

THEMATIC RESEARCH · ACCOUNTING, AI INFRASTRUCTURE & INDEX EARNINGS

The Depreciation Question

The Buyers of AI Compute Book Their Servers Over Five to Six Years. The Silicon Is Superseded Every Two. What Six Income Statements Look Like on a Three-Year Clock — and Why Almost Nobody Has Published the Arithmetic

TON618 CAPITAL RESEARCH · AS OF JULY 17, 2026 ·

Every figure is drawn from the companies' own SEC filings and is as-of dated. Fact ("the filing says...") is separated from estimate ("our model implies...") throughout. The restatement is a model, and §7 states plainly which of its inputs are observed, which are inferred, and which one is fitted. Independently red-teamed against primary filings twice before release; §8 documents the sourcing and the verification standard.

§0 The Argument

\$221bn

UNTAKEN DEPRECIATION, 2026–28E, ACROSS FIVE COMPANIES — ON THEIR OWN GUIDED CAPEX

The claim. Between 2020 and 2025, Microsoft, Alphabet, Amazon, Meta, Oracle and CoreWeave extended the assumed useful life of servers across **thirteen separate disclosed changes** — twelve of them quantified, worth roughly **\$24bn of first-year earnings benefit**. Thirteen changes; thirteen in the direction that raised reported earnings. **One correction the other way** — Amazon's, in February 2025, citing AI. Restated on a three-year clock, FY2025 earnings were **4.4% to 17.2% overstated** per company, and on guided capex the gap reaches **~\$221bn across 2026–28**. The accounting is legal, disclosed and audited. That is not the question. The question is what these income statements look like on the clock the hardware actually keeps.

Depreciation schedules are the quietest line in a 10-K and, right now, the most consequential. The four hyperscalers alone have guided to roughly **\$710bn of capital expenditure in 2026** on stated guidance. Oracle, on a May fiscal year, has already reported **\$55.7bn for FY2026** — 83% of its own revenue — and guides FY2027 higher again. CoreWeave guides **\$31–35bn against \$5.1bn of revenue**. All of it buys infrastructure whose core component, the GPU, is superseded by a new architecture roughly every two years.

We are not alleging fraud. Every change here was disclosed, justified by a useful-life study, and signed off by an auditor. Taken one at a time, each is defensible. Taken together, the pattern is the finding.

Why these six. The four hyperscalers are the *mildest* form of this problem, and on their own they invite an easy dismissal: four to nine per cent is a rounding error on an empire. Oracle and CoreWeave are in scope because they are the same problem with the cushions removed — the same five-to-six-year schedules, run against progressively thinner earnings. And they are the right two: Oracle is the only company beyond the big four

spending at hyperscaler scale — \$55.7bn of FY2026 capex — and the fifth name in the estimate that put this question into public circulation (§3); CoreWeave is the largest publicly traded pure-play AI cloud, the one company whose GPUs are the entire business rather than a line in a segment. The other large buildouts are private and file no 10-K to restate. They are not here as villains. They are the dose-response curve that shows what the estimate is actually carrying, and §5 sets out what they answer that the hyperscalers alone cannot. One scope note: CoreWeave is lossmaking — percentages fail against a negative base — so it is presented on actuals and sits outside the \$221bn forward figure, which covers the five names with earnings.

~\$221bn

Untaken depreciation 2026–28E, pre-tax (band \$192–240bn)

4.4–17.2%

FY2025 net income overstatement, per company

\$303bn

Steady-state server replacement run-rate vs ~\$62bn carried today


13 vs 1

Disclosed life extensions since 2020 vs corrections the other way

§1 What Happened

Useful lives went up as the hardware cycle went the other way.

The sequence matters. Amazon moved first — servers from three years to four in January 2020, a defensible true-up in the CPU era — and by the end of 2021 Microsoft, Alphabet and Meta had all followed. Through 2022–24 the group ratcheted higher: Microsoft and Alphabet straight to six years in 2023, Amazon to five and then six, Meta to five in two steps, Oracle and CoreWeave joining as their own buildouts began. Then, between mid-2024 and February 2025, the consensus cracked: Oracle reached six years and Meta 5.5 — while Amazon, owner of the largest fleet of all, went the other way.



