

Research Analysts

Pierre Ferragu

+1 646 681 4616
pierre@newstreetresearch.com

Antoine Chkaiban

+44 20 7375 9126
antoine@newstreetresearch.com

Benjamin Harwood

+1 929 379 4139
ben@newstreetresearch.com

Derek Mogull

+1 929 297 6941
derek@newstreetresearch.com

Asia Tech Specialist

Peter Vogel

+1 646 703 3211
peter@newstreetresearch.com

Sales

Heather Broffman

heather@newstreetresearch.com

Ethan Lacy

ethan.lacy@newstreetresearch.com

Daniel Gilroy

daniel.gilroy@newstreetresearch.com

Michael Chambers

michael@newstreetresearch.com

Dimitri Livchits

dimitri@newstreetresearch.com

Mark Franks

mark@newstreetresearch.com

Hyperscale & Cloud: continued capex revisions; visibility improving materially for 2027 & 2028

23 March 2026

We publish today the Hyperscale & Cloud section of our 30th Tech Infrastructure Quarterly Bible.

In this fast-paced world, it is important to regularly review extensively what happened across our coverage to update our convictions and strengthen our expertise. Getting fundamental and secular views is great, but often easy. Good research is also about thoroughly reviewing short-term industry dynamics and investor sentiment.

We publish the Bible section by section. For each segment, we wrap up the current context (recent developments and key controversies), what we have learnt in the quarter with

underlying data and analyses, a summary of what we expect next, and our updated convictions. In the coming weeks we will publish sections on, Enterprise IT, Foundries, Telecom Infrastructure, Industrials and PCs.

See the following links for the [Smartphones](#), [Ride-sharing & Delivery](#), [Memory](#), [Semicap](#), [Industrial](#), [Telecom Infrastructure](#), [Autos](#), and [PCs](#) sections.

There is monumental work behind this report. For each segment, we track all company comments and key reported numbers, industry data, and the evolution of consensus expectations. All the work is accessible in excel spreadsheets on request.

Cloud Infrastructure – The context

After surging 70% in 2018, cloud capex¹ flattened in 2019 with inventory & capacity absorption. Between 2020 and 2022, spending re-accelerated at a 25% CAGR, driven by the infrastructure demands of remote work. This was followed by another period of consolidation in 2023, with flattish capex. Growth resumed, with capex up 55% for 2024 and expected up another 35% in 2025. In 2Q25, hints emerged of sustained spending momentum into 2026.

Cloud capex is inherently lumpy, influenced by a small number of buyers and large-scale deployments.

Spending has increasingly shifted toward technology, and in particular AI clusters, with Generative AI reaching an inflection point in 2023. Demand for general-purpose infrastructure weakened, with slow product ramps & delays at Intel, a deceleration in public cloud revenue growth, and customers prioritizing GPUs. General-purpose infrastructure spending started to slowly recover over back half of 2024, though remains subdued compared to explosive growth in AI-related infrastructure .

As earnings season approached, investor focus centered on capex visibility for this year and next, how that translates into growth for the various XPU's, and the rollout of Ethernet for AI.

¹ = Top 7 Hyperscalers: Google, Amazon, Microsoft, Meta, Baidu, Alibaba, Tencent.

Cloud Infrastructure – What have we learnt this quarter?

Continued Capex revisions.

Hyperscale revenues¹ accelerated to 16% YoY, 14% QoQ. Public cloud growth in the mid-30% with material AI contribution. Capex² grew 61% YoY, 21% QoQ. Strong 2026 guides lift expectations for 2026/27 by 28/26%, to 58%/10% YoY growth, vs. 63% in 2025 and 55% in 2024. Compute demand to exceed capacity through 2026.

Replatforming healthily for Microsoft Amazon & Google, who absorb increased capital intensity with cash costs growing slower than revenues. Meta is in overdrive, squeezing margins.

DC compute: Strong visibility on 2027 already.

DC compute up 20% QoQ and guided up another 15% QoQ in 1Q26. Nvidia already has visibility on >\$1tn Blackwell/Rubin from 4Q25 to the end of 2027 and Broadcom guided AI semis to >\$100bn in FY27 (Nov-27), up >90% YoY. AMD DC GPUs set to grow mid-teens QoQ excluding China.

Memory: capacity squeeze leads to record margins & increased capex

All vendors are sold out for 2026; Hynix now signaling multi-year constraints and increasing further Capex. As we expected, AI demand growth has squeezed capacity across the market driving prices and gross margins to record levels.

Networking: Constrained; ethernet ramping; optics visibility into 2028

AI Ethernet ramping, with 3/4 large Arista-powered AI programs growing beyond 100k GPU clusters, 4th to reach 100k this year. Cisco guided FY26 (Jul-26) AI orders up 2x to >\$5bn. Nvidia networking up 3.5x, with strong Spectrum-X momentum. AI optics constrained, visibility now stretching into 2028.

General purpose infrastructure benefitting from AI, also constrained

Server CPU revenues up mid-teens QoQ at both AMD and Intel, with AMD taking 40bps of share. Both mentioned agentic AI workloads driving CPU demand, and capacity constraints.

1 = Top 7 Hyperscale companies: Google, Amazon, Microsoft, Meta, Baidu, Alibaba, Tencent

2 = Quarterly capex numbers exclude Baidu and Tencent

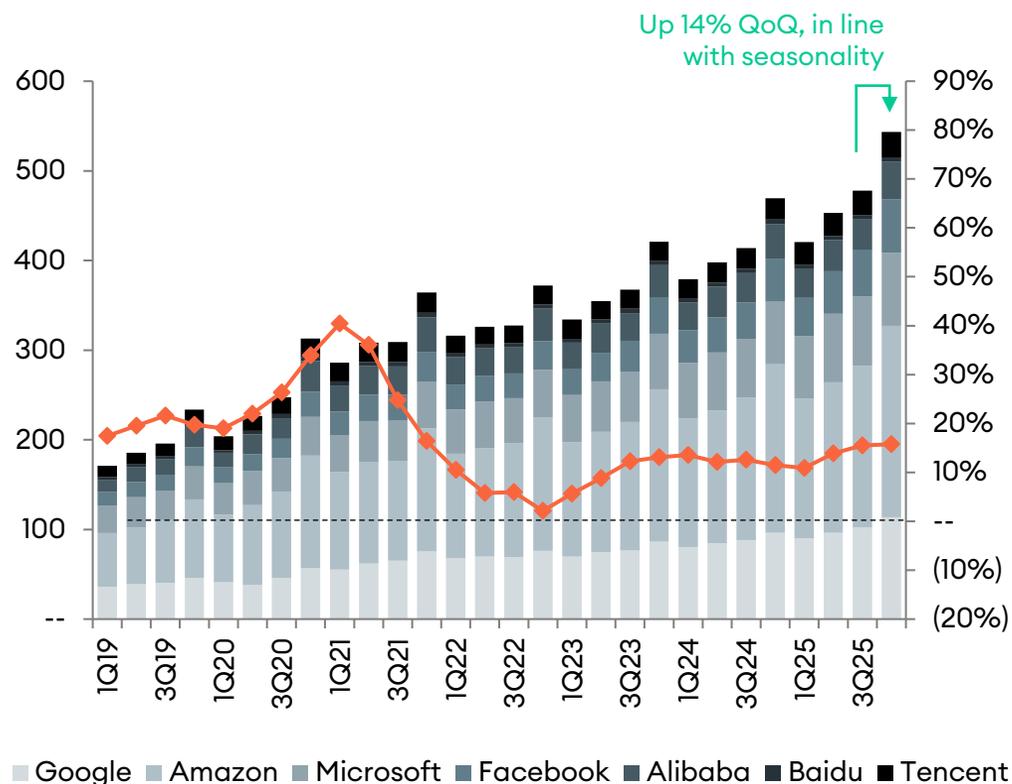
3 = Intel DCAI, Nvidia Datacenter, AMD Datacenter, Marvell Datacenter.

Hyperscale revenue & capex grew 16% & 61% YoY, respectively

Revenue growth stable at +16% YoY

LHS: Top 7 hyperscalers¹ revenues (US\$bn)

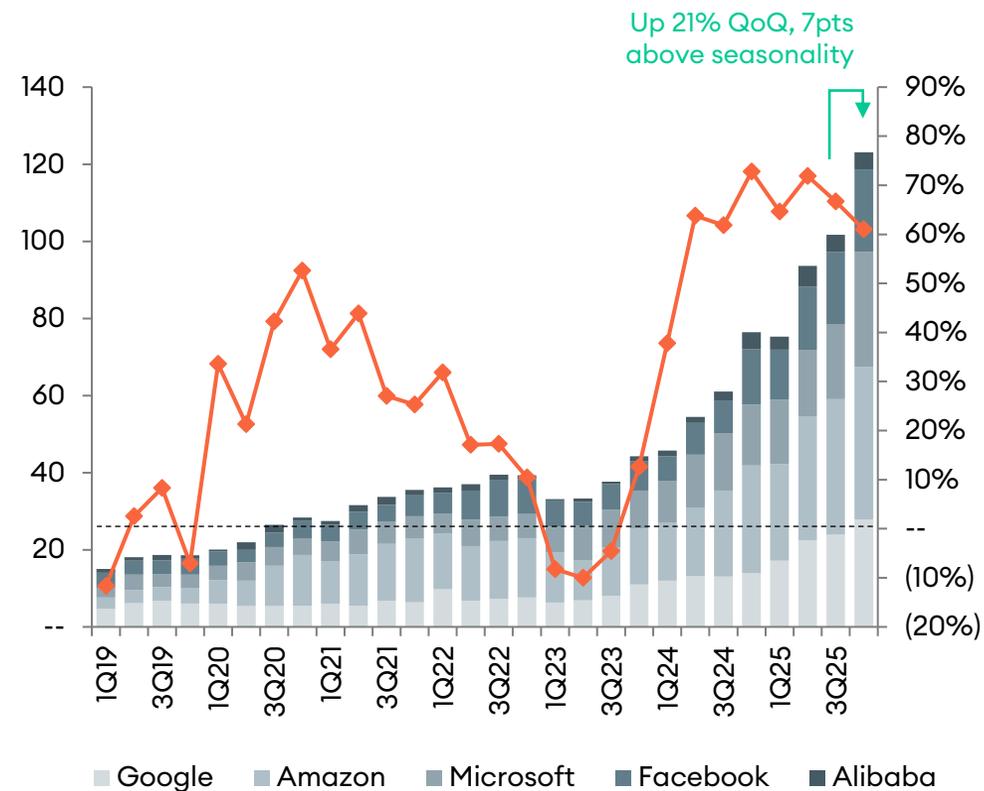
RHS: YoY Growth (%)



