updated information and usage patterns obtained subsequent to year-end" [10]. Amazon gave the reason for its 2025

shortening and its early retirements in the same words: "an increased pace of technology development, particularly in the area of artificial intelligence and machine learning" [3]. Technology change is named among the reasons for both directions.

Early retirement. Amazon also decided, mostly in the fourth quarter of 2024, to retire some servers and networking equipment early [3]. It recorded about $920 million of accelerated depreciation and related charges for that quarter, and its February 2025 filing expected the accelerated depreciation to continue into 2025 and lower that year's operating income by about $0.6 billion [3].

How filers set the estimate. Alphabet relies on "historical asset performance, expected technology advancements, and our future infrastructure deployment plans" [2]. Nebius weighs obsolescence, refresh cycles, utilization and retirement patterns [10], and IREN reviews useful lives at each reporting date [12]. Oracle writes that its useful lives "could be shortened should our cloud strategies change" [8], and CoreWeave states that it cannot guarantee its useful-life estimates are accurate [9].

2. What a change in useful-life estimate records, and what it does not

Under U.S. accounting rules, a change in the estimated useful life or salvage value of a long-lived asset is a change in accounting estimate, applied prospectively [16]. BDO's guide adds that such a change is not accounted for by restating earlier financial statements or by reporting pro forma amounts for earlier periods, and the entity discloses its effect on income from continuing operations and net income, with per-share amounts, for the current period [17].

The filings follow that pattern. Each change in Finding 1 comes with one year's effect and a per-share amount [1, 3, 4, 5, 6, 9], and Nebius gave an expected effect for its first year [10]. None states the effect on later years.

A disclosed effect depends on its asset base. Amazon's February 2025 filing estimated that the 2025 shortening would lower 2025 operating income by about $0.7 billion, measured on the affected assets in property and equipment at 31 December 2024 [3]. Its 2025 annual report gave the effect for the year as $1.4 billion of added depreciation and amortization and $1.0 billion less net income, primarily in its AWS segment [4]. Meta measured its 2025 effect on servers and network assets placed in service as of 31 December 2024 [5]. A first-year figure therefore describes the assets it was measured on, and assets added to the class later change the amount.

The arithmetic of a longer life. A $1,000 asset depreciated straight-line with no salvage value costs $250 a year over four years and about $167 a year over six (derived). The total expensed over the asset's life is $1,000 in both cases (derived). A longer life lowers the yearly charge and moves expense into later years.

What it does not record. A useful life can differ from economic or physical life [16], and a change of estimate reports no market value for the hardware. Impairment is reported on its own line: Meta recorded impairment losses on property and equipment of $237 million, $288 million and $738 million for 2025, 2024 and 2023 [5].

3. Tax life under MACRS, kept apart from book life

IRS Publication 946 for 2025 returns describes the Modified Accelerated Cost Recovery System (MACRS). Asset class 00.12, Information Systems, "Includes computers and their peripheral equipment used in administering normal business transactions and the maintenance of business records, their retrieval and analysis", with a class life of 6 years and a recovery period of 5 years under both the General Depreciation System (GDS) and the Alternative Depreciation System (ADS) [18]. The ADS table lists "Computers and peripheral equipment" at 5 years [18]. The publication does not name GPU servers.

GDS 5-year property uses the 200% declining balance method, and the half-year convention treats property "as placed in service or disposed of at the midpoint of the year" [18]. Twice the straight-line rate for five years is 40% a year, and the half-year convention takes half of that, 20%, in the first year (derived).

The publication's "What's New for 2025" says that P.L. 119-21, "commonly known as the One Big Beautiful Bill Act, reinstated the 100% special depreciation allowance for certain qualified property acquired and placed in service after January 19, 2025", and that a taxpayer can elect a 40% allowance instead [18]. The Section 179 limit for tax years beginning in 2026 is $2,560,000, reduced once property placed in service exceeds $4,090,000 [18].

A full first-year tax deduction and a six-year book life can describe the same server. Neither measures how long the hardware works.

