THEMATIC RESEARCH · ACCOUNTING AND VALUATION

Capital Expenditure, Depreciation and the Quality of US Earnings

profit now, cost later

When a technology company buys AI chips, the cost reaches its income statement over five or six years; the chipmaker books the profit immediately. This report measures how much that timing difference, and other features of AI-era accounting, are flattering S&P 500 earnings, and when the bill arrives.

Global Markets Research · Equity Strategy Desk Prepared by ANDRA MARKETS

the claim, tested

Coverage summary. A common argument among sceptical fund managers is that the rise in US equities rests on an accounting asymmetry: the largest technology companies capitalise their AI spending and depreciate it over several years, while their suppliers record the same spending at once as revenue and profit. We test the claim with company filings and our own model. It is correct in its mechanism: the five largest spenders will invest about US$800 billion this year, against depreciation of only about US$185 billion on our estimate, while the quarterly net income of the main AI suppliers is about US$83 billion higher than a year ago. It is also correct that several companies have extended the useful lives of their servers, adding billions to reported profits. But it overstates the effect on the index: on our model, using four-year rather than six-year lives would reduce S&P 500 earnings by about 1 to 3 per cent in 2026 to 2028, not wipe out the rally. The larger issues for earnings quality in 2026 are elsewhere: unrealised gains on stakes in AI laboratories, which added about 12 per cent to second-quarter earnings, and a growing gap between reported profits and cash. The accounting is legal; the timing is the risk: depreciation is about to rise sharply.

Top-5 capex 2026	Supplier net income	AI stake gains, Q2	Hyperscaler NI vs FCF
~$800bn	**+$83bn**	**~$150bn**	**$600bn / $121bn**
vs depreciation ~$185bn	*per quarter vs a year ago*	*12% of Q2 S&P 500 EPS*	*trailing 12 months*

1.0 Executive Summary

Key takeaways

mechanism: real

The asymmetry is real. Under US accounting rules, a server bought for US$100 is recorded as an asset and expensed over its useful life, now five and a half to six years at the largest companies, while the supplier records the sale immediately. With spending growing by 50 to 80 per cent a year, the annual cost recognised by buyers is a fraction of the revenue recognised by sellers: we estimate that in 2026 the top five spenders will invest about 4.3 times their depreciation charge.

useful lives: +$36n each

Useful lives have been stretched. Microsoft, Alphabet, Amazon, Meta and Oracle all lengthened server lives between 2021 and 2025; each change added about US$0.4 to 3.9 billion to annual operating profit, and Meta's 2025 change alone added US$1.00 to its earnings per share. Amazon is the only one to have reversed part of it, citing the pace of AI.

Exhibit A. Earnings quality scorecard: what flatters S&P 500 profits in 2026

Factor	Mechanism	Size (estimate)	Reverses when
Capex vs depreciation timing	buyer spreads cost, seller books profit	capex ~4x depreciation	depreciation catches up (2027-28)
Longer useful lives	servers depreciated over 5.5 to 6 years	up to ~$4bn a year per change	lives shortened (as Amazon did)
Unrealised gains on AI stakes	mark-to-market of OpenAI, Anthropic stakes	~$150bn in Q2 (12% of EPS)	valuations fall
Off-balance-sheet leases	costs as operating leases, outside capex	~$0.8 to 1.1tn commitments	leases commence
Vendor financing and gains at suppliers	suppliers invest in customers	Nvidia: $23.7bn gains in H1	customers fail to pay

Andra Markets framework; sizes from company filings, FactSet, LSEG via Reuters, Goldman Sachs, JPMAM and our model (Sections 4 to 8). Colour convention (charts only): green marks investment and capacity, red marks costs, gaps and risk.

But the effect on the index is modest. Our model shows that four-year lives would lower S&P 500 earnings by about 1.2 per cent in 2026, 1.8 per cent in 2027 and 2.7 per cent in 2028; even three-year lives would cut them by less than 5 per cent. The rally is not an accounting illusion, but its profits are front-loaded.

The bigger questions lie elsewhere. Unrealised gains on stakes in AI laboratories, the gap between reported profit and free cash flow, and a wave of depreciation that Goldman Sachs expects to reduce S&P 500 earnings growth by about 5 percentage points in 2027 and to neutralise AI's contribution by 2028.