Three facts break the symmetry of "reasonable people can disagree."

Amazon reversed — and it cost double what it guided. In February 2025 Amazon cut the useful life of a subset of its servers from six years back to five, took **\$920m** of accelerated depreciation on early-retired equipment,

and guided to a **\$0.7bn** hit to 2025 operating income from the life change, plus a further **\$0.6bn** from the early retirements carrying into 2025. The life change alone realised **\$1.4bn of additional D&A and \$1.0bn of net income — double Amazon's own guide for that item**. The FY2025 10-K says why: the \$0.7bn guide covered only equipment in service at 31 December 2024, while the realised effect also captured assets *"acquired during the year ended December 31, 2025."* Amazon kept buying the very assets it had just decided wear out faster.

"These two changes above are due to an increased pace of technology development, particularly in the area of artificial intelligence and machine learning."

AMAZON, FY2024 FORM 10-K (0001018724-25-000004) — ON SHORTENING SERVER LIVES AND RETIRING EQUIPMENT EARLY


The asymmetry is visible inside a single filing. In the same disclosure block where Amazon shortened server lives (–\$0.7bn), it *extended* the useful life of heavy warehouse equipment from ten years to thirteen (**+\$0.9bn**). Lives lengthen where the estimate is soft and shorten only where the technology forces the issue.

The supplier ships a new architecture roughly every two years. NVIDIA's datacentre cadence — Ampere (2020), Hopper (2022), Blackwell (2024), Rubin (2026) — is a clean 24 months, each generation with step-function gains in performance per watt. In a power-constrained datacentre the binding input is the megawatt, not the dollar: when a new chip does multiples of the work per watt, the old fleet's economics decay whether or not the silicon still functions. Useful life is an economic estimate, not a mechanical one.

§2 The Restatement

We rebuild each company's server depreciation from its actual capex vintages (10-K XBRL, FY2015–25), straight-line with a half-year convention, applying each disclosed life change prospectively to the in-service fleet — the GAAP treatment. We then re-run the identical fleet on a constant three-year life.

The server share of capex is taken from the companies' own PP&E disclosures, not assumed. Amazon and Meta each break out a clean servers line; Alphabet discloses that *"approximately 60% of technical infrastructure assets were comprised of servers and network equipment."* Dividing the annual change in disclosed server gross by capex yields the share directly. That share is **not constant** — Meta's own series runs 27% (2021) → 46% (2023) → 58% (2024) → 43% (2025) as the AI build displaced data-centre construction. For 2025 the three disclosed readings cluster tightly: **Alphabet 41%, Amazon 45%, Meta 43%.**



The test that matters — and where it does not apply

Each company disclosed what its life extension was worth. A restatement model that cannot reproduce those figures is not measuring anything.

DISCLOSED CHANGE	MODEL	DISCLOSED	RESIDUAL	INDEPENDENT TEST?
MSFT FY23 — servers + network 4→6	\$3.86bn	\$3.70bn	+4%	Yes
GOOGL 2023 — servers 4→6	\$3.43bn	\$3.90bn	–12%	Yes
AMZN 2022 — servers 4→5, network 5→6	\$2.80bn	\$3.60bn	–22%	Yes
ORCL FY25 — servers + network 5→6	\$0.66bn	\$0.73bn	–11%	Yes
META 2025 — to 5.5 yrs	\$2.83bn	\$2.92bn	–3%	No — fitted

Corrected 31 July 2026: the Amazon comparator was published as \$3.10bn, giving a –10% residual. See Corrections in §8.