4. Physical life: failure rates in large GPU clusters

Research clusters on A100 GPUs. Researchers at Meta analyzed 11 months of data from two research clusters, "with 4 million jobs and over 150 million A100 GPU hours" [19]. RSC-1, a general cluster of 16,000 GPUs, had 6.50 failures per thousand node-days, and RSC-2, an 8,000-GPU cluster for vision work, had 2.34 [19]. RSC-1's GPUs were swapped at about three times RSC-2's rate [19]. Mean time to failure was 7.9 hours for 1,024-GPU jobs against 47.7 days for 8-GPU jobs [19]. These are job failure rates on two clusters, not hardware lifetimes.

A pre-training run on H100 GPUs. The 2024 report The Llama 3 Herd of Models describes a run on up to 16,000 H100 GPUs [20]. In a 54-day snapshot of pre-training it recorded 466 job interruptions, 47 scheduled and 419 unexpected [20]. About 78% of the unexpected interruptions were attributed to confirmed or suspected hardware issues, and GPU issues accounted for 58.7% of them [20]. Effective training time stayed above 90% [20].

A100 and H100 compared. A study of the Delta system, 1,056 A100 and H100 GPUs observed over 2.5 years and 11.7 million GPU hours, found H100 memory resilience worse than A100's, with 3.2 times lower per-GPU mean time between memory errors, and H100 hardware resilience significantly better on critical components [21]. Its authors estimate that "significant overprovisioning of 5% is necessary to handle GPU failures" [21].

One survival analysis. The Titan supercomputer ran for nearly seven years, with most of its power from 18,688 GPUs [22]. Its operators analyzed "over 100,000 years of GPU lifetimes during Titan's 6-year-long productive period" [22]. GPU failures began to rise in mid 2016, and about 11,000 GPU assemblies were replaced from late 2016 through much of 2017, traced to "a resistor on the GPU circuit board (not the GPU chip itself) due to silver sulfide corrosion" [22]. About 11,000 of 18,688 is about 59% (derived). The authors found GPU reliability "dependent on heat dissipation to an extent that strongly correlates with detailed nuances of the cooling architecture and job scheduling" [22]. Titan's GPUs were one machine's parts, so the study shows how lifetimes can be measured, not how long current GPUs last [22].

Failure in loan terms. A credit agreement of 27 August 2026 filed with Nscale's S-1 requires the borrower to hold spare GPU servers equal to set percentages of the contracted VR200 and GB300 servers, with a new level 12, 9 and 6 months before the earliest warranty expiry and another at expiry [14]. From 12 months before that date, the borrower calculates a trailing 12-month "GPU Failure Rate" [14]. Above a threshold it reserves cash for spares or obtains additional warranty coverage, and after expiry it keeps spares of at least the greater of that rate and a floor [14]. Every spares and failure-rate percentage is redacted [14]. Oracle's risk factors make the same link, warning of "unplanned repair, replacement or refreshment costs, including in some cases outside applicable warranty coverage" [8].

5. Economic life: published rents for older GPU classes

Generations and support. NVIDIA announced the Tesla V100, its first Volta-based data center GPU, when it launched Volta on 10 May 2017 [24]. Microsoft's documentation gives 30 September 2025 as the retirement date of its NCv3 series, which runs on Tesla V100 GPUs [26], about eight years and four months later (derived). Google Cloud ended support for K80 GPUs on 1 May 2024, after NVIDIA announced their end of support in June 2022 [27]. In CUDA 13.0, the cuFFT library "Removed support for Maxwell, Pascal, and Volta GPUs", and cuSPARSE dropped support for pre-Turing architectures [25]. Software support can end while the hardware still runs.

Older classes on a current price list. Lambda's pricing page lists these prices per GPU-hour for 8-GPU instances, prices that "take effect November 16, 2026, at 00:00 UTC" [29]. They are posted prices (Listing), not prices paid.

GPU class, 8x instance	Price per GPU-hour [29]	Ratio to HGX H100 (derived)
NVIDIA HGX B200	$7.69	1.68
NVIDIA HGX H100	$4.59	1.00
NVIDIA HGX A100 80 GB	$3.19	0.69
NVIDIA HGX A100 40 GB	$1.99	0.43

GPU class, 8x instance	Price per GPU-hour [29]	Ratio to HGX H100 (derived)
NVIDIA Tesla V100 16 GB	$0.79	0.17

A GPU class announced in 2017 [24] is on a price list for November 2026, at about a sixth of the H100 rate (derived).