Capex decelerated slightly to +61% YoY

LHS: Top 5 hyperscalers² Capex (US\$bn)

RHS: YoY Growth (%)



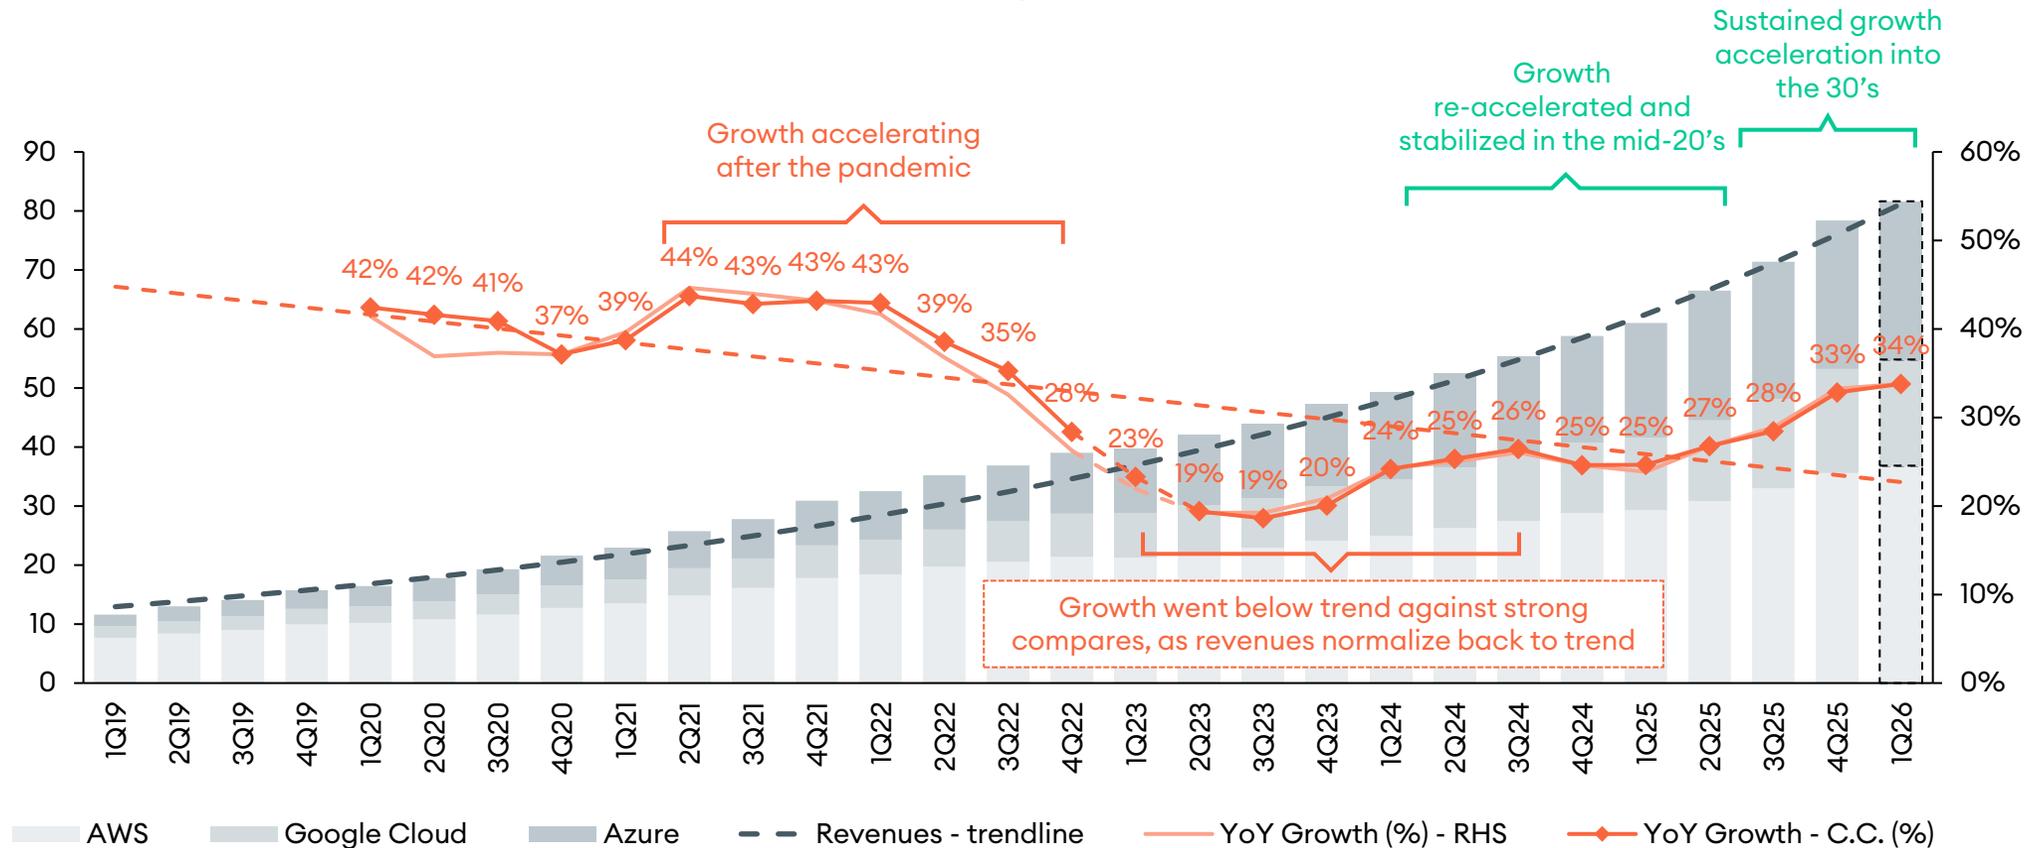
Source: Corporate reports and NSR analysis.

1 = Top 7 Hyperscale Companies: Google, Amazon, Microsoft, Meta, Baidu, Alibaba, Tencent.
2 = Quarterly capex numbers exclude Baidu and Tencent.

Public cloud growth accelerating: +33% YoY in 4Q25; expected +34% in 1Q26

Growth sustaining at 47% (in C.C.) at GCP, 38% at Azure, and accelerating 2pts at AWS from to 26%

Public cloud revenues, revenues trendline (LHS, US\$bn), and YoY growth (RHS, %)