Four of these five are genuine out-of-sample checks. Meta's is not, and we will not present it as one. Every input to the Microsoft, Alphabet, Amazon and Oracle rows — life schedule and server share alike — was fixed from the filings before the comparison was run; three land within 12% of a hard disclosed fact, and Amazon lands at –22%, the widest residual in this note. That residual runs in the conservative direction — the model reconstructs *less* benefit than Amazon itself disclosed, so it understates rather than overstates the effect this note measures. Meta is different: its filings disclose only a *range* for the pre-2025 life ("Four to Five years"), and we chose the value within that range — 4.65 years — that reproduces its \$2.92bn disclosure. A flat 5.0 misses by 47%. The range is a real constraint, and 4.65 implies roughly two-thirds of the fleet at five years, consistent with Meta's "a majority" language. But the parameter was *solved for*, so Meta's –3% residual measures nothing except that the solver worked.


The discipline throughout is to defer to the disclosure rather than to the fit. A model tuned until its residuals vanish is measuring its own solver: forcing these five to zero requires implied server shares of 30–73%, which contradict the companies' own asset mix. We take the share from the PP&E note and report the error that remains — except where, as with Meta, the disclosure is a range wide enough to fit inside, which we flag rather than exploit.


FY2025, restated

FY2025 — REPORTED VS RESTATED ON A THREE-YEAR SERVER CLOCK

COMPANY	SERVER SHARE	NET INCOME, \$BN	RESTATED, \$BN	EPS	RESTATED EPS	OVERSTATEMENT
Microsoft (FY Jun-25)	45%	101.8	97.0	13.64	13.00	+4.9%
Alphabet	41%	132.2	126.6	10.81	10.35	+4.4%
Amazon	45%	77.7	71.1	7.17	6.56	+9.3%
Meta	43%	60.5	55.5	23.49	21.56	+9.0%
Oracle (FY May-25)	44%	12.4	10.6	4.34	3.70	+17.2%

Aggregate after-tax earnings overstatement across the five: **\$23.8bn in 2025**. Oracle's FY2026 — reported June 2026 — runs at **+25.8%**, EPS \$5.83 restated to \$4.63.






§3 Forward: The Gap Compounds Before Capex Even Grows

Because a three-year schedule burns each vintage twice as fast as a six-year one, the record 2025–26 vintages dominate the forward gap. On guided capex the five companies will recognise **roughly \$221bn less depreciation across 2026–28** than the hardware clock implies — even with capex frozen at 2026 levels, and even before the finance-leased equipment our cash-basis series excludes. That is \$51bn in 2026 rising to \$92bn by 2028, on flat spending.

For scale against the one estimate in wide circulation: Michael Burry put suppressed depreciation at **\$176bn** across 2026–28 in November 2025, over a five-company perimeter that included Oracle. We do not adopt his counterfactual — secondary accounts describe his assumed life variously as two or three years and we could not confirm it from a primary source — so we cite the figure for magnitude only and rely on none of it. Our estimate is built bottom-up from filings and is not directly comparable.


Price your own clock

"Three years" is a judgment, and it is the one place a reader can reasonably refuse to follow us. So we make the choice explicit rather than load-bearing.



The most useful anchor is not our opinion at all: **Amazon's own disclosed useful life for servers was four years from January 2020 until January 2022** — not a bear-case assumption, but the number Amazon itself certified to its auditors, in the era *before* the AI build made the refresh cycle faster. Run the restatement on that four-year clock — longer than our central case — and **the 2026-28E gap is still \$124bn**. A reader who thinks the AI cycle is *faster* than the CPU era it replaced cannot consistently hold that servers deserve a longer life now than they got in 2021.

§4 The Cleanest Form of the Argument




This needs no vintage model and no view on the hardware cycle at all — only guided capex and the disclosed server share. **At steady state with flat capex, annual depreciation equals annual replacement capex under any useful-life assumption:** a company carrying L vintages, each expensing $1/L$ of its cost, recognises exactly one year of capex per year whether L is three or six. The life assumption does not change the destination. It changes only the speed of arrival — three years to converge, or six.

That reframes the whole question. The five are spending toward a **~\$303bn annual server replacement run-rate and currently carry ~\$62bn of server depreciation**. That \$241bn gap is not avoided by a six-year schedule; it is *deferred* by it — and deferred longer, because a longer life keeps reported depreciation below replacement cost for more years while capex compounds. Either AI revenue grows into that expense line, or margins meet it on the way down.