2.0 The Mechanism

One dollar, two income statements

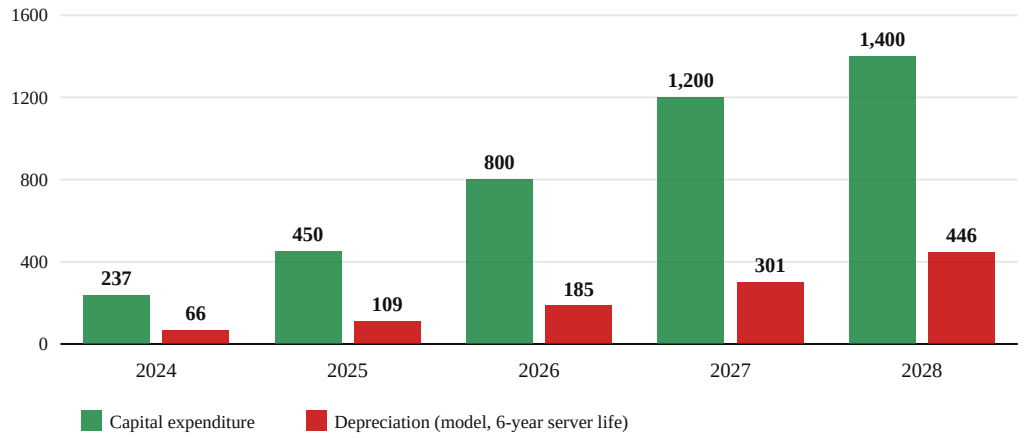
Consider a hyperscaler that buys US$60 billion of servers and US$40 billion of buildings and power equipment. In the first year, with servers depreciated over six years and buildings over twenty, it records about US$6 billion of depreciation (half a year of each); the rest stays on the balance sheet. Its suppliers, Nvidia, Broadcom, Micron, Dell and the builders, record US$100 billion of revenue in the same year and, at Nvidia's current 75 per cent gross margin on chips, a large share of it as profit. Across the economy, profits rise long before the cost appears. This is not a manipulation: it is how investment has always been accounted for, and the same logic applies to factories or power plants. The issue is one of degree and timing: spending on AI is growing so fast, and the assets are so short-lived, that the gap between cost incurred and cost recognised is unusually large.

3.0 Useful Lives

Five companies, one direction

Between 2021 and 2025 every large cloud company extended the period over which it depreciates servers, in most cases from three or four years to six. The disclosed effects, taken from their annual reports, are summarised below. The justification offered was that software improvements allowed servers to be used longer. The AI cycle challenges that logic: Nvidia now launches a new generation every year (Blackwell, Rubin in 2026, Rubin Ultra in 2027), and Amazon in 2025 shortened part of its fleet back to five years "due to the increased pace of technology development, particularly in the area of artificial intelligence and machine learning", a change that raised its depreciation by US$1.4 billion. Microsoft has gone the other way for buildings, extending their lives from 15 to 25 years from fiscal 2027, although it says the profit effect is minimal.

Exhibit 1. Top-5 hyperscalers: capital expenditure vs modelled depreciation ($bn)



gap: ~$600bn in 2026

depreciation x2.4 by 2028

Sources: capex from company reports (2024, 2025 approximate) and Goldman Sachs forecasts (2026 to 2028). Depreciation is an Andra Markets model: 62% of capex in servers over 6 years, 38% in buildings and power over 20 years, half-year convention. Modelled 2026 depreciation (~$185bn) compares with about $167bn reported over the last twelve months.

Exhibit B. Changes in server useful lives and their disclosed effects

Company	Year	Change	Disclosed effect
Alphabet	2021	servers 3 to 4 years	depreciation −$2.6bn; net income +$2.0bn
Microsoft	FY2023	servers 4 to 6 years	operating income +$3.7bn; EPS +$0.40
Amazon	2022	servers 4 to 5 years	D&A −$3.6bn; net income +$2.8bn
Alphabet	2023	servers 4 to 6 years	depreciation −$3.9bn; net income +$3.0bn
Amazon	2024	servers 5 to 6 years	D&A −$3.2bn; net income +$2.5bn
Oracle	FY2025	servers 5 to 6 years	expenses −$0.7bn; EPS +$0.20
Meta	2025	servers to 5.5 years	depreciation −$2.9bn; EPS +$1.00
Amazon	2025	subset 6 back to 5 years	D&A +$1.4bn; EPS −$0.10
Microsoft	FY2027	buildings 15 to 25 years	“minimal” profit benefit

Meta: $1 of EPS

only reversal

Source: company annual reports (Forms 10-K) and Microsoft’s Q4 FY26 earnings call (29 July 2026). Effects are for the first year after the change.