Posted price cuts. AWS announced on 5 June 2025 reductions to On-Demand prices from 1 June 2025, and to Savings Plan purchases after 4 June 2025, of up to 45% for P5, up to 26% for P5en and up to 33% for P4d and P4de instances running Amazon Linux [28]. These are posted prices (Listing).

6. Re-contracting and multi-year terms for older GPUs

Contract length. CoreWeave's committed contracts had a weighted-average duration of about five years at 31 December 2025, and its remaining performance obligations were $60.7 billion [9]. Its risk factors say that "any inability to redeploy components of our existing infrastructure to extend past their contracted life could significantly affect our business" [9].

What CoreWeave told investors in August 2026. On its second-quarter call on 11 August 2026, CoreWeave's management said "pricing for prior generation SKUs is at or above where it was years ago" and that "we remain largely sold out of prior generations of NVIDIA GPUs in addition to the current SKUs" [30]. It described "a typical five year contract", after which "we are free to recontract that cloud infrastructure or offer it to the market" [30]. It said: "we recently signed an A100 contract that extends into 2029 at an attractive price. As a reminder, this SKU was introduced in 2020" [30]. A contract into 2029 for a GPU introduced in 2020 spans about nine years from introduction (derived). The chief financial officer said capacity coming up for renewal is "a very limited part of our fleet", with older-generation average selling prices "higher or at levels that we've seen about a year ago" [30]. These are management statements on a call, without prices or counterparties, so they are reported, not verified.

The opposite risk, in another filing. IREN's annual report for fiscal 2026 warns that "customer demand may concentrate on the newest generations of hardware. As a result, the economic useful life and residual value of our deployed equipment may be shorter or lower than we anticipate, pricing for services delivered on older generations of hardware may decline" [12].

How credit research frames it. KBRA's research of 22 June 2026 places the key credit issue beyond whether GPUs or TPUs remain physically functional after the original contract term, in "whether customers would continue using the capacity at an attractive price" [31]. It covers "workload suitability, same-customer extension potential, third-party demand, chip fungibility, switching costs", "power cost per unit of compute, and refresh strategies", across "contract maturity, extension, replacement customer, and refresh scenarios" [31]. The press release gives no figures.

Loan terms tied to life. Lambda's $926 million senior secured term loan B, announced 27 August 2026, matures on 31 December 2030 and amortizes fully on a schedule aligned with the contracted cash flows and the useful life of the GPU infrastructure [32]. It is secured by the GPU servers and related infrastructure it funds, rated Baa2 by Moody's and priced at SOFR + 3.00% [32]. From announcement to maturity is about four years and four months (derived). Nscale's credit agreement defines a limited parent guarantee "with respect to a 'residual value' guarantee and indemnity", with an RVG Period of 60 days after a tranche termination date [14].

7. Resale evidence by evidence class

Reported sale, without a price. HIVE Digital's management discussion for the period ended 31 December 2024 says it operated 4,000 NVIDIA A-series GPUs, mainly A40s with some A6000 and A5000 cards, and 96 H100 GPUs [33]. It "elected to sell some A-Series GPUs which are not being used for HPC" and recorded a gain on the sale of cards that were fully depreciated [33]. It adds that "The revenue generated from the operation of these GPU cards combined with the proceeds on sale, exceeded the original purchase price" [33]. This is a reported sale by the seller, with no price, quantity or buyer. It shows used GPUs sold above zero book value, and nothing about the price level.

Indicative levels from a residual value firm. American Compute's report of 15 June 2026 states "We have transaction data covering over 600,000 secondary-market units" and that "A GPU's useful life can reach eight years, well beyond the typical financing term" [34]. It says investors and lenders "can also book residuals, even north of 10% of equipment cost over five years", and that "Most lenders agree that over a three-year term a GPU server should fetch more than 5% of original cost. They debate whether the right number is 10% or 60%" [34]. The same page invites readers to discuss "structuring residual value insurance solutions for a transaction", and says its transaction data is the firm's proprietary data [34]. Each figure is the firm's own statement of a level, at most an indicative price, and none is an observed sale.

An analogy from mining hardware. IREN's impairment of its mining hardware used fair value "determined using Level 3 inputs, based on quoted market values and related adjustments for similar assets" [12], the kind of input a fair-value test of used hardware can rest on.

No verified transaction price for used A100, H100 or B200 systems was found for this review.