ROIC, honestly

The rebuttal that usually arrives at this point is that returns on capital already absorb all of this and look fine. They do — and the finding cuts against us, so we publish it: restating depreciation cuts NOPAT and accumulated invested capital roughly proportionally, so **FY2025 ROIC levels move by well under a point** — Microsoft 29.7% reported vs 30.0% restated; Alphabet 24.8% vs 24.5%; Amazon 16.3% vs 15.9%; Meta 24.4% vs 23.8%; Oracle 16.4% vs 15.0%, the largest move of the five and still only 1.4 points. (Simplified invested capital: equity plus long-term debt less cash.) Anyone attacking this thesis on "but ROIC is fine" grounds is answering the wrong question. The damage is not in the ratio — it is in the earnings power the ratio is computed from.


Where it does bite is the *incremental* return on the AI build itself:



§5 Same Assumption, Thinner Cushion

If the note stopped at the four hyperscalers, the fair rebuttal writes itself: *a mid-single-digit haircut on an empire's earnings is not a scandal*. Oracle and CoreWeave are the answer to that rebuttal, and the answer is not that they are more aggressive. **It is that they are not aggressive at all — and it happens to them anyway.**

Every company in this note books servers over five to six years. Oracle's six is Alphabet's six, to the year. CoreWeave's six is Alphabet's six. And yet the distortion runs from 4.4% to a sign flip. The tell is Amazon: it carries the **shortest** life of the group at five years, and is **more** distorted than Alphabet on the longest at six. The schedule is not what varies.



What varies is how much weight the same estimate has to carry. Alphabet spends $0.7\times$ its net income on capex; the assumption is a rounding error on the empire. Oracle spends $3.3\times$ its net income, and the identical six-year number produces a 26% distortion. CoreWeave has no earnings base at all, and the same six years decides the sign of its operating line. Rank the five by capex-to-earnings and you rank them, almost exactly, by distortion — nine of ten pairs in order.

That is the finding, and it points forward rather than backward. **Oracle is what a hyperscaler looks like when capex outruns earnings; CoreWeave is what it looks like at the limit.** The hyperscalers' 4–9% is not evidence that the estimate is immaterial. It is evidence that they still have a cushion — and their capex-to-earnings ratios are all rising.

Oracle — the direction, not the level

Oracle is not more aggressive than its peers. It is tied with Alphabet. Verified against each company's latest 10-K, the current server lives are: Alphabet six years (*"we depreciate servers and network equipment generally over a period of six years"*), **Oracle six years**, Meta 5.5, Amazon five. Microsoft discloses no server-specific life at all — only a two-to-six-year computer-equipment range. Any claim that Oracle uniquely stretches lives does not survive the filings, and we will not make it.


What *is* supportable is the cadence. **Oracle extended twice in three years** — servers four to five effective FY2023 (worth \$434m of total operating expenses), then servers and networking five to six effective FY2025 (worth \$733m of operating expenses and \$573m of net income, \$0.20 per diluted share). It arrived at the ceiling in the same window Amazon was retreating from it.

And the fleet carrying that six-year assumption is almost entirely untested: **accumulated depreciation is just 18.5% of gross PP&E** (\$22.7bn of \$122.7bn), the computer and network book **nearly doubled in a single year** (\$30.3bn to \$59.6bn), and \$40.0bn more sits in construction in progress. Oracle has barely begun to run the hardware it is asserting will last six years.

The consequence is arithmetic. Oracle's earnings are small relative to its capex, so the same assumption moves them further: **FY2025 net income 17.2% overstated, FY2026 25.8%**. That is the largest distortion of any profitable company here, and it comes not from an unusual policy but from an ordinary policy applied to an extraordinary capex ramp.