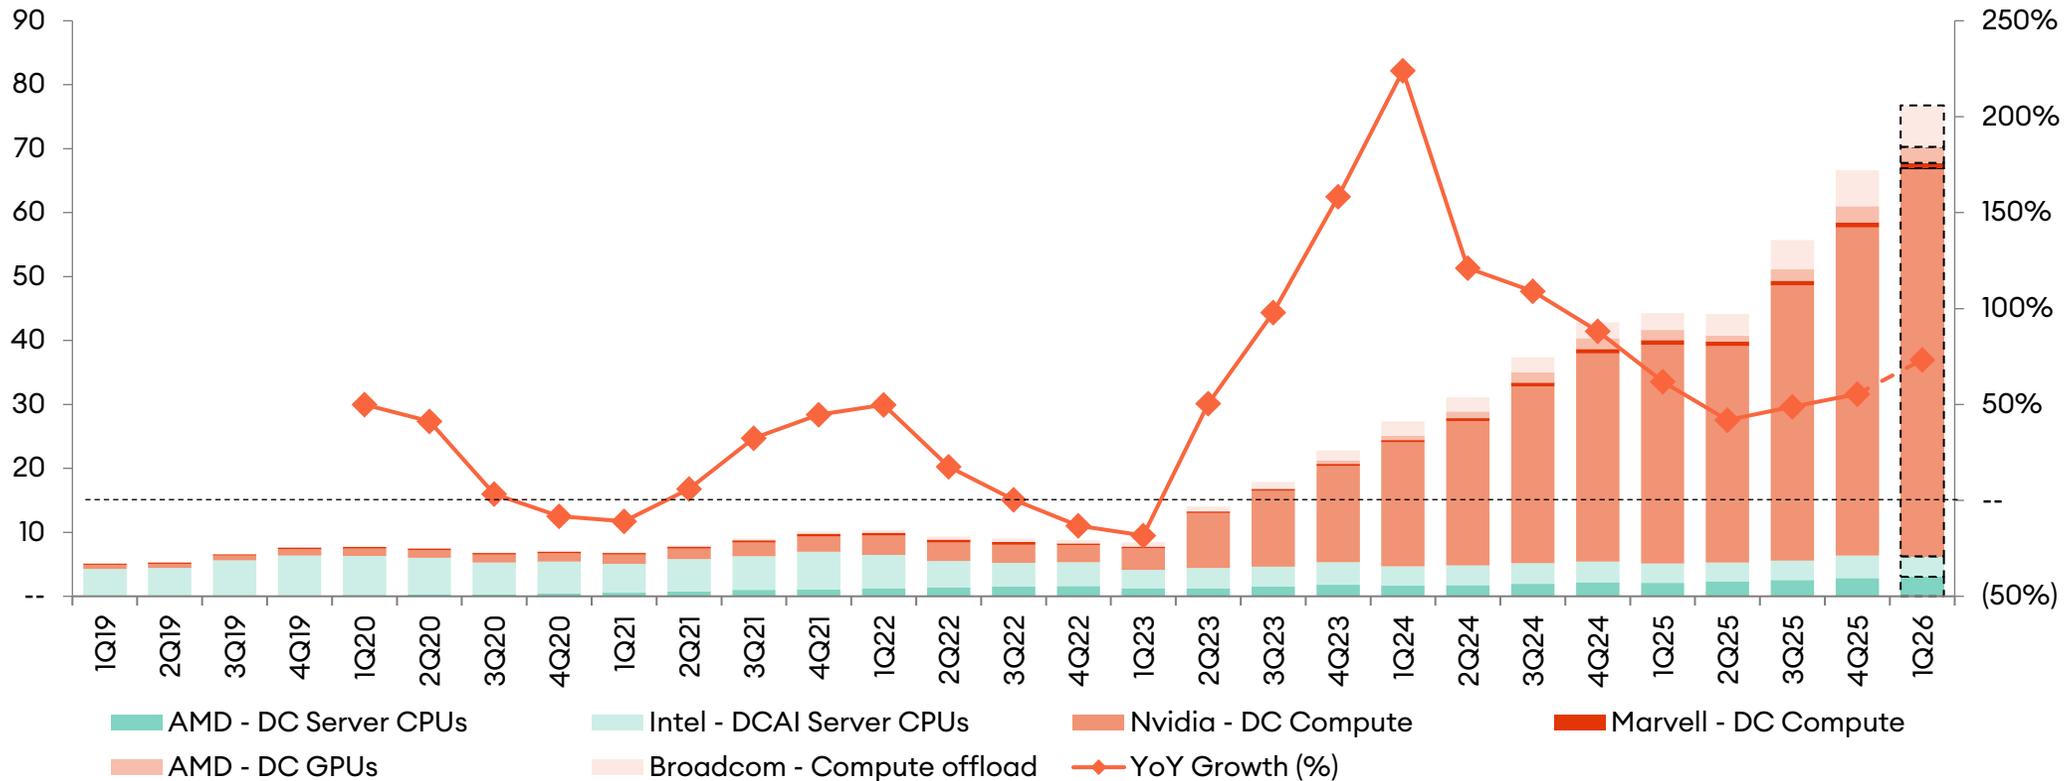


Source: FactSet (post-event consensus), Corporate Reports and NSR estimates and analysis.

DC compute up 55% YoY (+20% QoQ) in 4Q25, expected +73% YoY (+15% QoQ) in 1Q26

XPUs up 61% in 4Q25 (+20% QoQ) and CPUs up 17% (+14% QoQ); XPUs up 80% in 1Q26 (+17% QoQ) and CPUs up 21% (-2% QoQ)

LHS: Datacenter semi revenues (US\$bn) RHS: YoY Growth (%)

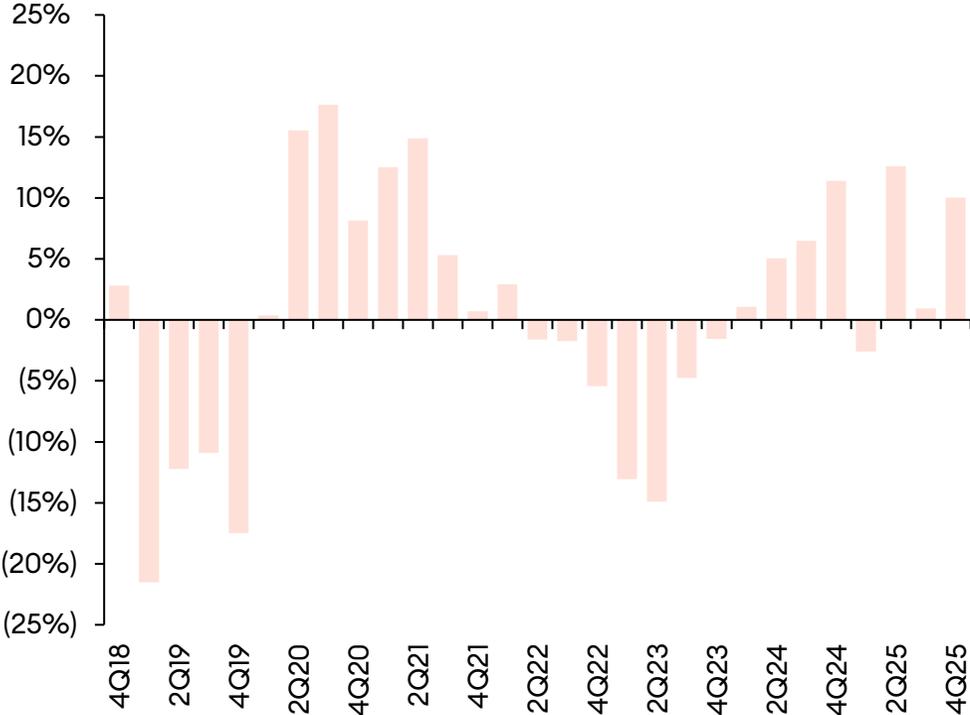


Source: Corporate reports and NSR analysis.

Capex 10% above consensus expectations, datacenter compute 3% above

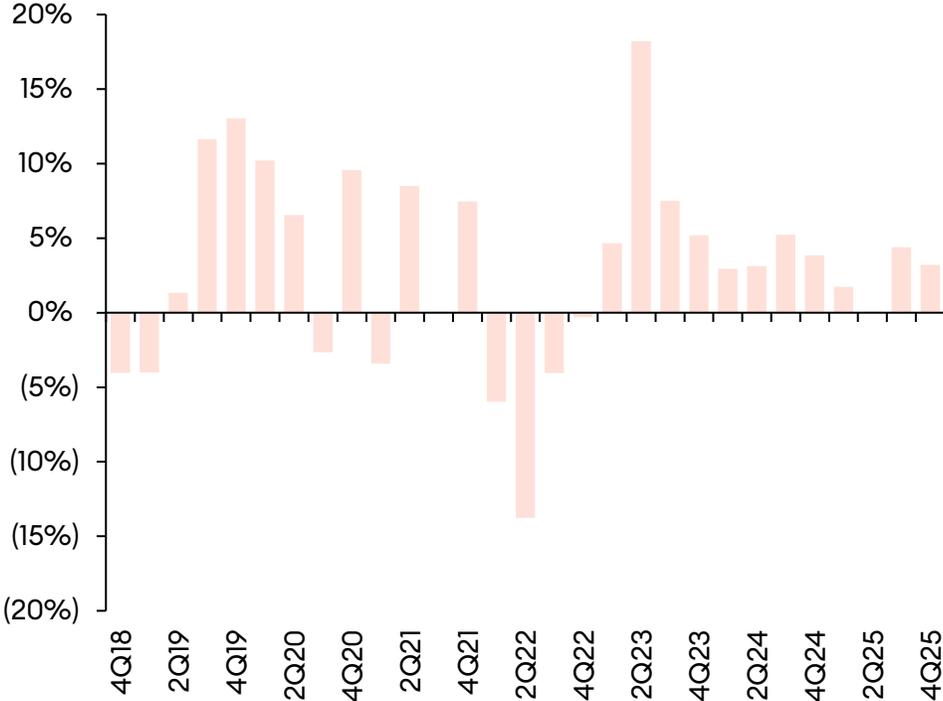
Hyperscale capex 10% above consensus expectations

Hyperscale¹ capex vs. consensus expectations (%)



DC compute semi revenues 3% above expectations

DC semi² revenues vs. consensus expectations (%)



1 = Top 7 Hyperscale companies: Google, Amazon, Microsoft, Meta, Baidu, Alibaba, Tencent. Quarterly capex numbers exclude Baidu and Tencent.

2 = Intel DCAI, Nvidia Datacenter, AMD Datacenter, Marvell Datacenter

Source: Corporate reports and NSR analysis.

Cloud Revenues: Strength across both AI and traditional workloads

Microsoft: traditional infra above expectations, AI in line

*“Results were ahead of expectations, driven by **better than expected growth in our core infrastructure business**, primarily from our largest customers. **Azure AI services revenue was generally in line with expectations** and this quarter demand again exceeded supply across workloads”*

Microsoft 2FQ26 Earnings Call – 28 January 2026

Amazon: Acceleration across traditional & AI workloads

*“We added \$2.6 billion in quarter-over-quarter revenue and AWS now has an annualized revenue run rate of \$142 billion. This **acceleration was driven by both core and AI services as customers** continue to modernize their infrastructure and migrate workloads to the cloud.”*

*“Some of it is for our core workloads, **which are non-AI workloads because they’re growing at a faster rate than we anticipated**. But most of it is in AI.”*

Andrew Jassy, CEO, Amazon – 5 February 2026

Google: Enterprise AI accelerating, traditional strong

*“GCP’s performance was driven by **accelerating growth in enterprise AI products**, which are generating billions in quarterly revenues”*

*“**Core GCP was also a meaningful contributor to growth** due to strong demand for infrastructure and other services such as cybersecurity and data analytics.”*

Ruth Porat, CFO, Google – 4 February 2026

Source: Company Reports, Press Reports and NSR analysis.

Cloud Revenues: Constraints gating AI growth, capacity rapidly ramping

Microsoft: rapidly adding capacity but constraints persist

*“Our customer **demand continues to exceed our supply** [...] we’re working as hard as we can to add capacity as quickly as we can [...] All up, we **added nearly 1 gigawatt of total capacity this quarter alone.**”*

*“And while we’re accelerating the amount of capacity we’re bringing online, **we will continue to balance Azure revenue growth with the growing needs across our first-party apps and AI solutions, our own R&D efforts and the end-of-life server replacements.** Therefore, we now expect to be capacity constrained through at least the end of our fiscal year.”*

Microsoft 2FQ26 Earnings Call – 28 January 2026

Amazon: Capacity to double by end of 2027.