8. How the Rillor Compute Index methodology treats hardware prices

The Rillor Compute Index page states that RCI Hardware "uses published listings, quotes and reported sales of GPU systems. Each observation is tagged with its evidence class, and the classes stay separate in calculation" [35]. The series measures "GPU system prices by comparable grade" in USD per system, and air-cooled and liquid-cooled builds are separate grades [35]. RCI Rent-to-Value is "Annualized rent over hardware price for the same GPU class, with payback and residual-value context" [35]. The page lists financing diligence among its uses: "Borrowers and lenders review hardware prices and rent-to-value for the systems behind a lease or loan" [35].

Implications

For lenders and credit analysts. A borrower's book life is a management estimate its own filings describe as subject to change [8, 9, 10], and seven changes in four years show it moves [1, 3, 4, 5, 6, 9, 10]. Physical risk appears in covenants as spares and trailing failure rates tied to warranty expiry [14]. KBRA's residual question, whether customers keep using the capacity at an attractive price [31], is answered in this review only by management statements [30] and indicative levels [34].

For finance teams at operators. A change in estimate is disclosed for one year, on a stated asset base [3, 4, 5, 17]. Teams comparing filers can read the asset line as well as the number: five filers report servers with network or other equipment as one class [2, 4, 5, 7, 8], and three state a life for GPUs by name [10, 12, 13], while CoreWeave states one for technology equipment [9]. The tax schedule is separate and front-loaded [18].

For buyers weighing ownership. Older classes still appear on price lists [29] and in new contracts [30], and AWS cut posted prices for older and newer instances in 2025 [28]. Software support ends on its own schedule [25], and one cloud retired its V100 series about eight years after the V100 was announced [24, 26]. A plan built on a six-year life can be tested against rents for the same class, each tagged with its evidence class.

Limits of this analysis

- Alphabet, Amazon, Meta, Microsoft and Oracle state lives for servers with network or other equipment as one class, so this report gives no GPU-specific life for them [2, 4, 5, 7, 8].
- The useful-life policies come from annual reports, registration statements and Nebius's Q2 2026 shareholder letter [1-13]. A change disclosed only in another quarterly report is outside this review.
- Failure studies measure job failures, interruptions, memory errors and replacements on specific clusters [19, 20, 21]. The only survival analysis of GPU lifetimes found covers Titan's GPUs [22].
- Every failure-rate threshold and spares percentage in the Nscale credit agreement is redacted [14].
- Lambda's prices are posted prices that take effect 16 November 2026 [29], and KBRA's findings come from a press release with no figures [31].
- CoreWeave's statements on re-contracting and pricing are management statements on an earnings call, with no prices [30].
- No listing prices for used GPU systems were captured for this report, and no verified transaction was found. HIVE's sale has no price [33], and American Compute's figures rest on proprietary data [34].
- The tax section restates IRS Publication 946 for 2025 returns only and is not tax advice [18].

Sources

1. Alphabet Inc., Annual Report on Form 10-K for the year ended 31 December 2023. U.S. Securities and Exchange Commission, filed 31 January 2024, viewed 8 November 2026. https://www.sec.gov/Archives/edgar/data/1652044/000165204424000022/goog-20231231.htm

2. Alphabet Inc., Annual Report on Form 10-K for the year ended 31 December 2025. U.S. Securities and Exchange Commission, filed 5 February 2026, viewed 8 November 2026. https://www.sec.gov/Archives/edgar/data/1652044/000165204426000018/goog-20251231.htm

3. Amazon.com, Inc., Annual Report on Form 10-K for the year ended 31 December 2024. U.S. Securities and Exchange Commission, filed 7 February 2025, viewed 8 November 2026. https://www.sec.gov/Archives/edgar/data/1018724/000101872425000004/amzn-20241231.htm

4. Amazon.com, Inc., Annual Report on Form 10-K for the year ended 31 December 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

5. Meta Platforms, Inc., Annual Report on Form 10-K for the year ended 31 December 2025. U.S. Securities and Exchange Commission, filed 29 January 2026, viewed 8 November 2026. https://www.sec.gov/Archives/edgar/data/1326801/000162828026003942/meta-20251231.htm

