// WHITE PAPER · PUBLISHED 16 SEPTEMBER 2026

Depreciation and Useful Life: The Chip Outlives Its Economics

Why two of the largest buyers of server hardware reached opposite conclusions in the same month, and how to set the estimate yourself

Part 4 of 10 in *The Money Language Series*. Each piece stands on its own. The others: Part 1 · Part 2 · Part 3 · Part 5 · Part 6 · Part 7 · Part 8 · Part 9 · Part 10.

Jason Newell · MAX Research Collective · Report WPR-BIZ-005 · Published 16 September 2026
Licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0), https://creativecommons.org/licenses/by/4.0/
Original: https://www.maxresearchcollective.com/articles/depreciation-and-useful-life.html
Cite as: Newell, J. (2026). Depreciation and Useful Life: The Chip Outlives Its Economics. MAX Research Collective, WPR-BIZ-005. https://www.maxresearchcollective.com/articles/depreciation-and-useful-life.html

In January 2025, Meta stretched the estimated useful life of most of its servers and network assets to five and a half years, and expected that one change to cut 2025 depreciation by roughly $2.9 billion. In the same month Amazon went the other way, pulling part of its server fleet from six years back to five and citing the pace of development in AI and machine learning. The reversal cost Amazon about $700 million of 2025 operating income, plus around $920 million of accelerated depreciation on gear retired early (the second figure is a one-time charge rather than a run rate).

Two buyers, one class of hardware, one month, and a swing of billions in reported earnings. Nothing physical separated the two fleets. What separated them was a judgment about how long the boxes would stay worth running.

That judgment is the subject of this piece. Useful life is an estimate, and estimates are where the money sits.

How depreciation works

You buy an asset, and you don't expense it all at once; the cost is spread across the years you expect the thing to be productive, which turns one cash event into a line that reappears on the income statement every year until the schedule runs out. That is depreciation.

Straight-line is the usual book method: cost minus expected salvage value, divided by useful life, charged evenly. A $600,000 server estate on a five-year life with no salvage runs $120,000 a year until the schedule is exhausted. Accelerated methods front-load the charge instead, so more of the cost lands early and less lands late, and they're common for tax and rare for book.

Tax uses a separate system called MACRS (the Modified Accelerated Cost Recovery System), which prescribes recovery periods by asset class. Computer equipment usually sits in the five-year class. You don't pick that number; the code does. That's the first place the book answer and the tax answer come apart.

	BOOK DEPRECIATION	TAX DEPRECIATION
Purpose	Tell investors what happened	Work out what you owe
Life	Your estimate of productive life	Prescribed by asset class under MACRS
Method	Usually straight-line	MACRS, often accelerated
Changeable?	Yes, as a change in estimate	Only inside the code's rules
Special rules	None	Bonus depreciation and Section 179

Which number is the real one? Both are, each for its own reader. Running two schedules on one asset is ordinary practice rather than a sign that somebody made a mistake. Auditors are used to it.

Newell, J. (2026). MAX Research Collective, WPR-BIZ-005. CC BY 4.0

The tax side turned generous

The One Big Beautiful Bill Act, signed 4 July 2025, made 100 percent bonus depreciation permanent for qualified property both acquired and placed in service after 19 January 2025; that one provision killed the phase-down which had walked bonus from 100 percent in 2022 to 80, then to 60, and which was scheduled to reach 40 percent in 2025 and zero by 2027.

Qualified property covers tangible assets with a recovery period of twenty years or less, which takes in servers, storage, networking gear and off-the-shelf software. That is most of a data centre.

Section 179 expensing rose alongside it. The limit went from $1 million to $2.5 million, while the phase-out threshold moved from $2.5 million to roughly $4 million. Both figures are indexed for inflation.

Three details catch people out:

- Both dates have to clear 19 January 2025. Acquired after, and placed in service after; meeting one and missing the other gets you nothing. Acquisition turns on when you became legally committed, so purchase orders and contract dates matter more than the delivery truck.
- Section 179 can't create a net operating loss, and bonus depreciation can. A profitable company often prefers the targeting Section 179 gives it (the election is made property by property), while a company running a loss gets more out of bonus.
- States don't all conform. Many decouple from the federal bonus rules, so your federal answer and your state answer can differ by a wide margin.

For the first time in several years, the tax code is friendly to buying hardware. A $2 million colocation build can be deducted in full in year one for tax while it depreciates over five years for book: one asset, two schedules, both filed, both right. The cash effect lands now and the reported-earnings effect arrives slowly, and that gap is the most common omission in the cloud-versus-colocation comparisons I've reviewed. Finance sees the slow effect on the slide and the fast one only in the bank

balance, which is how a case that clears the hurdle rate on cash still loses the room on reported margin.

The hyperscaler drift

Public useful-life disclosures show how much swing sits inside the estimate; the numbers are audited, the changes are dated, and the arithmetic is checkable by anyone who wants to.