GPU life: 2 to 6 years?

How long do AI chips really last? The evidence is mixed. Rental prices for Nvidia’s H100 at hyperscalers fell from a median of about US$9.34 an hour in the second half of 2024 to US$6.26 a year later, according to Silicon Data, and an index of market prices stood at US$2.65 on 11 October; Amazon cut its own H100 prices by 44 per cent in 2025. Michael Burry argues that the economic life is two to three years, and Jim Chanos three to four. On the other side, CoreWeave’s contracts average about five years, older A100 chips remain fully booked, and Bernstein estimates that GPUs remain profitable to operate for about six years. Microsoft’s chief executive has said he does not want to carry four or five years of depreciation on a single generation, a reported comment that captures the dilemma.

H100 rent: −33% from peak

4.0 The Depreciation Wave

The bill arrives in 2027-28

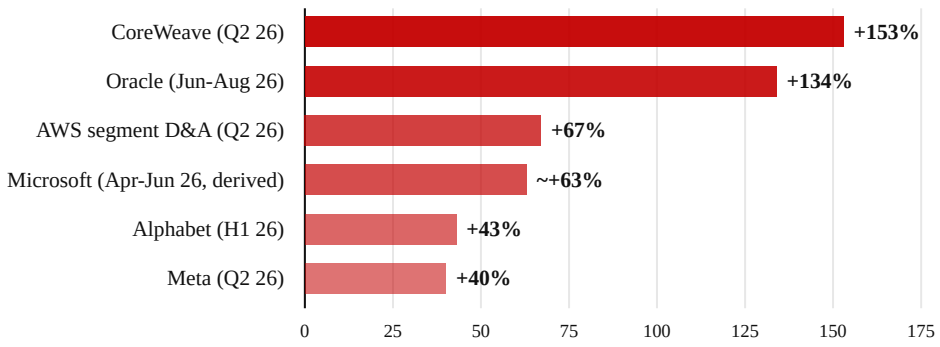
MSFT dep. +63%

Depreciation is already accelerating: Microsoft’s rose 56 per cent in fiscal 2026 and about 63 per cent in the June quarter, Alphabet’s 43 per cent in the first half, Meta’s 40 per cent and Oracle’s 134 per cent in their latest quarters. Much more is in the pipeline: assets under construction, which do not depreciate until they enter service, account for 32 per cent of Oracle’s gross property, 29 per cent of Alphabet’s (US$122.8 billion) and 27 per cent of Meta’s. On our model, the top five’s depreciation rises from about US$185 billion in 2026 to about US$300 billion in 2027 and US$450 billion in 2028 even with six-year lives. Goldman Sachs estimates that hyperscaler depreciation will subtract about 5 percentage points from S&P 500 earnings growth in 2027 and that by 2028 it will fully offset the earnings boost from AI

$122bn not yet depreciating

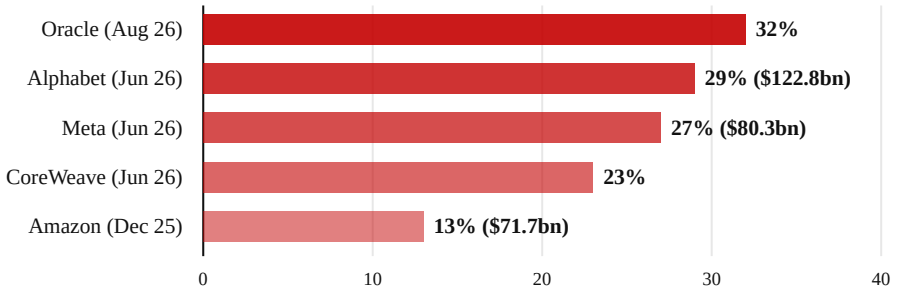
spending.

Exhibit 2. Depreciation growth, latest reported period (% year on year)



Source: Forms 10-Q and 10-K; Microsoft June quarter derived from annual and nine-month figures. Definitions differ (Amazon reports total D&A).

Exhibit 3. Assets not yet depreciating: construction in progress as % of gross property and equipment



Source: Forms 10-Q and 10-K; Andra Markets calculation. Microsoft does not disclose construction in progress; it reports $329bn of leases signed but not yet commenced.

