

Back to Research

GPU Depreciation Uncertainty in AI Infrastructure

February 4, 2026

The dominant accounting lever in AI infrastructure P&L is the estimated useful life of servers/network equipment. Small changes in life assumptions can shift billions of dollars of depreciation across periods, without changing cash economics.

Across the issuer set, management rationales cluster around (i) software and operational efficiencies, and (ii) hardware/data center design improvements that purportedly extend service life. Quantified impacts are often disclosed when the change is material.

Capital markets are increasingly willing to lend against GPU-heavy infrastructure using bankruptcy-remote entities and collateral packages. The next step is to standardize residual value and

remarketing assumptions to support repeatable GPU-backed ABS issuance.

Accounting Useful Life vs Economic Useful Life

Accounting useful life is a management estimate of the period over which a fixed asset provides economic benefit for financial reporting. Economic useful life reflects the real-world period over which the asset remains competitive and monetizable (e.g., via redeployment, resale, or continued service revenue). In fast-cycling GPU markets, these concepts can diverge materially: an accelerator can remain operational while its economic rent collapses due to new-chip performance, power efficiency, or software stack advantages.

Because most issuers use straight-line depreciation for data center equipment, increasing useful life mechanically reduces the periodic depreciation charge and increases reported operating income, while leaving EBITDA and cash flow unchanged. The relevant investor question is not whether longer lives are ‘right’ in the abstract, but whether disclosures provide enough evidence to underwrite (i) service-life reality, (ii) retirement behavior, and (iii) residual value realizations.

For example, assume a $10.0B GPU/server fleet placed in service today, depreciated straight-line with zero salvage value. At a 4-year life, annual

depreciation is $2.50B. At a 6-year life, annual depreciation is $1.67B. That $0.83B/year difference flows through operating income and EPS even if unit economics and cash spend are identical. If management later reverses the estimate or retires assets early, the ‘missing’ depreciation can reappear abruptly via higher future depreciation, accelerated depreciation, or impairments.

Effective date	Asset class*	Change in Useful Life	Quantified Impact*	Ratio
FY2023	servers; certain network equipment	4→6 (servers); 5→6 (certain network)	Depreciation ↓ $2.0B (6M ended 6/30/2023); net income ↑ $1.5B (6M)	No fu
an 1, 2024	servers	5→6	D&A ↓ $2.4B (9M ended 9/30/2024); net income ↑ $1.9B (9M)	Conti softw
an 1, 2025	servers and networking equipment (subset)	6→5 (subset)	Expected D&A ↑ $1.3B in 2025	Ongoi adjust
FY2023 7/1/2022)	server and network equipment	4→6	Depreciation expense shifted: $3.7B from FY23 to future periods (mgmt remark)	
FY2025	servers and network assets	prior range 4–5 → 5.5	Depreciation ↓ $2.3B (9M ended 9/30/2025)	Softw advan
FY2025 Q1 period ended /31/2024)	certain cloud computing equipment	5→6	D&A ↓ $193M; net income ↑ $197M (3M ended 8/31/2024)	Reflec
FY2024	server and network equipment; GPUs	4** (servers/network)	no change quantified	Young estim
FY2025	HPC hardware	5**	no change quantified	Usefu critic

*as disc

**first ye

How issuers try to reduce depreciation volatility

The strategies below are not mutually exclusive; issuers often combine them. The analyst’s job is to connect disclosure signals to an economic hypothesis about service life, retirement behavior, and residual value realizations.

Strategy 1: Fleet segmentation

Apply different useful lives to different cohorts (e.g. by generation, workload class, or deployment model).

This strategy limits the 'blast radius' of estimated changes and allows targeted updates as technology cycles evolve.

Examples:

- Amazon: effective 1/1/2025, a subset of servers and networking equipment moved from 6 to 5 years (subset framing).
- Alphabet: 'certain network equipment' explicitly carved out alongside servers.

Strategy 2: Operational life extension

Extending service life by shifting older hardware down the performance stack (training → inference), improving utilization software, or optimizing power/cooling to keep assets viable longer.