PROVIDER	2017-2019	2020-2022	2023-2024	2025
AWS	3 years	4 then 5 years	6 years	6 years, with a portion moved back to 5
Google Cloud	4 years	4 years	6 years	6 years
Azure	4 years	4 years	6 years	6 years
Meta	3 years	4 to 5 years	4 to 5 years	5.5 years

The industry walked from three or four years to six over about five years, which roughly doubled the assumed life of a server. AWS moved first and the others followed. Why did all three land on six? The stated reasons held up: workloads diversified, software got more efficient, and older hardware kept earning longer than anyone assumed when the three-year schedules were written.

The timing was also convenient. Every year added to a schedule shifts billions from expense into earnings, and the additions arrived while capital spending on AI infrastructure was climbing hard.

Then Amazon partially reversed, which is the entry worth sitting with. It's the only case in the table where a company volunteered a charge, and the reason it gave was the speed of AI hardware development.

The GPU question

If you're buying accelerators, you're making a useful-life call in the one area where the evidence genuinely conflicts. Both camps have audited numbers behind them. Neither one is bluffing.

The case for five or six years rests on the value cascade. A GPU's economic life doesn't end the day it stops being the best training chip available. It steps down to fine-tuning, then to high-value inference, then to batch and internal work that nobody times with a stopwatch; demand for inference capacity is growing faster than demand for frontier training, which keeps the lower rungs of that ladder busy long after the top rung has moved on. CoreWeave underwrites six years on exactly this reasoning.

The case for three or four years rests on power. Silicon generations now arrive about once a year rather than once every two, and each one brings a large gain in performance per watt. Power and cooling cost roughly the same per rack unit whatever generation fills the rack, so an older GPU — same draw, less output — goes uneconomic well before it stops working. Nebius runs four years and Lambda runs five (both disclose the figure, so the gap is a real disagreement rather than a rumour).

Both arguments are right about different things. The question an estimate has to answer is narrower than either of them: how long will this rack earn more than it costs to power? The chip keeps working, and the economics degrade before the hardware does, so useful life should track the crossover rather than the failure.

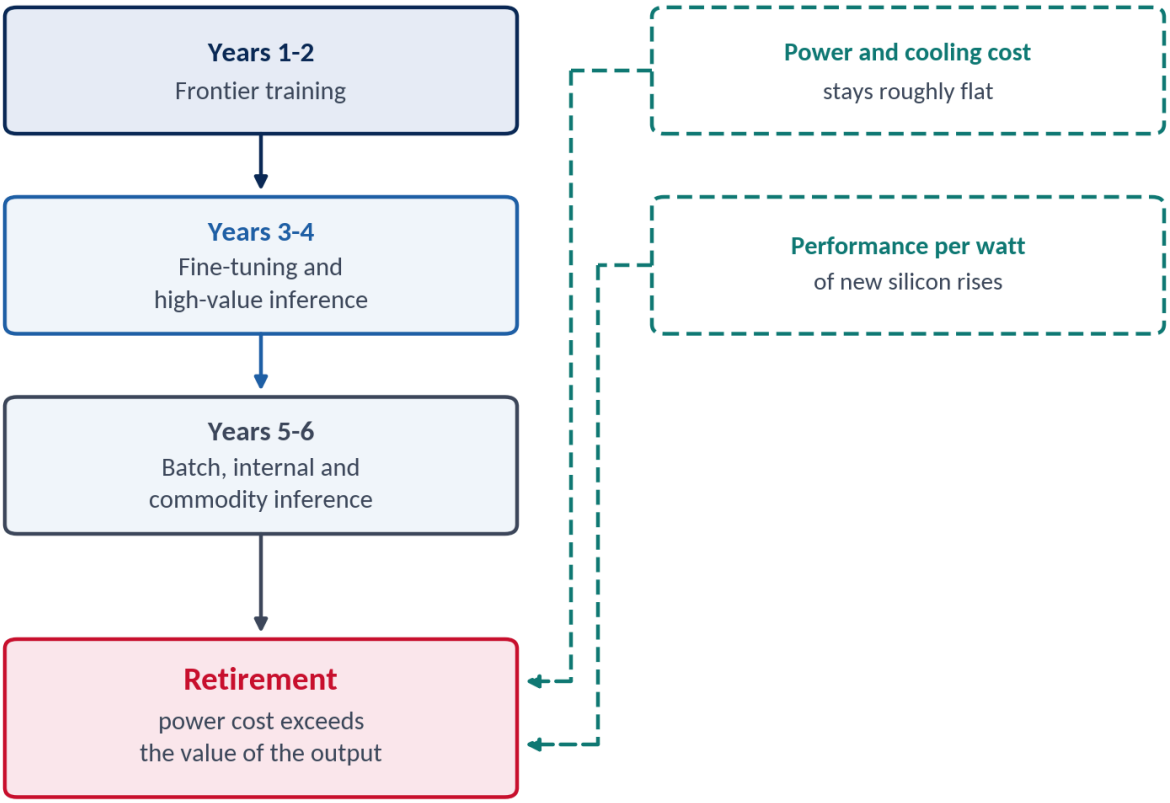


Figure 1. The cascade and the two things that end it. Each stage is lower-value work on the same silicon, and retirement arrives where the flat cost line crosses the falling value line.