*“Yet I think every provider would tell you, including us that **we could actually grow faster if we had all the supply that we could take.**”*

*“If you look, in the **last 12 months, we added 3.9 gigawatts of power.** Just for perspective, **that's twice what we had in 2022, when we were in \$80 billion annual run rate business.** **We expect to double it again by the end of 2027.** We **added 1.2 gigawatts of power in Q4, just quarter-over-quarter.**”*

Andrew Jassy, CEO, Amazon – 5 February 2026

Google: Capacity constraints through 2026.

*“We’ve been supply constrained even as we’ve been ramping up our capacity [...] I do **expect to go through the year in a supply constrained way.**”*

Sundar Pichai, CEO, Google – 4 February 2026

Source: Company Reports, Press Reports and NSR analysis.

Cloud Capex: 2026 guided up across the board, short vs. long lived assets mix stable

Meta: 2026 capex guided to grow 73% YoY at mid-point

“We anticipate 2026 capital expenditures, including principal payments on finance leases to be in the range of \$115 billion to \$135 billion, with year-over-year growth driven by increased investment to support our Meta Superintelligence Labs efforts and core business.”

Susan Li, CFO, Meta – 28 January 2026

Google: 2026 capex guided to double, Asset mix stable

“The vast majority of our [4Q25] CapEx was invested in technical infrastructure, with approximately 60% of that investment in servers and 40% in data centers and networking equipment. “

“For the full year 2026, we expect CapEx to be in the range of \$175 billion to \$185 billion, with investments ramping over the course of the year.”

Anat Ashkenazi, CFO, Alphabet – 4 February 2026

Microsoft: Asset mix stable

“Capital expenditures were \$37.5 billion, and this quarter roughly two-thirds of our CapEx was on short-lived assets, primarily GPUs and CPUs. [...] As we work to close the gap between demand and supply, we expect the mix of short-lived assets to remain similar to Q2.”

Amy Hood, CFO, Microsoft – 28 January 2025

Amazon: 2026 capex guided to grow 56% YoY, driven by AWS

“We expect to invest about \$200 billion in capital expenditures across Amazon, but predominantly in AWS, because we have very high demand, customers really want AWS for core and AI workloads, and we’re monetizing capacity as fast as we can install it.”

Andrew Jassy, CEO, Amazon – 5 February 2026

Source: Company Reports, Press Reports and NSR analysis.

AI drives productivity gains at hyperscalers

Google: AI drives productivity and efficiency gains

“Given the increase in our CapEx investments in recent years, we expect the growth rate in 2026 depreciation to accelerate in Q1 and meaningfully increase for the full year.”

“when we invest at these amounts, we look at how we can ensure that we are the most efficient with every dollar that goes towards our technical infrastructure [...] using AI within the business to drive daily operations, it can be all the way from the engineering team to small teams within our back office.”

Anat Ashkenazi, CFO, Alphabet – 4 February 2026

Amazon: Offsetting depreciation with efficiencies & cost reduction

“It certainly has a headwind from the investments in AI and the depreciation on that CapEx. But we also work very hard to offset that with efficiencies and cost reductions.”

Brian Olsavsky, CEO, Amazon – 5 February 2026

Microsoft: Azure efficiency gains

“Gross margin percentage decreased year-over-year, driven by continued investments in AI [...] partially offset by efficiency gains in Azure.”

Amy Hood, CFO, Microsoft – 29 January 2026

Meta: Optimizing workflows using AI tooling

“Since the beginning of 2025, we've seen a 30% increase in output per engineer with the majority of that growth coming from the adoption of agentic coding”

“We're investing in AI-native tooling, so individuals at Meta can get more done [...] projects that used to require big teams now be accomplished by a single very talented person.”

Mark Zuckerberg, CEO, Meta – 28 January 2026

Source: Company Reports, Press Reports and NSR analysis.

Compute – AI: Agentic AI inflection extends visibility on demand

Nvidia: Agentic AI inflecting, demand visibility elongating

*“We have now seen the **inflection of agentic AI** and the usefulness of agents across the world in enterprises everywhere. You're seeing **incredible compute demand** because of it.”*

*“purchase commitments also increased significantly, as we have strategically secured inventory and capacity to meet demand beyond the next several quarters. This is further out in time than usual and **reflects the longer demand visibility we have**”*

*“we have strong confidence in visibility **of \$1 trillion plus of Blackwell plus Rubin [through end of 2027]**”*

Nvidia 4FQ26 Earnings Call & GTC – 25 February & 17 March 2026

AMD: Entering multi-year demand super cycle for AI compute

*“We are **entering a multiyear demand super cycle** for high-performance and AI computing”*

*“In addition to our multi-generation **partnership with OpenAI to deploy 6GW of Instinct GPUs**, we are in **active discussions with other customers on at-scale multiyear deployments** starting with Helios and MI450 later this year.”*

*“Based on the strength of our EPYC and Instinct road maps, we are **well positioned to grow Data Center segment revenue by more than 60% annually over the next three to five years** and scale our AI business to tens of billions in annual revenue in 2027.”*

Lisa Su, CEO, AMD – 3 February 2026

Broadcom: FY27 (Nov-27) AI semis guided to >\$100bn, >90% YoY

*“Today, in fact, we have **line of sight to achieve AI revenue from chips, just chips, in excess of \$100 billion [in FY27] [...]** we are saying getting **close to 10 gigawatts.**”*

Hock Tan, CEO, Broadcom – 4 March 2026

Source: Corporate reports and NSR analysis.

Compute – General Purpose: Demand rebounds at last, leading to supply constraints

Intel: AI workloads driving CPU demand, and significant constraints; DCAI growth in 2026

“The continuing proliferation and diversification of AI workloads is placing significant capacity constraints on traditional and new hardware infrastructure, reinforcing the growing and essential role CPUs play in the AI era.”

“Demand for traditional servers continue to be very strong, and we are focused on ramping available capacity to support the meaningful uptick we are seeing

“[...] likelihood of a strong year of growth for DCAI”

Intel 4Q25 Earnings Call – 23 January 2026

AMD: Agentic AI driving surging CPU demand; increasing supply; multi-year CPU cycle.

“Demand for EPYC CPUs is surging as agentic and emerging AI workloads require high-performance CPUs to power head nodes and run parallel tasks alongside GPUs.”

“[CPU] demand continues to be strong... we’re working with our supply chain partners to increase supply”

“we think the overall server CPU TAM is going to grow, let's call it, strong double digits in 2026, just given the – as we said, the relationship between CPU demand and overall AI ramp”

“really this multiyear CPU cycle, and we continue to see that.”

Lisa Su, CEO, AMD – 3 February 2026

Source: Corporate reports and NSR analysis.

Memory – AI: 2026 HBM sold out, supply constraints extending to 2027 already

Micron: 2026 HBM sold out

“our HBM for 2026 is sold-out in terms of volume.”

Sanjay Mehrotra, CEO, Micron – 17 December 2025

Hynix: HBM multi-year contracts hard to accommodate with constraints

“Even as we maximize production, we cannot meet HBM demand 100%”

“Customers now prefer multiyear contracts [for HBM], but capacity constraints make it difficult to accommodate all customer requests”

Song Hyun Jong, President, SK Hynix – 28 January 2026

Samsung: Customers seeking HBM supply orders for 2027 and beyond

“All production-ready capacity is currently fully booked with customer POs.”

“Major customer demand for HBM in 2026 still exceeds available supply from us.”

“For volumes in 2027 and beyond, major customers are seeking to finalize supply discussions as soon as possible to secure supply.”

Jaejune Kim, EVP Memory, Samsung – 28 January 2026

Source: Corporate Reports and NSR analysis.

Memory – Others: Strong demand & capacity constraints for the foreseeable future

Hynix: Conventional server DRAM to grow above overall market in 2026

“Demand increased significantly, not only for HBM, but also for conventional server memory [...] to support AI workloads, requirements for memory density and bandwidth in general purpose servers are rapidly increasing [...] demand for server DRAM and enterprise SSDs is expected to grow structurally at a pace well above the overall market growth [in 2026]”

Song Hyun Jong, President, SK Hynix – 28 January 2026

Samsung: Significant memory shortage for the time being

“We expect a significant shortage of supply relative to demand to continue across all product categories, whether it's HBM, conventional DRAM or NAND with tight undersupply conditions expected for the time being.”

Jaejune Kim, EVP Memory, Samsung – 28 January 2026

Micron: Memory supply to remain tight for foreseeable future

“2026 supply on HBM will be tight and non-HBM DRAM will be tight as well.”

“Over the last few months, our customers' AI data center build-out plans have driven a sharp increase in demand forecasts for memory and storage [...] We believe that the aggregate industry supply will remain substantially short of the demand for the foreseeable future”

Micron 1FQ26 Earnings Call – 17 December 2025

Source: Corporate reports and NSR analysis.

Networking – AI: Ethernet ramping at Arista, Cisco and Nvidia

Arista: 4 AI titan ramping beyond 100k GPU clusters; 2027 outlook strengthening

“Three of [AI titan ethernet customers] have already deployed a cumulative of 100k GPUs and are now growing from there[...] Our fourth customer [...] fully expect them to get there this year”

“first thing that happens in the CapEx is they got to build the data centers and the – get the power and get all of the GPUs and accelerators, and the network comes – lags a little. So demand is going to be very good. But whether the shipments exactly fall into 2026 or 2027”

Jayshree Ullal, CEO, Arista – 12 February 2026

Cisco: AI orders guided up 2x to >\$5bn in FY26 (Jul-26)

“Given the strong demand for our Silicon One systems and optics, we now expect to take AI orders in excess of \$5 billion and to recognize over \$3 billion in AI infrastructure revenue from hyperscalers in FY 2026.”