5.0 How Much Is at Stake

An Andra Markets sensitivity

To size the effect, we re-run our depreciation model with shorter server lives and compare the after-tax difference with aggregate S&P 500 earnings, which we estimate at about US$3.2 trillion in 2026 and US$3.7 trillion in 2027 from FactSet's bottom-up forecasts. With five-year lives, earnings would be lower by less than 1 per cent in 2026 and 2027; with four-year lives, by 1.2 per cent in 2026, 1.8 per cent in 2027 and 2.7 per cent in 2028; with three-year lives, the most pessimistic view, by 2.2, 3.4 and 4.7 per cent. These are meaningful amounts, between US$40 and 190 billion a year, but they are an order of magnitude smaller than the 33 per cent growth in S&P 500 earnings expected this year. For individual companies the effect is much larger: Burry estimates that by 2028 Oracle's earnings would be overstated by 26.9 per cent and Meta's by 20.8 per cent.

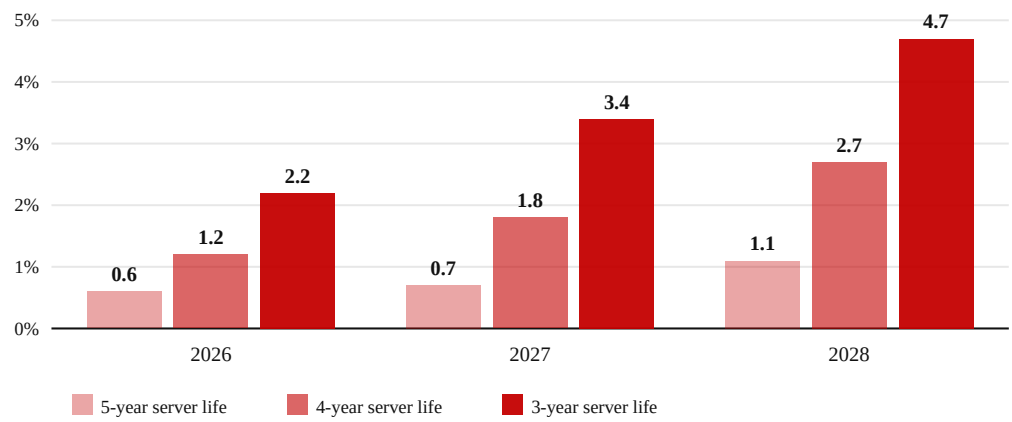
6.0 The Seller Side

Where the profits appear

The other half of the asymmetry is visible in suppliers' accounts. Nvidia earned US$59.7 billion of net income in its latest quarter, 126 per cent more than a year earlier, on a gross margin of 75 per cent; Micron's quarterly profit rose about twelvefold to US$37.7 billion as memory prices soared; Broadcom's more than tripled to US$13.1 billion. Combining the main AI suppliers in the S&P 500, quarterly net income is about US$83 billion higher than a year ago, which annualises to about US$330 billion, roughly 10 per cent of expected 2026 index earnings. Goldman Sachs estimates that AI investment accounts for nearly half of S&P 500 earnings growth this year, and JPMorgan's Michael Cembalest finds that a small group of AI-related stocks has generated the majority of the index's earnings growth since 2024. The concentration is visible in sector data: information technology will contribute about 32 per cent of third-quarter index earnings, and its expected earnings growth falls from 65 to 24 per cent excluding semiconductors.

Exhibit 4. Reduction in S&P 500 earnings if servers were depreciated over shorter lives (% of index earnings)

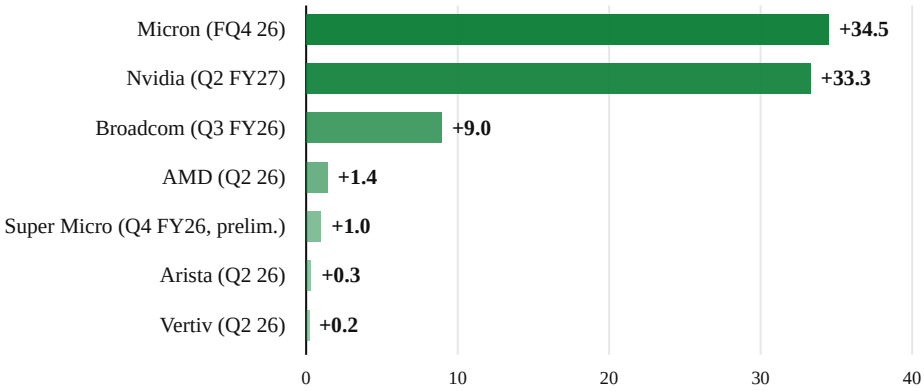
a few %, not the rally



Andra Markets model: top-5 hyperscaler capex (company data and Goldman Sachs forecasts), 62% in servers, base case 6 years, 17% tax rate. S&P 500 aggregate earnings from FactSet bottom-up EPS ($363.41 for 2026, $420.59 for 2027) times an index divisor of about 8.87bn; 2028 assumes +10%. Illustrative.