Getting it wrong, in both directions

The error runs in both directions, and the two columns cost different things:

LIFE SET TOO LONG	LIFE SET TOO SHORT
Reported earnings look better now	Reported earnings look worse now
The asset carries value it doesn't have	The asset is understated on the balance sheet
Later: impairment or accelerated depreciation, in a lump	Charges land early, and cash is unaffected
Investors and auditors start doubting your estimates generally	Conservative, and rarely questioned
A strategy change gets expensive to announce	There is room to move

The asymmetry is what matters. Setting life too short costs you optics. Setting it too long costs optics and credibility together, and the bill arrives later, in a lump, usually in a quarter that was already going badly. Amazon's $920 million accelerated depreciation charge — one quarter, disclosed in full, reason named — is what the expensive direction looks like when it's handled well.

A mid-market company I advised carried a five-year book life on its storage arrays and a hardware refresh policy of four years. Nobody had put the two numbers side by side. Every refresh threw off a write-off that surprised the finance team, and every time it got filed as a one-off. The pattern was structural. Either the depreciation schedule matches the refresh policy, or somebody writes down why the two differ.

What I actually do

1. Set useful life from the refresh policy rather than from a template. If the fleet gets replaced on a four-year cycle, the life is four years. If finance wants six, the refresh policy is the thing that has to change.
2. Review the estimate every year rather than only at purchase. A change in useful life is a change in accounting estimate, applied forward, and it's an ordinary

action taken in the open. Discovering the same fact through an impairment is the expensive version of it.

3. Model power separately from hardware on anything GPU-heavy. Across five years the electricity is often the larger number, and it's the one that decides retirement.
4. Bring the tax answer and the book answer to the same slide. Section 179 and bonus depreciation can make an owned-hardware case work on cash terms while it still looks unattractive on reported earnings, and whoever signs needs both columns.
5. Check the state conformity position before quoting a tax benefit. Quoting the federal number and finding the state decoupling afterwards has happened once, which was once more than it should have.
6. Never quietly stretch useful life to hit a number. It works exactly once, and it's the sort of thing a diligence process is built to find.

Glossary

TERM	DEFINITION
Accelerated depreciation	Bigger charges early in an asset's life and smaller ones later
Bonus depreciation, Section 168(k)	Deducting the full cost of qualified property in year one for tax
Change in accounting estimate	Revising a judgment such as useful life, applied going forward rather than backward
Depreciation	Spreading the cost of a physical asset across its useful life
Impairment	Writing an asset down because its value dropped
MACRS	The US tax depreciation system, with recovery periods prescribed by asset class
Placed in service	The date an asset is ready and available for use
Salvage value	What you expect the asset to be worth when the schedule ends
Section 179	An election to expense qualifying purchases, capped, and unable to create a loss
State conformity	Whether a state follows the federal depreciation rules
Straight-line	An equal charge in each year of the schedule
Useful life	The number of years an asset is expected to stay productive
Value cascade	Hardware moving down a chain of less demanding workloads as it ages

Newell, J. (2026). MAX Research Collective, WPR-BIZ-005. CC BY 4.0

References

- IRC Section 168(k) and Section 179, as amended by the One Big Beautiful Bill Act, enacted 4 July 2025
- CLA, "Section 179 and Bonus Depreciation, Examples and Tax Strategies"
- CBIZ, "Cost segregation and bonus depreciation benefits under the OBBBA"
- Texas Society of CPAs, "Bonus Depreciation and the OBBBA"
- Amazon.com Inc. Form 10-K, fiscal 2024, useful life disclosure
- Meta Platforms Inc. Form 10-K, fiscal 2024, useful life disclosure
- theCUBE Research, "Resetting GPU Depreciation"
- Deep Quarry, "Amazon revises server lifespan amid AI shift"

Best next read

Part 5, Placement Economics: Cloud, Colo, and the Return of Owned Hardware. You now know what owning hardware costs on paper. Part 5 is where you decide whether to own it at all.

The full series

PART	TITLE
1	CapEx, OpEx, and EBITDA: The Same Dollar, Two Different Stories
2	Reading the Financials: The Primer for Everyone Who Never Had To Care
3	Capitalize or Expense: Your Jira Board Is Now Audit Evidence
4	Depreciation and Useful Life: The Chip Outlives Its Economics, which is this one
5	Placement Economics
6	Lift-and-Shift vs Re-Architect
7	Vendor Economics
8	AI Spend: Three Budgets Pretending to Be One
9	AI and Headcount
10	Outsourcing After Section 174A

This series is an introduction for technology leaders. It is accounting and tax context rather than accounting or tax advice. Treatment is fact-specific. Model it with your CFO and auditor before you commit.

CAREER & LEADERSHIP WPR-BIZ-005

// ABOUT THE AUTHOR

Jason Newell leads organizations and people in technology, and is a researcher, deep technologist and builder — from a CRAY to agentic AI. He founded MAX Research Collective to build AI for regulated industries that runs where its owner has control, with healthcare as its strongest focus. About Jason → Connect on LinkedIn →