6. Microsoft Corporation, Annual Report on Form 10-K for the fiscal year ended 30 June 2023. U.S. Securities and Exchange Commission, filed 27 July 2023, viewed 8 November 2026. https://www.sec.gov/Archives/edgar/data/789019/000095017023035122/msft-20230630.htm

7. Microsoft Corporation, Annual Report on Form 10-K for the fiscal year ended 30 June 2026. U.S. Securities and Exchange Commission, filed 29 July 2026, viewed 8 November 2026. https://www.sec.gov/Archives/edgar/data/789019/000119312526323660/msft-20260630.htm

8. Oracle Corporation, Annual Report on Form 10-K for the fiscal year ended 31 May 2026. U.S. Securities and Exchange Commission, filed 22 June 2026, viewed 8 November 2026. https://www.sec.gov/Archives/edgar/data/1341439/000119312526277521/orcl-20260531.htm

9. CoreWeave, Inc., Annual Report on Form 10-K for the year ended 31 December 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

10. Nebius Group N.V., Annual Report on Form 20-F for the year ended 31 December 2025. U.S. Securities and Exchange Commission, filed 30 April 2026, viewed 8 November 2026. https://www.sec.gov/Archives/edgar/data/1513845/000110465926052948/nbis-20251231x20f.htm

11. Letter to Shareholders, Q2 2026. Nebius Group, 12 August 2026, viewed 8 November 2026. https://assets.nebius.com/assets/4462517b-ce83-41f2-96ed-f2ac1bc06a05/SHLQ226%20(1).pdf-0

12. IREN Limited, Annual Report on Form 10-K for the fiscal year ended 30 June 2026. U.S. Securities and Exchange Commission, filed 27 August 2026, viewed 8 November 2026. https://www.sec.gov/Archives/edgar/data/1878848/000187884826000052/iren-20260630.htm

13. Nscale Ltd, Registration Statement on Form S-1. U.S. Securities and Exchange Commission, filed 18 September 2026, viewed 8 November 2026. https://www.sec.gov/Archives/edgar/data/2110365/000119312526395475/ck0002110365-20260918.htm

14. Credit Agreement dated as of 27 August 2026 among Nscale Ward County Borrower SPV, LLC and others (Exhibit 10.23 to Nscale Ltd Form S-1). U.S. Securities and Exchange Commission, filed 18 September 2026, viewed 8 November 2026. https://www.sec.gov/Archives/edgar/data/2110365/000119312526395475/ck0002110365-ex10_23.htm

15. Comment letter to CoreWeave, Inc. on its Draft Registration Statement. SEC Division of Corporation Finance, 13 January 2025, viewed 8 November 2026. https://www.sec.gov/Archives/edgar/data/1769628/000000000025000329/filename2.txt

16. Property, plant, equipment and other assets guide, 4.2 Determining the useful life and salvage value of an asset. PwC Viewpoint, 15 May 2026, viewed 8 November 2026. https://viewpoint.pwc.com/dt/us/en/pwc/accounting_guides/property_plant_equip/property_plant_equip_US/chapter_4_depreciati_US/42_determining_the_u_US.html

17. SEC Reporting Matters: Accounting Changes and Error Corrections. BDO USA, September 2024, viewed 8 November 2026. https://www.bdo.com/getmedia/0ef2400c-81fe-4710-801e-ad5630c50ff1/ASSR_Accounting-Changes-and-Error-Corrections.pdf?ext=.pdf

18. Publication 946 (2025), How To Depreciate Property. Internal Revenue Service, for use in preparing 2025 returns, viewed 8 November 2026. https://www.irs.gov/publications/p946 (Table B-1 also in the PDF edition, https://www.irs.gov/pub/irs-pdf/p946.pdf)

19. Revisiting Reliability in Large-Scale Machine Learning Research Clusters. A. Kokolis, M. Kuchnik, J. Hoffman and others, arXiv 2410.21680, version 2, 6 February 2025, viewed 8 November 2026. https://arxiv.org/abs/2410.21680

20. The Llama 3 Herd of Models. Llama Team, arXiv 2407.21783, version 3, 23 November 2024, viewed 8 November 2026. https://arxiv.org/abs/2407.21783

21. Story of Two GPUs: Characterizing the Resilience of Hopper H100 and Ampere A100 GPUs. S. Cui, A. Patke, H. Nguyen and others, arXiv 2503.11901, version 4, 10 December 2025, viewed 8 November 2026. https://arxiv.org/abs/2503.11901