Longer lives reduce depreciation per period; redeployment reduces early retirements and impairment triggers. However, there is a risk of 'zombie fleets' if utilization falls and there is a potential future catch-up via accelerated depreciation or impairments.

Examples:

- Microsoft: rationale explicitly ties life extension to software efficiencies and technology advances.
- Amazon: cites continuous improvements in hardware, software, and data center designs as the basis for 5→6 years.

Strategy 3: Front-loading pain

Recognizing obsolescence earlier via accelerated depreciation, early retirement/abandonment, or held-for-sale classification.

This strategy clears out 'dead' capacity and reduces the risk of later large impairments; can improve the credibility of longer lives on the remaining fleet. However, it leads to a near-term margin hit; investors will ask whether retirements are concentrated in specific generations/workloads.

Examples:

- Meta: held-for-sale classification for certain data center assets and an impairment loss on data center assets held for sale.
- Amazon: discloses derecognition of build-to-suit assets after construction period (lease accounting mechanics can change depreciation profile).

Strategy 4: Risk shifting via structure (SPV!)

Moving asset ownership or financing leverage into bankruptcy-remote entities and/or relying on operating/finance leases, financing obligations, project finance loans, or SPVs that may be consolidated or unconsolidated depending on control and economics.

This strategy shifts residual risk away from the operating company, converts depreciation into lease/interest expense, and can align financing with contract cash flows. However, it clearly leads to hidden leverage risks (guarantees, take-or-pay obligations, service covenants), refinancing risk, and complexity that can impair transparency.

Examples

- CoreWeave: references a special purpose vehicle tied to an OpenAI contract while separately stating it had no off-balance sheet arrangements as of 9/30/2025.
- Meta: announced third-party participation in data center development (contribution to a third party for co-developing data centers) in its 10-Q disclosure; later press releases show larger JV-style financing structures.

Strategy 5: Residual value realization

This is the strategy that I believe will begin to become more prevalent; although it has yet to be used, exercising residual value protection products allows public clouds to immediately refresh their accelerator fleets, while also giving their balance sheets a backstop for useful life estimates.

Credible residual value reduces impairment risk and supports tighter financing terms while also reinforcing useful-life estimates. However, today, there is increased "slippage risk" as secondary markets are still immature and have yet to be tested with large inflows of HPC equipment.



The Aircraft Leasing Analogy

Aircraft finance evolved around three features: (i) standardized appraisals and residual value curves, (ii) deep remarketing infrastructure, and (iii) financing structures (EETCs/ABS) that rely on predictable lease cash flows and credible recovery values.

In my opinion, AI infrastructure is moving in that direction, especially for GPU-dense fleets, because the scale of capex demands capital markets solutions beyond corporate balance sheets. However, there are things that have to change in order to close this gap.

What already exists:

- Bankruptcy-remote SPV owning GPU servers (or GPUs) with perfected security interests and robust collateral reporting.
- Long-duration capacity leases or take-or-pay compute contracts (ideally with investment-grade or highly collateralized counterparties).

What needs more time:

- Independent valuation / residual reference curves (SKU-level) and remarketing waterfalls (sale, redeploy, part-out).
- Credit enhancement sized to technology-cycle risk (OC, liquidity reserve, trigger-based amortization).

GPU Asset-Backed Securities are Next

Those looking to finance tomorrow's data center projects must begin to look beyond what has "worked" thus far and instead innovate by using strategies that have previously succeeded elsewhere.

GPU-backed ABS unlocks a scalable capital stack for AI infrastructure that doesn’t depend on corporate balance sheets keeping pace with ever-growing capex. It turns depreciation and residual-value uncertainty into explicit, priced risk, which lowers the penalty for refreshing fleets and reduces “trapped” hardware exposure.

The net effect is faster deployment, quicker upgrade cycles, and materially more compute online than equity and unsecured debt markets alone can finance.

Abundance needs an exchange.