CoreWeave — the weakest defence

CoreWeave needs no share assumption. It discloses an explicit *Technology equipment* line — \$20.9bn gross at end-2025, up from \$9.1bn and \$1.3bn — so additions come straight from disclosed gross. Its reported depreciation covers every PP&E category, not only GPUs, so we model each disclosed category at its disclosed life rather than comparing a GPU-only figure against a whole-fleet one. That tracks **CoreWeave's reported depreciation within 6–16% across 2023–25, with no share assumption and no fitted parameter**. The model runs slightly hot — most plausibly because a half-year convention over-credits service on a fleet whose additions are heavily back-loaded within the year — so the restatement below is **scaled to CoreWeave's reported depreciation base**, the conservative direction.




Restated on a three-year clock, **FY2025 GPU depreciation goes from ~\$2.2bn to ~\$4.3bn**. That converts a reported operating result of **~\$0.05bn into roughly ~\$2.2bn**, and a net loss of **~\$1.17bn into about ~\$3.3bn**. CoreWeave's near-breakeven operating line is not a result; it is a life assumption. D&A already runs at 48% of revenue. On guided 2026 capex the restatement delta alone is **~\$4.7bn — close to the company's entire 2025 revenue of \$5.1bn**. It carries \$21.4bn of total debt and \$8.2bn of operating-lease liabilities against \$3.3bn of equity.

The steelman here is real and we take it seriously. CoreWeave discloses **\$60.7bn of remaining performance obligations** at end-2025, up 302%, with a tail running to **84 months** — *longer* than the six-year life it is defending. Roughly a fifth of that backlog is contracted beyond year four. Assets matched to committed, partly prepaid revenue are a genuine argument, and six years is the same number Alphabet uses.

Three things break it. First, **the cascade defence does not transfer**: a hyperscaler retiring an A100 from frontier training moves it to internal inference and keeps extracting value; CoreWeave has no captive demand sink, because its GPUs *are* the product. Year-six value requires an external customer who will pay rent on six-year-old silicon while the frontier resets every two. Second, **the estimate is stale**: six years was set effective 1 January 2023, on a largely A100 fleet worth \$1.3bn gross, and has never been revisited through H100, H200 and Blackwell — the FY2025 10-K simply restates the 2023 paragraph against a book sixteen times larger. Third, **the defence is correlated, not independent**: Customer A is **67% of 2025 revenue and 68% of receivables**. The contracted backlog that justifies the six-year life and the residual value that justifies it are not two supports. They are one support, counted twice.

The comparison that lands hardest is Amazon. The company with the *strongest* waterfall shortened to five years and took a \$1.4bn hit to do it. CoreWeave, with the *weakest*, sits a full year longer. **The firm least able to redeploy old GPUs carries the most optimistic life**. The number to challenge is not Oracle's. It is CoreWeave's.

§6 Why This Is an Index Question



Step back from six income statements and one structure remains: many buyers, one seller, and opposite clocks. **NVIDIA books \$216bn of FY2026 revenue the moment the silicon ships; the five buyers in our forward model recognise \$99bn of server depreciation over the same window** — and would recognise \$150bn on a three-year clock. The aggregate P&L of the AI complex recognises the profit now and the cost later. That is not a conspiracy; it is what happens when a fast-cycling asset is sold by one party and capitalised by another on a slow schedule.

The four extending hyperscalers earned roughly **\$372bn in FY2025** — **about 61% of Mag7 net income** (\$608bn including NVIDIA's FY26, Apple and Tesla), which in turn is roughly a third of S&P 500 market capitalisation (~32.5% as of July 2026) and, per FactSet, contributed about a third of the index's 2025 net-earnings growth. Across the five restated names — Oracle sits outside the Mag7 but well inside the S&P 500 — **~\$23.8bn of FY2025 after-tax earnings exists only by grace of the schedule**; on guided capex that figure roughly doubles by 2027, and the distortion concentrates exactly where the index's growth narrative lives. A market paying a premium multiple for high-quality, high-ROIC earnings is compounding the multiple on the flattered base.