Exhibit 5. AI suppliers: increase in quarterly net income vs a year earlier ($bn, latest quarter)

memory + GPUs = most of it



Source: company results (Forms 8-K), quarters ending June to September 2026; Andra Markets calculation. Nvidia's figure includes $7.8bn of gains on equity investments in the quarter. Dell excluded (not verified against filing).

7.0 Profits vs Cash

The gap widens

NI $600bn, FCF $121bn

The clearest test of earnings quality is to compare reported profit with cash. Over the past twelve months, the five largest spenders reported combined net income of about US$600 billion but generated only about US$121 billion of free cash flow, and two of them, Amazon and Oracle, had negative free cash flow. Part of the gap is the investment itself, which is the point of capitalisation; but part comes from non-cash gains: Alphabet recorded a US$77.1 billion unrealised gain and Amazon US$53.4 billion of other income, mainly from its Anthropic stake, in the second quarter. LSEG data show that S&P 500 earnings grew 52 per cent in the second quarter, but only about 33 per cent excluding gains on AI stakes, and Goldman Sachs estimates that about US$150 billion of equity-investment income at large technology companies made up 12 per cent of second-quarter index earnings per share. In 2026, the largest single distortion in reported profits is not depreciation but mark-to-market gains on private AI companies.

Amazon, Oracle: FCF < 0

8.0 The Circular Element

Gains on both sides

Nvidia: $23.7bn gains in H1

The feedback loops reinforce each other. Hyperscalers sell computing power to AI laboratories, invest in those same laboratories, and record gains when the laboratories raise money at higher valuations, money that is then spent on computing power. Suppliers do the same: Nvidia's net income in the first half of its fiscal year included US$23.7 billion of gains on equity investments, including stakes in customers such as OpenAI and xAI. None of this is improper, and the gains are disclosed, but it means that a part of reported profit on both sides depends on

stake gains = customers' valuations

private valuations that can reverse. As BofA's Savita Subramanian put it, "mark-to-market gains can turn into losses just as fast".

Exhibit C. Net income vs free cash flow, trailing 12 months ($bn)

Company	Net income	Operating cash flow	Capex	Free cash flow
Amazon (to Jun 26)	135.3	161.4	173.0	−11.6
Alphabet (to Jun 26)	244.2	185.7	132.4	53.3
Meta (to Jun 26)	68.1	130.3	89.3	41.0
Microsoft (FY to Jun 26)	133.8	182.9	116.0	67.0
Oracle (to Aug 26)	18.9	46.9	75.7	−28.7
Total	**600.3**	**707.2**	**586.4**	**120.9**

Source: SEC XBRL data (Forms 10-Q and 10-K); Andra Markets calculation. Free cash flow is operating cash flow minus purchases of property and equipment and excludes finance leases. Amazon's own definition, net of proceeds from asset sales and incentives, gives about −$7.6bn.

9.0 Historical Parallels

Telecoms, not WorldCom

The comparison most often made is with the telecom and fibre boom of 1997 to 2001. Telecom capital spending more than doubled between 1997 and 2000 (estimates of the peak range from about US$83 billion to US$120 billion), equipment suppliers such as Lucent and Nortel financed their own customers, and when demand fell short Lucent lost US$16.2 billion in 2001 and Nortel's value fell from US$136 billion to US$14 billion. The famous accounting scandal of that era, WorldCom, is a different matter: it booked more than US$3.8 billion of ordinary operating costs as capital spending, which was fraud. Today's capitalisation of servers is legitimate; the parallel is with the telecom cycle itself, in which real profits at suppliers were followed by write-downs at buyers once the useful life of the capacity proved shorter than expected. In scale, AI capital spending is already about 1.2 to 1.3 per cent of US GDP, above the telecom peak of about 1 per cent, though still below the railway booms of the nineteenth century.