Chuck Robbins, CEO, Cisco - 11 February 2026

Nvidia: Networking up 3.5x YoY to \$11bn; strong Spectrum-X momentum

“Networking [...] was a standout this quarter, generating \$11bn in revenue, up more than 3.5x YoY. Demand for our scale up and scale out technologies reached record levels, both growing double digit sequentially, driven by strong adoption of NVLink, Spectrum-X Ethernet and InfiniBand.”

“Momentum is strong with our Spectrum-X Ethernet scale up and scale across networking, as customers work to unify distributed data centers into integrated giga scale AI factories.

Colette Kress, CFO, Nvidia – 25 February 2026

Source: Corporate reports and NSR analysis.

Networking – Others: Solid order growth in traditional networking

Arista: AI clusters driving front-end cloud infra upgrades.

*“I’m actually telling you we might do, oh, north of \$3 billion [AI networking revenue] this year and growing, right? **That number definitely includes the frontend** as it’s tied to the backend GPU clusters, and it’s an all Ethernet, all AI system for agentic AI applications.”*

Jayshree Ullal, CEO, Arista - 12 February 2025

Cisco: Enterprise product orders up 40% QoQ on broad-based strength

*“Excluding hyperscalers, product orders were up 10% year over year, demonstrating the broad-based demand we see for our technology globally. **Enterprise product orders [accelerated 4pts] up 8% year over year in Q2 [up 40% QoQ]**, with strength across our entire networking portfolio.”*

Chuck Robbins, CEO, Cisco – 11 February 2026

Source: Corporate reports and NSR analysis.

Networking – Optics: Visibility on demand and capacity constraints extending to 2028

Marvell: FY27 (Jan-27) interconnect guide increased 20pts to 50% YoY

*“we now see our fiscal 2027 DC revenue growing by 40% year-over-year. We expect all our key product lines in DC to be stronger than our prior outlook. **Notably, we expect our interconnect business to grow more than 50% year-over-year, well above our prior expectation of 30% growth.**”*

Matthew Murphy, CEO, Marvell – 5 March 2026

Coherent: 2026 booked out; 2027 filling up; constraints through at least next year; visibility into 2028

*“So, number one, **bookings being booked out through the rest of this calendar year [...] calendar 2027 is filling very, very quickly [...]** getting really good, detailed long-term forecasts from our big customers [...] **we're getting forecasts that go out into calendar 2028**, which again is great for visibility.”*

*“I don't foresee the supply/demand getting back in balance this calendar year. I **don't think it happens next calendar year.**”*

James Anderson, CEO, Coherent – 4 February 2025

Lumentum: Optics demand far outstripping supply, driven by AI.

*“The strain on our fab is high, right. **We're very much sold out in our high power laser fab [...]** Even as **we've added this 20% additional capacity, the demand/supply imbalance has increased.**”*

*“**1.6T is definitely stronger** than we felt 90 days ago.”*

*“**as we enter calendar year 2027, it should go up from there** in terms of what we see in our backlog and in terms of our revenue.”*

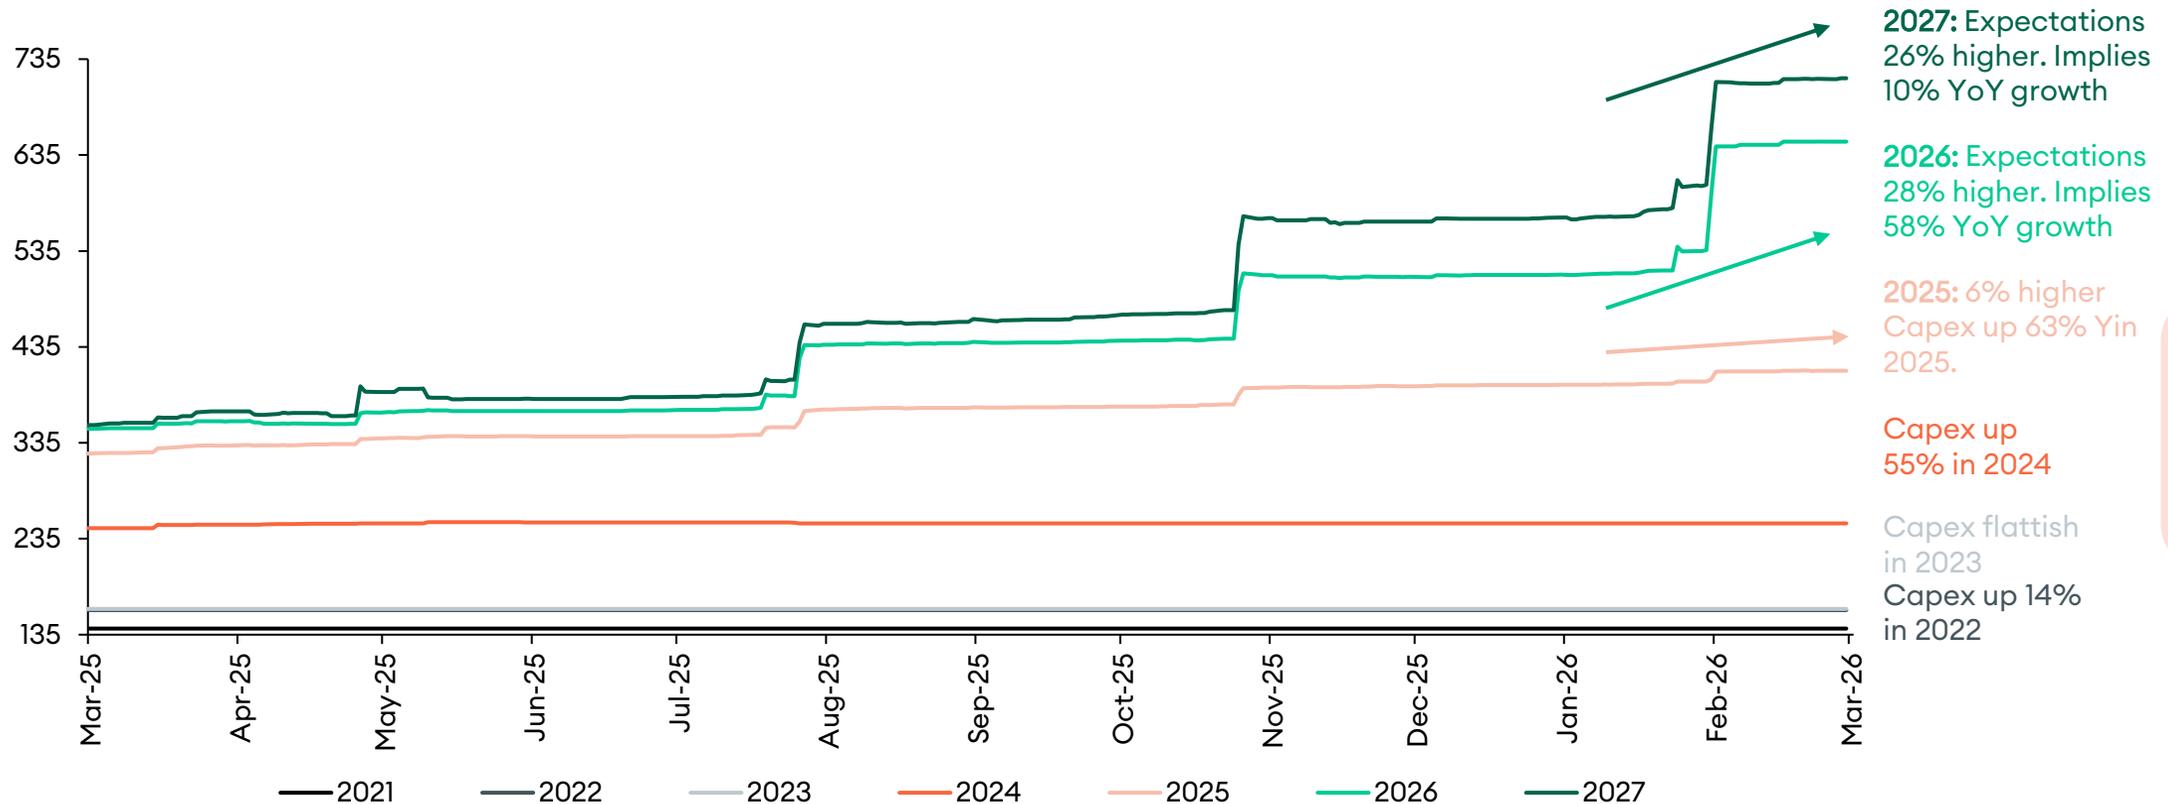
Michael Hurlston, CEO, Lumentum, 3 February 2026

Source: Corporate reports and NSR analysis.