The steelman, and why it fails at the margin

Their case. Old GPUs cascade to inference and internal workloads; A100s from 2020 still run at high utilisation, and six years of *revenue-generating* service is documented. Much of the fleet is still CPUs, storage and network gear, where five to six years is genuinely right — and servers are under half of capex on the companies' own disclosures. Depreciation is non-cash; free cash flow already bears the full capex. Auditors re-test these lives annually against utilisation and failure data, and keep passing them.

Our response. The waterfall argument was calibrated on the CPU era; the marginal dollar now buys GPU systems in a *power-constrained* datacentre, where keeping old silicon powered is an opportunity cost, not an asset. We grant the mix point and price it — that is precisely why the restatement uses the disclosed 41–45% server share rather than the whole capex line; the gap survives the concession. Amazon, the largest operator of the waterfall, is the one that shortened lives citing AI, and its realised cost ran double the guide. And "the market prices FCF" concedes the point: if FCF is the honest number, the premium multiples on *reported earnings* rest on the flattered line.

Auditors test whether machines are *in service*, not whether the service is economic. Thirteen changes, ~\$24bn, one direction, during the greatest hardware-obsolence cycle in computing history. Pattern, not fraud — but pattern.

What to watch

- **Useful-life disclosures in each January–February 10-K cycle.** Any second company following Amazon down re-rates the whole group's schedule credibility.
- **GPU rental spot rates and the secondary market for H100/A100 systems** — the observable market price of "useful life."
- **Depreciation growth vs capex growth in quarterly filings.** The wedge is now large enough to track in real time.
- **The first capex deceleration.** When spending flattens, the deferred depreciation from record vintages keeps arriving, and the earnings "beat" machine runs in reverse.

The question, answered

Restated on the clock the hardware actually keeps, the five profitable names in this note earned **\$23.8bn less in FY2025 than they reported**, and on their own guided spending will defer **roughly \$221bn more** of recognition into 2026–28. None of that requires the schedules to be dishonest — only long, and staying long while the asset gets shorter. Thirteen extensions in one direction tell you which way the accounting wanted to go; Amazon's single reversal tells you which way the hardware is going. The distance between those two statements is the depreciation question — and the index, for now, is priced as if the answer were zero.

§7 What This Model Is, and Is Not

A restatement is a model. These are the places it can be wrong, stated before anyone asks.

- **Microsoft and Alphabet disclose no servers line.** Microsoft's share is inferred from a blended "computer equipment and software" line; Alphabet's rests on a single 60% footnote. Only Amazon and Meta break servers out cleanly — and Amazon introduced its line only in the FY2025 10-K, so no pre-2024 history exists for it.
- **Three of five disclose a life range, not a point** (MSFT 2–6 yrs, AMZN 5–6, META 5–5.5). Any single-life-per-company chart is a simplification. **Meta's blended pre-2025 life is fitted** — inside its disclosed 4-to-5-year range, but chosen within it to reproduce the \$2.92bn disclosure. Its validation residual is therefore not an independent test. The other four are.
- **The server share is probably understated.** $\Delta(\text{disclosed server gross}) \div \text{capex}$ nets out retirements and disposals, because write-offs reduce the gross balance against which new additions are measured. 2025 was a heavy-disposal year at both Amazon and Meta by their own disclosure. If so the true share exceeds 41–45% and every figure here is conservative — but we have not quantified the bias and claim no credit for it.
- **Oracle's share derivation is noisier than its peers'** because of construction-in-progress timing: the numerator captures only assets that have left CIP, while the denominator includes CIP spend not yet transferred. Oracle's CIP grew \$16.5bn to \$40.0bn in FY2026 alone — about 42% of that year's capex — which is the most likely reason its disclosed share swings between 35% and 82% across years. We use its own year-by-year series rather than smoothing it.
- **Oracle's disclosed benefits use an opening-balance convention**, so they exclude servers acquired during the year and understate the full policy effect. We validate Oracle on that same basis, to match. Its FY2023 figure is operating-expense-only: no net income or EPS figure exists for it.
- **Basis.** The series is cash PP&E and excludes finance-leased equipment, which is also depreciated — conservative. Microsoft and Meta guide on a finance-lease-inclusive basis and are scaled to the cash basis;