Exhibit D. AI build-out vs the telecom boom

Feature	Telecom 1997 to 2001	AI 2023 to 2026
Capex peak (% of US GDP)	~1%	~1.2 to 1.3% (2025)
Supplier profit surge	Lucent 1999: $4.8bn profit	Nvidia FY26: $120bn net income
Vendor financing	Lucent, Nortel lent to carriers	Nvidia, Broadcom, AMD invest in and back customers
Accounting	WorldCom: $3.8bn costs capitalised (fraud)	legal capitalisation; lives extended
Demand	fibre capacity far exceeded traffic	contracted backlogs ~$2.35tn; revenue still small vs capex
Aftermath	Lucent −$16.2bn (2001); Nortel −90%	to be determined

Sources: American Affairs (2020); Congressional Research Service (RS21253); Bloomberg via Advisor Perspectives (2025); company filings.

10.0 Implications for Investors

What to adjust, what to watch

Three practical conclusions follow. First, valuations based on earnings flatter the AI spenders relative to valuations based on cash: the forward price-to-earnings ratio of the S&P 500 is 19.1, but on free cash flow the largest spenders look much more expensive than their earnings suggest. Second, earnings growth will slow mechanically in 2027 and 2028 as depreciation rises, even if revenue keeps growing; consensus expects S&P 500 earnings growth to fall from 32.7 per cent this year to 15.8 per cent in 2027, and Goldman's estimate of a 5-point

depreciation drag suggests that forecast could still be optimistic. Third, the suppliers' profits are real but cyclical: they depend on spending continuing to grow, not merely to remain high. A flat capex year would cut supplier earnings while buyers' depreciation keeps rising, which is the combination that made the telecom bust so severe.

Exhibit E. Earnings-quality indicators to monitor

Indicator	Latest	Warning sign
Capex / depreciation (top 5)	~3.5x (trailing)	falls as depreciation catches up
Free cash flow vs net income	$121bn vs $600bn	more companies with negative FCF
Construction in progress / gross PP&E	13 to 32%	rising: future depreciation building
Useful-life changes	Amazon shortened (2025)	other companies follow
Unrealised gains in EPS	~12% of Q2 EPS	turn into losses
GPU rental prices	H100 ~$2.65/hr index	continued falls: shorter economic life
Supplier orders	Dell AI backlog $95bn; GE Vernova $176bn	order growth slows

Andra Markets selection; latest values from this report's sources (Dell from a secondary summary).

The third-quarter results of Alphabet, Microsoft, Meta and Amazon at the end of October will update all of these indicators: watch depreciation growth, any change in useful lives, the size of investment gains, and capital spending guidance for 2027. Nvidia's results, expected around 18 November, will show whether supplier profits are still accelerating.

verdict: the accounting is legal and the rally is real, but profits are front-loaded. watch cash, not EPS

Methodology and Source Notes

Method. This report tests a common claim about AI-era accounting using company filings (Forms 10-K, 10-Q and 8-K) and published research. The depreciation sensitivity is an Andra Markets model with simplified assumptions (62% of capex in servers, 20-year life for buildings and power, half-year convention, 17% tax rate, no construction lag) and is illustrative. S&P 500 aggregate earnings are estimated from FactSet bottom-up EPS. Andra Markets uses AI tools made by Anthropic, whose stake valuations are discussed in this report. Handwritten notes and highlights are informal commentary by the desk. This note is analysis, not investment advice.

Principal sources. SEC EDGAR filings of Microsoft, Alphabet, Amazon, Meta, Oracle, CoreWeave, Nvidia, Broadcom, Micron, Arista, Super Micro, Vertiv, GE Vernova, Eaton and AMD (2021 to 2026); Microsoft Q4 FY26 earnings call (29 July 2026); Alphabet Q4 2025 earnings call; FactSet Earnings Insight (9 October 2026) and FactSet Insight (2026); Goldman Sachs (23 September 2026) and via TKer and TechTimes; LSEG via Reuters/Kitco (August 2026); JPMorgan Eye on the Market (8 October 2026); JPMAM (10 September 2026); Silicon Data; Ornn; AWS; Michael Burry via StreetInsider; Jim Chanos via Benzinga/Yahoo; Bernstein via Tiger; Congressional Research Service; American Affairs; Advisor Perspectives.

Disclaimer

This document is produced by **Andra Markets** Global Markets Research for informational and educational purposes only. It is not investment, legal, accounting or tax advice, and it is not an offer or solicitation to buy or sell any security or financial instrument. Figures are drawn from public sources believed to be reliable but are not independently verified. Multiples are illustrative calculations; analyst views are reported for reference only. Views are those of the desk at publication and may change without notice. Past performance is not indicative of future results. Andra Markets accepts no liability for any loss arising from the use of this material.