22. GPU Lifetimes on Titan Supercomputer: Survival Analysis and Reliability. G. Ostrouchov, D. Maxwell, R. A. Ashraf and others, SC20, IEEE, November 2020, full text at OSTI, viewed 8 November 2026. https://www.osti.gov/biblio/1771896

23. NVIDIA Corporation, Annual Report on Form 10-K for the fiscal year ended 25 January 2026. U.S. Securities and Exchange Commission, filed 25 February 2026, viewed 8 November 2026. https://www.sec.gov/Archives/edgar/data/1045810/000104581026000021/nvda-20260125.htm

24. NVIDIA Launches Revolutionary Volta GPU Platform, Fueling Next Era of AI and High Performance Computing. NVIDIA Newsroom, 10 May 2017, viewed 8 November 2026. https://nvidianews.nvidia.com/news/nvidia-launches-revolutionary-

volta-gpu-platform-fueling-next-era-of-ai-and-high-performance-computing

25. CUDA Toolkit Release Notes. NVIDIA, last updated 5 November 2026, viewed 8 November 2026. https://docs.nvidia.com/cuda/cuda-toolkit-release-notes/index.html

26. NCv3 sizes series. Microsoft Learn, last updated 28 July 2026, viewed 8 November 2026. https://learn.microsoft.com/en-us/azure/virtual-machines/sizes/gpu-accelerated/ncv3-series

27. NVIDIA K80 end of support. Google Cloud, last updated 6 November 2026, viewed 9 November 2026. https://cloud.google.com/compute/docs/eol/k80-eol

28. Pricing and usage model updates for Amazon EC2 instances accelerated by NVIDIA GPUs. AWS What's New, 5 June 2025, viewed 8 November 2026. https://aws.amazon.com/about-aws/whats-new/2025/06/pricing-usage-model-ec2-instances-nvidia-gpus/

29. GPU cloud pricing. Lambda, prices effective 16 November 2026, viewed 8 November 2026. https://lambda.ai/pricing

30. CoreWeave, Inc., Q2 2026 Earnings Call, Corrected Transcript. CoreWeave investor relations, call of 11 August 2026, viewed 8 November 2026. https://s205.q4cdn.com/133937190/files/doc_financials/2026/q2/CRWV-US-CORRECTED-TRANSCRIPT-CoreWeave-Q2-2026-Earnings-Call-11August2026.pdf

31. KBRA Releases Research: AI Compute Financings, Evaluating Re-Leasing Risk for GPUs and TPUs. KBRA, 22 June 2026, viewed 8 November 2026. https://www.kbra.com/publications/KNzgvHxx/kbra-releases-research-ai-compute-financings-evaluating-re-leasing-risk-for-gpus-and-tpus

32. Lambda closes $926 million senior secured term loan B facility, backing GPU deployment for an investment-grade customer. Lambda, 27 August 2026, viewed 8 November 2026. https://lambda.ai/blog/lambda-closes-926-million-senior-secured-term-loan-b-facility

33. HIVE Digital Technologies Ltd., Management's Discussion and Analysis for the three and nine months ended 31 December 2024 (Exhibit 99.2 to Form 6-K). U.S. Securities and Exchange Commission, filed 11 February 2025, viewed 8 November 2026. https://www.sec.gov/Archives/edgar/data/1720424/000106299325002116/exhibit99-2.htm

34. GPU Residual Value: Underwriting a $3 Trillion Opportunity. Bernie Margulies and AC Research, American Compute, 15 June 2026, viewed 8 November 2026. https://www.amcompute.com/blog/gpu-depreciation-residual-value-report-2026

35. Rillor Compute Index. Rillor Corporation, viewed 8 November 2026. https://rillor.com/markets/compute-index

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. GPU depreciation and useful life in filings and market evidence. 10 November 2026. https://rillor.com/insights/gpu-depreciation-and-useful-life

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.

Rillor Compute Index series are for research and planning use. They are not offered for settlement of any contract.

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 series, GPU classes and history you need, and how you plan to use them. Contact us at rillor.com/contact?type=index or contact@rillor.com. More on the Rillor Compute Index: rillor.com/markets/compute-index