Microsoft's calendar-to-fiscal-year conversion is an estimate. **The \$710bn headline and the \$630bn model basis are different things.**

- **Oracle and CoreWeave guide on company-defined non-GAAP bases.** Oracle's "~\$70bn FY2027" is *net cash outlay for capex*, a metric introduced only in Q4 FY26; its own reconciliation (\$55.663bn GAAP – \$3.345bn – \$4.592bn = \$47.726bn) gives a 0.857 factor, so we model FY2027 at the GAAP-equivalent ~\$81.6bn. Management indicated reported capex would run \$20–25bn higher (~\$90–95bn); that figure exists only in third-party transcripts — Oracle publishes none — so we do not rely on it. Oracle's slide labels the \$70bn "FY27+", a run-rate rather than an explicit single-year guide. CoreWeave's guided "capex" is its own defined measure; its FY2025 sits at **\$10.3bn cash / \$14.9bn guided-basis / \$21.1bn gross additions** — three numbers for one year.
- **CoreWeave's non-GPU book is no longer trivial** — \$3.66bn by end-2025, 17.5% of the Technology-equipment line — which is why we model every disclosed category and scale the restatement to its reported depreciation base rather than compare a GPU-only figure against a whole-fleet one. Its filings never say "GPU"; the category is "Technology equipment." Our GPU reading is an inference, not a quotation. CoreWeave is excluded from the forward aggregate and presented on actuals.
- **2027–28E hold capex flat** at 2026 guidance against companies guiding increases — conservative by construction.

§8 Source & Verification Note

Every company figure is from SEC filings via EDGAR (XBRL company facts and 10-K note text, FY2015–FY2026), and every life-change quote and dollar figure in this note was verified verbatim against the filing that carries it. The model and its inputs were then **independently red-teamed twice against primary filings** by an adversarial pass instructed to refute rather than confirm — re-pulling the filings itself rather than trusting our workings, and re-running the model end to end. Both passes reproduced every headline number and found no fabricated primary figure. Where those passes changed our inputs, the numbers published here are the corrected ones; §7 states the limitations that remain. One *misattributed* figure did survive both passes and is corrected below — a genuine Amazon disclosure attached to the wrong change and the wrong year. It is the class of error a verifier asking "does this number appear in the filings?" will confirm rather than catch, and we record it as such.

Secondary sources were used only to locate filings, never as the basis for a figure. Michael Burry's \$176bn is cited for magnitude only, with explicit non-reliance.

Corrections. *Version 1.1 — 31 July 2026 — §2 validation table, Amazon row.* As first published on 17 July 2026, that row compared the model's \$2.80bn against a disclosed benefit of "\$3.10bn" and reported a –10% residual, and the sentence beneath it stated that all four independent checks landed within 12%. **The \$3.10bn is the wrong figure.** It is Amazon's *anticipated 2024 operating-income impact* from the January 2024 extension of server lives from five years to six, disclosed in the FY2023 Form 10-K — not the January 2022 change (servers four to five, networking five to six) that the row tests. Amazon's FY2022 Form 10-K (0001018724-23-000004) discloses the 2022 change as "*a reduction in depreciation and amortization expense of \$3.6 billion and a benefit to net income of \$2.8 billion, or \$0.28 per basic share and \$0.28 per diluted share.*" Measured pre-tax, as every other row in the table is, the correct comparator is \$3.60bn and the residual is –22%. One coincidence is worth naming so no reader mistakes it for a match: the model's \$2.80bn equals Amazon's *after-tax* \$2.8bn exactly, but comparing a pre-tax model output with an after-tax disclosure is not a like-for-like test, and we do not claim it as one. **No other figure in this note changes.** The validation table is a test *of* the model, not an input *to* it: the ~\$221bn forward estimate, the 4.4–17.2% FY2025 overstatements, the clock sensitivities and every company table stand as published. The error was found on 31 July 2026 while preparing the Amazon single-name note in this series, which narrates Amazon's full life-change sequence and so forced each disclosure to be matched to its own year.