2026/27 expectations for hyperscale capex increased 28/26% this earnings season

Expectations now imply 58/10% growth in 2026/27, on the back of 63% growth in 2025 and 55% in 2024

Top 7 Hyperscale¹ Group Capex (US\$bn)

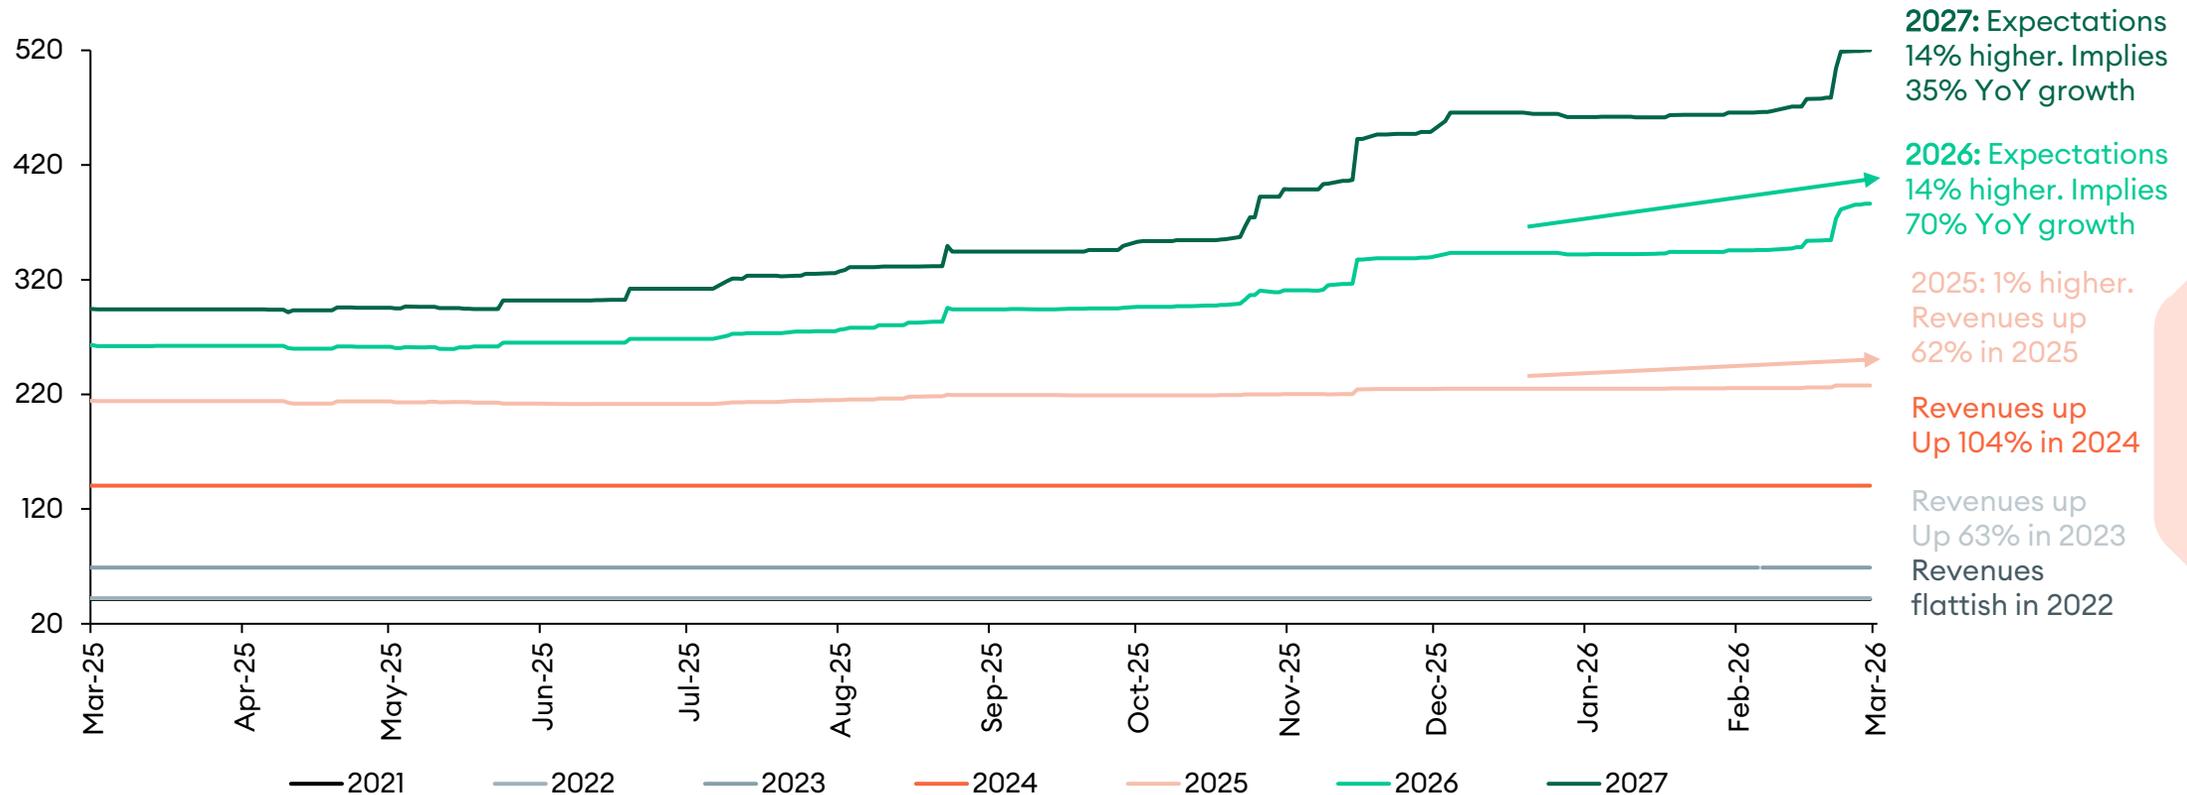


1 = Top 7 Hyperscale Companies: Google, Amazon, Microsoft, Baidu, Alibaba, Tencent
Source: Corporate reports and NSR analysis.

2026/27 expectations for DC semis up 14% this earnings season

Expectations now imply 70/35% growth in 2026/27, on the back of 62% growth in 2025 and 104% in 2024

Datcenter semi revenues¹ (US\$bn)

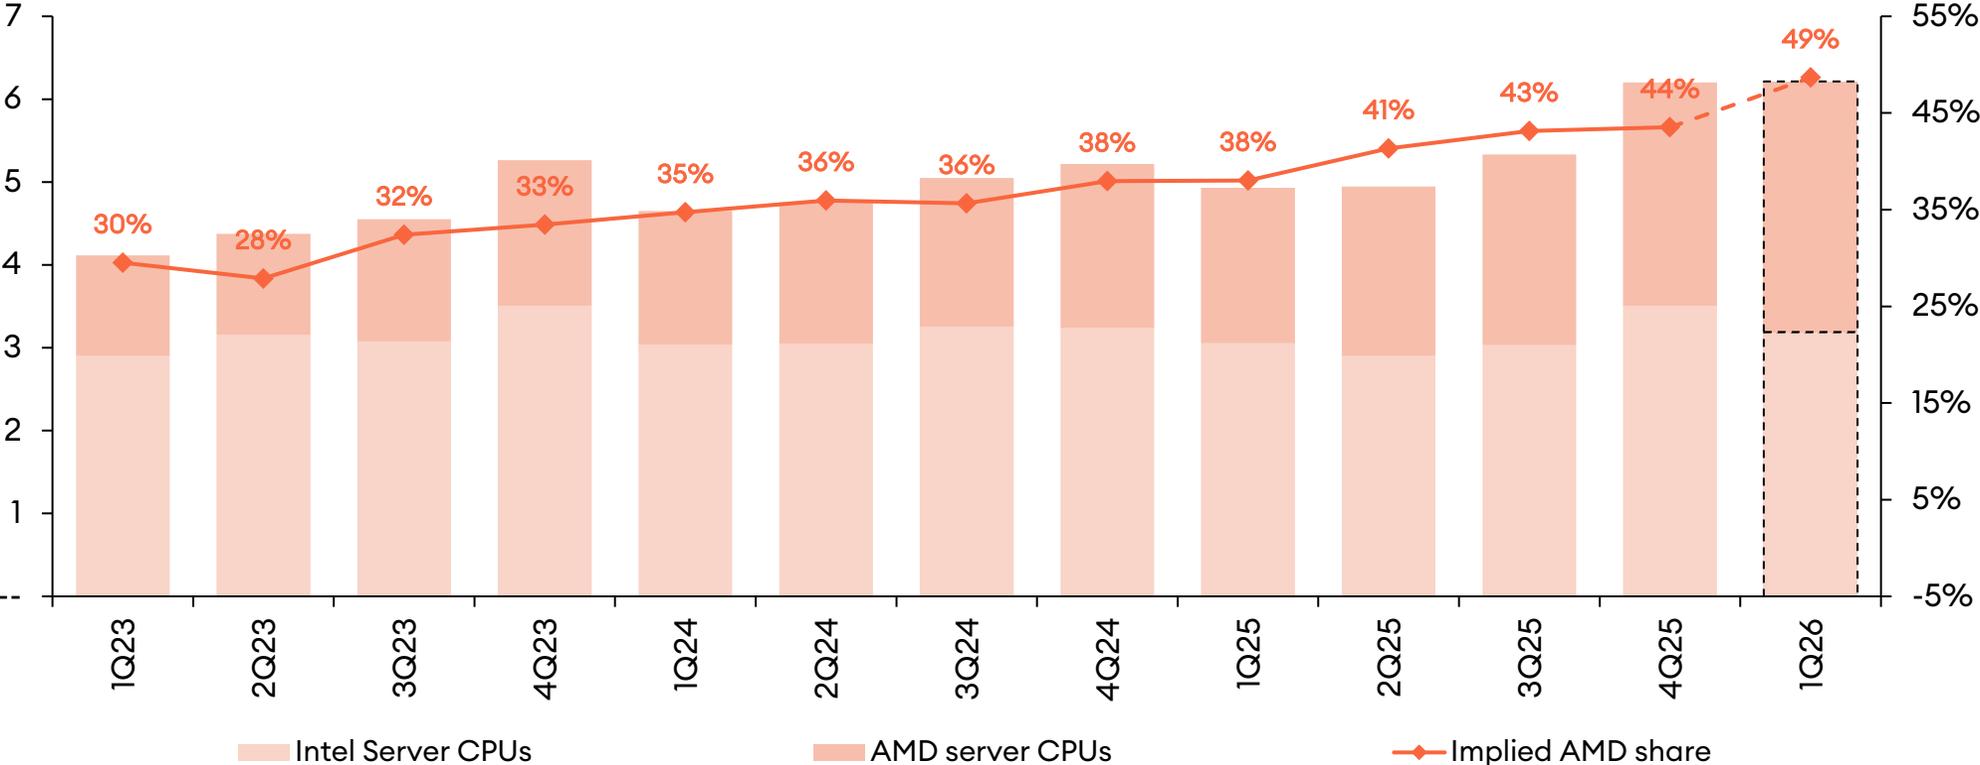


1 - Intel DCAI, Nvidia Datacenter, AMD Datacenter, Marvell Datacenter
Source: Corporate reports and NSR analysis.