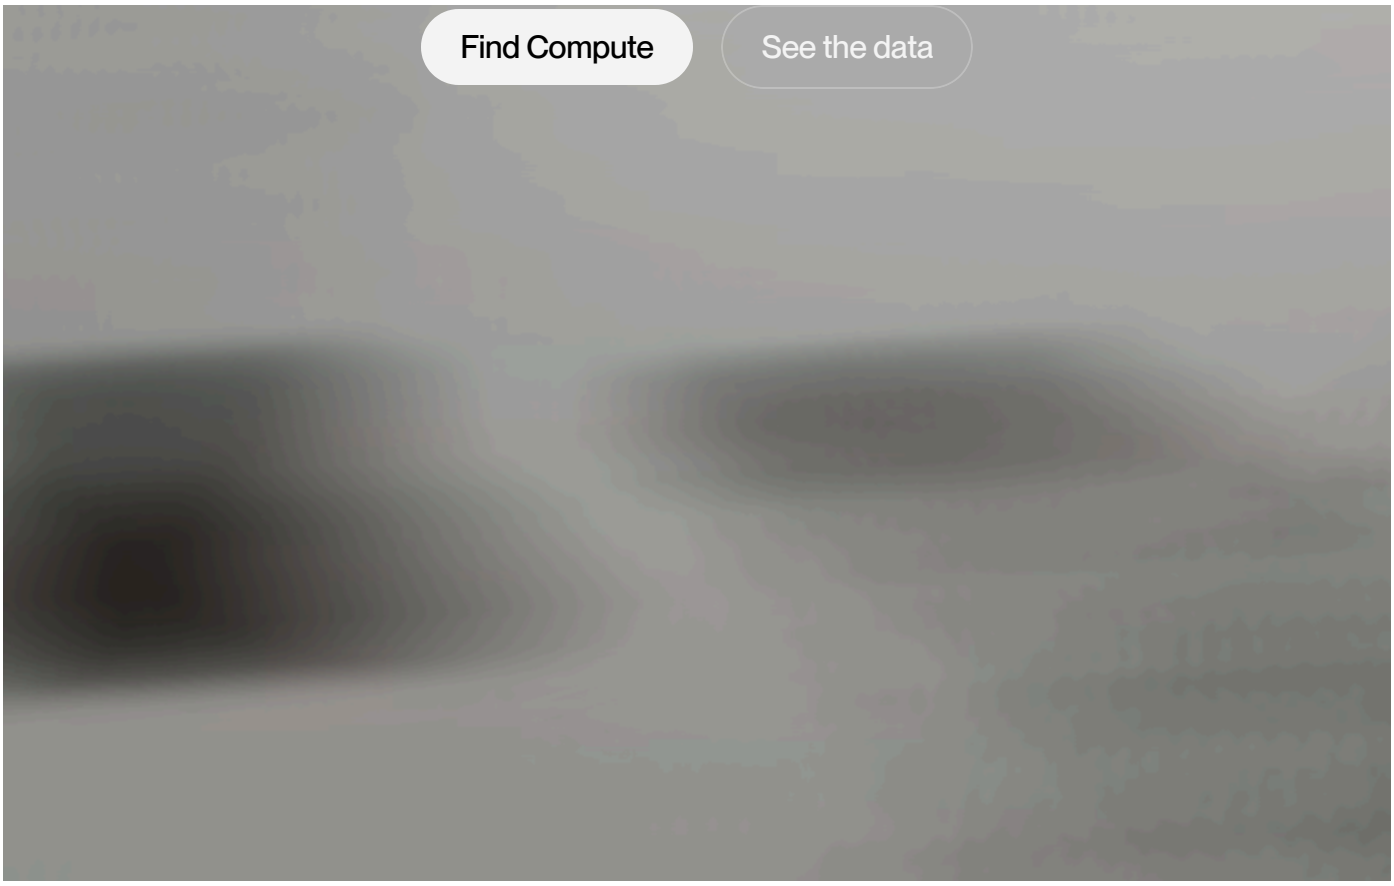

X
LinkedIn
Substack

Careers
About
Research
Docs
Trust Center
Regulatory

Privacy Policy
Terms of Use

Brand + Website by Haptic

©2026 Ornn AI, Inc. All rights reserved.