Filings relied upon. Microsoft FY2021 (0001564590-21-039151), FY2022 (0001564590-22-026876), FY2023 (0000950170-23-035122), FY2025 (0000950170-25-100235); Alphabet FY2020 (0001652044-21-000010), FY2022 (0001652044-23-000016), FY2023 (0001652044-24-000022), FY2025 (0001652044-26-000018); Amazon FY2021 (0001018724-22-000005), FY2023 (0001018724-24-000008), FY2024 (0001018724-25-000004), FY2025 (0001018724-26-000004); Meta FY2021 (0001326801-22-000018), FY2022 (0001326801-23-000013), FY2023 (0001326801-24-000012), FY2024 (0001326801-25-000017), FY2025 (0001628280-26-003942); Oracle FY2023 (0000950170-23-028914), FY2025 (0000950170-25-087926), FY2026 (0001193125-26-277521), Q3 FY26 8-K Ex-99.1 (0001193125-26-100148); CoreWeave S-1 (0001193125-25-058309), FY2025 10-K (0001769628-26-000104). Capex guidance as stated at the Q4-2025 and Q1-2026 prints: Microsoft \$190bn CY26; Alphabet \$180–190bn; Amazon ~\$200bn; Meta \$125–145bn; Oracle ~\$70bn FY27 net cash outlay (Q4 FY26 slides, 10 June 2026); CoreWeave \$31–35bn (CFO Nitin Agrawal, Q1 2026 call, 7 May 2026).

Reference prices. Closing prices, 17 July 2026: Microsoft \$401.10 · Alphabet \$354.46 · Amazon \$249.89 · Meta \$664.54 · Oracle \$124.27 · CoreWeave \$72.91 · NVIDIA \$207.40. Model data as of the same date. This note carries no target price and takes no position in any name.

§9 Disclosures

Information only. TON618 Capital. This report is for information purposes only. Nothing here is an offer to sell or a solicitation of an offer to buy any security, fund interest, or digital asset, and nothing here is personalized investment advice or a recommendation regarding any instrument.

Publisher's exclusion. All research is published solely as general, impersonal information of regular circulation. It is not tailored to the objectives or circumstances of any individual and is not issued in connection with compensation from any client. The Fund has no clients and distributes all research free of charge. On that basis it publishes in reliance on the publisher's exclusion from the definition of "investment adviser" under the Investment Advisers Act of 1940 (§202(a)(11)(D); cf. *Lowe v. SEC*, 472 U.S. 181 (1985)).

Registration & conflicts. TON618 Capital is not registered as an investment adviser or broker-dealer in any capacity. The Fund is a Bitcoin fund and may hold or transact in the securities or digital assets it discusses. This report discusses Microsoft (MSFT), Alphabet (GOOGL), Amazon (AMZN), Meta Platforms (META), Oracle (ORCL), CoreWeave (CRWV) and NVIDIA (NVDA); the Fund holds no position, long or short, in any of them, and has no economic interest in the price of any security named here. The Fund receives no compensation from any party in connection with its research.

Use of AI. Artificial intelligence is used in the creation of this research. All methodology and data integrity are reviewed and approved before publication by TON618 Capital's Chief Investment Officer, **Keyth Beck**; errors may nonetheless occur, and readers should verify independently.

CFA. This report was prepared to align with CFA Institute analytical standards (methodology only). CFA® and Chartered Financial Analyst® are registered trademarks owned by CFA Institute. That reference describes the analytical framework applied; it does **not** imply the report was prepared, reviewed, or authored by a CFA charterholder, and the report is **not** issued, reviewed, endorsed, certified, or approved by — nor affiliated with — CFA Institute.

Risk & feedback. Past performance is not indicative of future results. Digital assets and equities are volatile and may result in total loss of capital. Corrections and feedback are welcome — please direct them to CIO **Keyth Beck** at keyth@ton618capital.com.