AMD DC CPU share approaching 50%

AMD server share up 1pt QoQ to 44%, with we expect further gains in 1Q26

Intel vs AMD server CPU revenues (US\$bn) & market share (%)

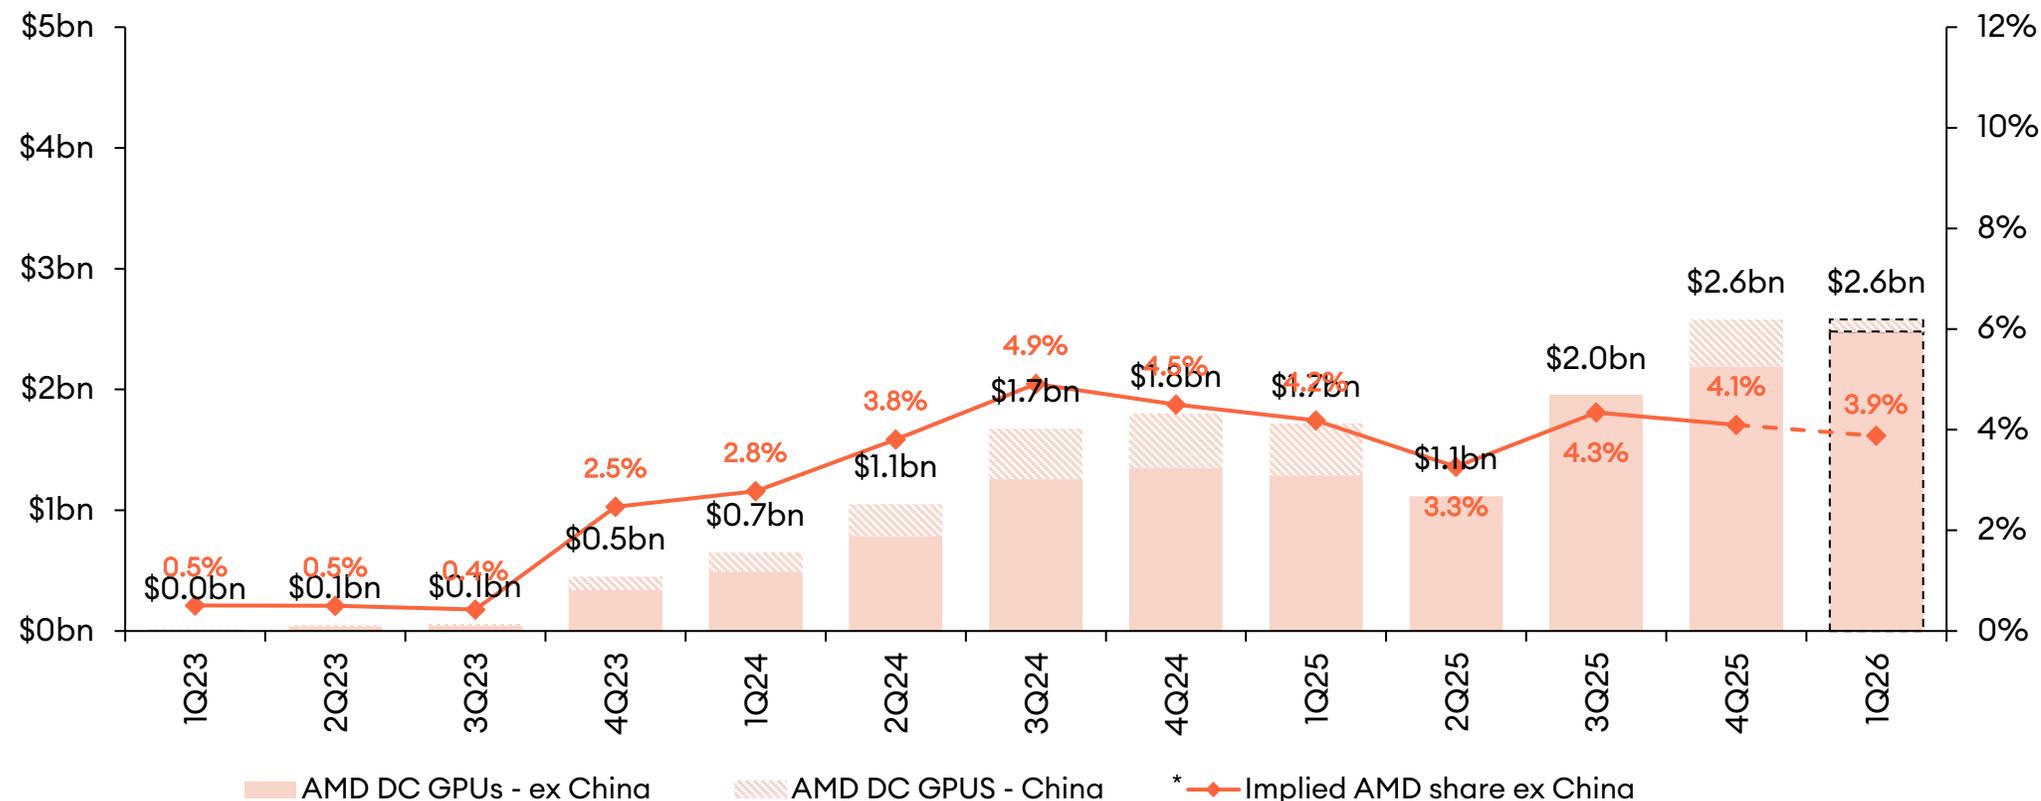


Source: Corporate reports, FactSet consensus and NSR estimates and analysis.

AMD DC GPU: Share (ex-China) broadly stable ahead of MI400 launch

AMD share ex China broadly stable at 4.1%; guides imply shares will remain broadly stable at 3.9% in 1Q26

AMD Datacenter GPU revenues (US\$bn) and market share** (%)



*we estimate China would have represented 40% of AMD datacenter GPU revenues in 2Q25, and represented 20-30% prior to that.

Source: Corporate reports and NSR estimates and analysis.

**based on NSR expectations for Nvidia in 2Q25 & beyond.

Hyperscaler replatforming: Microsoft doing ok, Google the best, Meta the worst

Against conventional wisdom (and stock reactions), we are concerned about Meta and see healthy transformation elsewhere

| | | YoY growth (%) | 2Q25 | 3Q25 | 4Q25 | 1Q26E ¹ | 2026E ¹ | |
|------------------|--|------------------------------|------|------|------|--------------------|--------------------|---|
| Microsoft | 2026 D&A up 56%, but cash costs growing only 6% vs. revenues 16%, making room for EBIT growing in line with revenues. | Revenues | 18% | 18% | 17% | 16% | 16% | Our replatforming framework is about measuring how increasing capital intensity drives revenue growth and efficiency improvement. |
| | | Total cost | 14% | 13% | 13% | 17% | 16% | |
| | | D&A | 75% | 21% | 62% | 66% | 56% | |
| | | Cash cost & SBC | 7% | 11% | 5% | 7% | 6% | |
| | | Operating income | 23% | 24% | 21% | 16% | 16% | |
| Google | D&A accelerating to +60% growth in 2026 from a low base. Cash costs growing 5pts slower than revenues make room for margin expansion. | Revenues | 14% | 16% | 18% | 19% | 17% | We can track the replatforming performance of hyperscalers, looking at how much the growth gap between cash costs and revenue growth makes room for increased depreciation. |
| | | Total cost | 14% | 13% | 16% | 19% | 16% | |
| | | D&A | 35% | 41% | 44% | 60% | 60% | |
| | | Cash cost & SBC ² | 12% | 11% | 14% | 15% | 12% | |
| | | Operating income | 14% | 22% | 23% | 20% | 19% | |
| Meta | Revenues up 24%, cash costs 36%, capex nearly doubling => margin squeeze. Meta spends more than it can afford. Investment for future performance or inefficient AI strategy? | Revenues | 22% | 26% | 24% | 30% | 24% | The 2026 outlook suggests Microsoft, Google and Amazon are on solid ground, while Meta largely overspends, with capex and opex growing too fast, squeezing margins. |
| | | Total cost | 12% | 32% | 40% | 45% | 39% | |
| | | D&A | 19% | 23% | 21% | 57% | 54% | |
| | | Cash cost & SBC | 10% | 34% | 45% | 43% | 36% | |
| | | Operating income | 38% | 18% | 6% | 8% | 3% | |
| Amazon | D&A growing 30-35% from a low base. Cash costs growing 3pts slower than revenues, enough to make room for margin expansion. | Revenues | 13% | 13% | 14% | 14% | 12% | |
| | | Total cost | 11% | 12% | 12% | 14% | 12% | |
| | | D&A | 26% | 25% | 25% | 36% | 32% | |
| | | Cash cost & SBC ² | 10% | 11% | 10% | 12% | 9% | |
| | | Operating income | 31% | 25% | 29% | 14% | 18% | |

Source: Company reports, VA Consensus and NSR estimates and analysis. (1) NSR expectations for Microsoft, VA consensus elsewhere. (2) Ex. one-off fines & charges.

Hyperscalers to nearly double AI capex again in 2026. Implications.

Hyperscalers' AI capex growth: 95/90% in 2025/26

The 4 hyperscalers have reported 4Q25. Based on their comments and guidance, are set to grow capex 73% and we estimate this implies AI capex growing 90%, vs. 95% last year.

| DC capex ¹ (US\$bn) | 2024 | 2025 | 2026 | Comments |
|--------------------------------|----------------|----------------|----------------|--|
| Microsoft | \$76bn | \$118bn | \$186bn | FY26 to accelerate vs. FY25. Assume 1FH27 too. |
| Google | \$53bn | \$91bn | \$180bn | Guide (4Q25 earnings). |
| Meta | \$37bn | \$72bn | \$125bn | Guide (4Q25 earnings). |
| Amazon | \$62bn | \$100bn | \$165bn | Guide (4Q25 earnings). |
| Total | \$227bn | \$380bn | \$656bn | |
| YoY growth (%) | | 67% | 73% | |
| Non-AI | \$80bn | \$92bn | \$110bn | |
| AI | \$148bn | \$288bn | \$546bn | |
| YoY growth (%) | | 95% | 90% | |

DC Servers/XPU growth: 60/>100% in 2025/26

Based on company comments, we estimate XPU revenues are set to double in 2026, in line with hyperscale capex. We see room for upside to this numbers as sparse comments point to the share of short duration assets to be stable to up in the overall mix.

| XPU ² (US\$bn) | 2024 | 2025 | 2026 | Comments |
|---------------------------|----------------|----------------|--------------------|------------------------|
| GPUs | \$107bn | \$164bn | >\$312bn | |
| Nvidia | \$102bn | \$157bn | >\$297bn | Jensen's guide |
| AMD/Others | \$5bn | \$7bn | >\$15bn | OpenAI ramp in 2H26 |
| ASICs | \$11bn | \$25bn | >\$69bn | |
| Broadcom | \$8bn | \$12bn | >\$36bn | Hock's guide |
| All other | \$3bn | \$13bn | >\$33bn | Trainium, Ascend, etc. |
| Total | \$118bn | \$190bn | >\$380bn | |
| YoY growth (%) | | 60% | >100% | |

We see room for XPU spending to more than double in 2026

Source: Company reports and NSR estimates and analysis

1 = Based on capex guides for Google/Meta/Amazon, NSR estimate for Microsoft, based on company comments.

Nvidia could be on a \$1tn run-rate business and at 2x consensus at the end of 2027

Nvidia now has visibility on \$1tn Blackwell/Rubin revenue through the end of 2027, implying a \$1tn run rate business exiting 2027

| NVIDIA DC revenues (US\$bn) | Total | 1Q25 | 2Q25 | 3Q25 | 4Q25 | 1Q26 | 2Q26 | 3Q26 | 4Q26 | 1Q27 | 2Q27 | 3Q27 | 4Q27 | 1Q28 | 2Q28 | 3Q28 | 4Q28 | 1Q29 | Comments | |
|-----------------------------------|-----------|-----------|-----------|-----------|--------------|------------|------------|------------|------------|------------|------------|------------|------------|------|------|------|------|------|---|---|
| Blackwell/Rubin fwd Orders as of: | | | | | | | | | | | | | | | | | | | | |
| 3Q25 | 500 | | | | 60 | 75 | 95 | 120 | 150 | | | | | | | | | | \$500bn disclosed at GTC DC (Oct-25). Allocated assuming stable 25% growth. | |
| | | | | | \$1tn | | | | | | | | | | | | | | | |
| 1Q26 | 869 | | | | | | 100 | 122 | 152 | 175 | 200 | 80 | 40 | | | | | | | \$1tn disclosed at GTC SJ (Mar-26). 2Q26-4Q26 quarters: minor increase vs. 2-4Q26 orders. 15% growth following 2 quarters. Minimal orders beyond. |
| HoH (%) | | | | | | | 5% | 2% | 1% | | | | | | | | | | | |
| 3Q26 | 1097 | | | | | | | | 155 | 180 | 210 | 222 | 230 | 100 | | | | | | All orders 4Q26-2Q27 minimally increased vs. orders as of 4Q25 (supply constraints). Single digit growth following 2 quarters. |
| HoH (%) | | | | | | | | | 2% | 3% | 5% | 178% | 475% | | | | | | | |
| 1Q27 | 1433 | | | | | | | | | | 215 | 228 | 240 | 245 | 250 | 255 | | | | All orders 2Q27-1Q28 minimally increased vs. orders as of 2 quarters ago. Low single digit growth beyond. |
| HoH (%) | | | | | | | | | | | 2% | 3% | 4% | 145% | | | | | | |
| 3Q27 | 1547 | | | | | | | | | | | | 247 | 250 | 255 | 260 | 265 | 270 | | All orders 4Q27-3Q28 minimally increased vs. orders as of 2 quarters ago. Low single digit growth beyond. |
| HoH (%) | | | | | | | | | | | | | 3% | 2% | 2% | 2% | | | | |
| Blackwell/Rubin revenues | | | | | 60 | 78 | 100 | 125 | 155 | 186 | 215 | 237 | 247 | | | | | | | |
| QoQ growth (%) | | | | | | 30% | 28% | 25% | 24% | 20% | 16% | 10% | 4% | | | | | | | |
| Other DC revenues | | | | | 2.3 | 2.5 | 3.2 | 4 | 5 | 6 | 6.9 | 7.6 | 7.9 | | | | | | | 4Q25 by difference. Grows in line with Blackwell/Rubin rev. beyond. |
| Total DC revenues | 39 | 41 | 51 | 62 | 81 | 103 | 129 | 160 | 192 | 222 | 244 | 255 | | | | | | | | >\$1tn run rate exiting 2027. |
| QoQ growth (%) | | 5% | 25% | 22% | 29% | 28% | 25% | 24% | 20% | 16% | 10% | 4% | | | | | | | | |
| Consensus DC revenues | | | | | 72 | 79 | 88 | 97 | 102 | 106 | 112 | 118 | | | | | | | | |
| NSR vs. consensus | | | | | 11% | 31% | 47% | 65% | 89% | 110% | 118% | 116% | | | | | | | | |

Source: Company reports, VA Consensus and NSR estimates and analysis.

Cloud Infrastructure – Implications for our coverage and beyond

Updated segment thesis

AI re-platforming. AI is increasing capital intensity and productivity. Microsoft Amazon & Google absorb increased capital intensity with cash costs growing slower than revenues. Meta is in overdrive, growing both capex and cash costs faster than revenues, squeezing margins.

Supply constraints across the supply chain: Public cloud growth is approaching 35% YoY, constrained by capacity addition. Similar comments at every step of the supply chain: compute, memory, networking, optics, final assembly & testing. This strengthens visibility on the medium term (see next paragraph) but creates some near-term risks on revenue momentum.

Visibility improving materially for 2027 and 2028. Nvidia's orders imply a \$1tn run rate business exiting 2027, 2x above consensus. Broadcom guided AI semis to >\$100bn in FY27 (Nov-27), likely up >90% YoY. The HBM and Optics supply chain have visibility extending to 2028. We expect material revisions to AI capex, as consensus models only 10% growth.

AI Ethernet ramping for large clusters. The opportunity will be large enough for Arista, Nvidia and Cisco to do well. InfiniBand will likely remain dominant in smaller clusters & amongst Enterprises.

Investment opportunities

Nvidia: Strong set-up in 2026 and 2027, adding to our best-idea list. **Buy, \$275 TP.**

AMD: OpenAI and Meta deals transformational. Best idea. **Buy, \$354 TP.**

Intel: A tough road ahead. The turnaround will take time. **Neutral, \$50 TP.**

Arista: demonstrates broad-based momentum. Well positioned for AI infrastructure. **Buy, \$175 TP.**

Broadcom: quality exposure, through compute ASICs and networking. **Buy, \$420 TP.**

In the broader leading-edge value chain, we prefer **TSMC (Buy, TWD 2,500)** over **Micron (Neutral, \$265).**

Disclosures

12 month historical recommendation changes are available on request

This report was produced by New Street Research LLP. 5 Brayford Square, London E1 0SG Tel: +44 20 7375 9111

Regulatory Disclosures: This research is directed only at persons classified as Professional Clients under the rules of the Financial Conduct Authority ('FCA'), and must not be re-distributed to Retail Clients as defined in the rules of the FCA.

This research is for our clients only. It is based on current public information which we consider reliable, but we do not represent that it is accurate or complete, and it should not be relied on as such. We seek to update our research as appropriate, but various regulations may prevent us from doing so. Most of our reports are published at irregular intervals as appropriate in the analyst's judgment.

This research is not an offer to sell or the solicitation of an offer to buy any security in any jurisdiction where such an offer or solicitation would be illegal. It does not constitute a personal recommendation or take into account the particular investment objectives, financial situations, or needs of individual clients.

All our research reports are disseminated and available to all clients simultaneously through electronic publication to our website.

© Copyright 2026 New Street Research LLP

No part of this material may be copied, photocopied or duplicated in any form by any means or redistributed without the prior written consent of New Street Research LLP.

New Street Research LLC is neither a registered investment advisor nor a broker/dealer. Subscribers and/or readers are advised that the information contained in this report is not to be construed or relied upon as investment, tax planning, accounting and/or legal advice, nor is it to be construed in any way as a recommendation to buy or sell any security or any other form of investment. All opinions, analyses and information contained herein is based upon sources believed to be reliable and is written in good faith, but no representation or warranty of any kind, express or implied, is made herein concerning any investment, tax, accounting and/or legal matter or the accuracy, completeness, correctness, timeliness and/or appropriateness of any of the information contained herein. Subscribers and/or readers are further advised that the Company does not necessarily update the information and/or opinions set forth in this and/or any subsequent version of this report. Readers are urged to consult with their own independent professional advisors with respect to any matter herein. All information contained herein and/or this website should be independently verified.

All research is issued under the regulatory oversight of New Street Research LLP.