

Compute market review: Q4 2025

Looking back at October to December 2025: buyers committed to GB300 systems at scale while powered buildings and grid capacity set the pace of deployment.

Published by	Rillor (Rillor Corporation)
Published	9 October 2026
Period covered	October to December 2025
Series	Compute market review
Pillar	Compute
Web version	rillor.com/insights/compute-market-review-q4-2025

Contents

[01](#) Summary[02](#) Background[03](#) Data and method[04](#) Findings[05](#) Implications[06](#) Limits of this analysis[07](#) Sources

Summary

Looking back at October to December 2025, the public record shows demand for AI compute meeting physical limits. Buyers committed to GB300 systems in large clusters and to non-NVIDIA accelerators by the gigawatt. Operators and the PJM grid operator reported that powered buildings and grid capacity set the pace at which that capacity could be delivered, while memory contract prices rose sharply.

- **H100 rents found a floor.** SemiAnalysis places the one-year H100 contract low at \$1.70 per GPU-hour in October [1]. Silicon Data's posted H100 index rose 10% in four weeks, from \$2.00 on 9 December to \$2.20 on 6 January [2].
- **GB300 entered contracts by name.** Microsoft Azure announced a production cluster of 4,608 Blackwell Ultra GPUs in GB300 systems for OpenAI workloads on 9 October [4]. AWS and OpenAI (\$38 billion) and IREN and Microsoft (\$9.7 billion) both named GB300 capacity in contracts signed on 3 November [13, 15]. NVIDIA said "cloud GPUs are sold out" [5].
- **Buyers widened their chip suppliers.** OpenAI agreed to 6 GW of AMD GPUs and 10 GW of accelerators built with Broadcom, and Anthropic to up to one million Google TPUs [7, 8, 11].
- **Financing took new forms.** Meta's Hyperion campus moved into an 80/20 joint venture with funds managed by Blue Owl, backed by leases and a residual value guarantee [9, 10]. Meta and Alphabet sold bonds with maturities to 2065 and 2075 [23, 24]. CoreWeave sold convertible notes and amended a term loan's covenants to match delayed deliveries [25, 26].
- **Buildings and power set the pace.** CoreWeave cut its 2025 capital spending guidance because a third-party developer was behind schedule [17]. Microsoft's chief executive described chips he could not plug in [36]. PJM's 2027/2028 capacity auction cleared at its cap and fell 6,623 MW short of its reliability requirement [37].
- **Memory joined the cost stack.** Fourth-quarter conventional DRAM contract prices rose 45% to 50% from the prior quarter, against an initial forecast of 8% to 13% [38, 39].

Background

AI compute reaches users in two ways. Buyers rent GPU capacity by the hour or by contract from hyperscale clouds and from neoclouds, the cloud providers built around GPU capacity. Or they own GPU systems, financed with equity, debt or customer prepayments. Rental prices, hardware prices, power and financing terms are linked: what a GPU earns on rent supports what an owner can pay for it and what a lender will advance against it.

By October 2025, two generations of NVIDIA hardware were in the market at once. Hopper GPUs (H100 and H200) made up most of the installed rental base. NVIDIA said "Blackwell sales are off the charts" [5], and the rack-scale GB300 NVL72, built on Blackwell Ultra, was moving into large production clusters [4]. CoreWeave's revenue backlog had nearly doubled in the third quarter, to \$55.6 billion at 30 September [16]. The fourth quarter tested whether contracted capacity could be delivered on schedule.

Data and method

Period. 1 October to 31 December 2025. Results that companies reported after the quarter are used for what they state about October to December, and each is labeled with its fiscal period. NVIDIA's fiscal fourth quarter ended on 25 January 2026, Oracle's fiscal second quarter on 30 November 2025, and Microsoft's fiscal second quarter on 31 December 2025. Two sources published after the quarter are used for specific facts: SemiAnalysis's April 2026 contract index [1] and TrendForce's February 2026 report of fourth-quarter memory prices [39].

Sources. Company press releases and SEC filings (Forms 8-K with their exhibits); earnings call transcripts and webcast pages; grid operator releases; published commentary from Silicon Data and SemiAnalysis on GPU rental prices; TrendForce's published memory price research; and reporting from trade and business publications where a primary document was not available. Every source appears in the numbered list at the end. We opened each source and took figures as stated.

Measures. Seven groups:

1. Rental prices: published H100 rental index levels and contract price surveys.
2. Accelerator supply: NVIDIA data center revenue, GB300 deployments and non-NVIDIA accelerator agreements.
3. Capacity contracts: announced agreements between compute providers, chip suppliers and buyers.
4. Financing: bonds, joint ventures, convertible notes, loan amendments and customer prepayments.
5. Delivery constraints: construction delays, statements about power, and the PJM capacity auction.
6. Memory: DRAM contract price forecasts and outcomes.
7. Capital spending: reported and guided capital expenditures of Alphabet, Amazon, Meta and Microsoft.

Treatment of figures. We keep posted rental rates, surveyed contract prices and reported deal values apart. Deal values are reported as their sources state them, including the words "up to" where they appear. Capital expenditure measures differ by company: Alphabet and Amazon report purchases of property and equipment; Meta reports capital expenditures including principal payments on finance leases; Microsoft reports capital expenditures including finance leases. Amazon's third-quarter guidance refers to cash capital expenditures as stated on its call.

Relation to the Rillor Compute Index. The Rillor Compute Index methodology v1.0 defines four series: RCI Rental, RCI Availability, RCI Hardware and RCI Rent-to-Value. This review keeps the same separation between rental, availability, hardware and financing questions. It reports third-party figures only and contains no RCI values.

Findings

1. Rental prices: a floor for the H100

Measure	Value	Source
One-year H100 contract low (SemiAnalysis index)	\$1.70 per GPU-hour, October 2025	[1]
Silicon Data H100 index, 9 December 2025	\$2.00 per GPU-hour	[2]
Silicon Data H100 index, 6 January 2026	\$2.20 per GPU-hour, up 10% in four weeks	[2]
A100 and B200 indexes over the same weeks	No significant change	[2]
Hyperscaler H100 median, July to December 2025	\$6.26 per GPU-hour	[3]
Neocloud H100 median, July to December 2025	\$3.33 per GPU-hour	[3]
Marketplace H100 median, July to December 2025	\$1.95 per GPU-hour	[3]

SemiAnalysis's one-year H100 index is built from monthly surveys of more than 100 market participants and checked against transaction data [\[1\]](#). It describes \$1.70 in October 2025 as the low from which prices rose, and says one-year contracts broke above \$2 per GPU-hour by late January 2026 [\[1\]](#).

Posted prices turned in December. Silicon Data reported that its H100 index rose from a \$2.00 baseline on 9 December to \$2.20 on 6 January, while its A100 and B200 indexes showed no significant change [\[2\]](#). It gave three groups of reasons. On supply, providers "prioritized their allocation policies toward reserved capacity and long-term contracts", tightness was most pronounced in Europe and Asia, and infrastructure still ran near capacity limits. On demand, it cited end-of-year training deadlines and budget cycles. On structure, it pointed to low elasticity of supply in the hourly market [\[2\]](#). It described a two-tier market in which spot and smaller buyers "bore the full pricing pressure" [\[2\]](#). Its January commentary gave a spread from about \$1.90 per GPU-hour at some specialist providers to over \$4 at hyperscalers for the same GPU [\[2\]](#).

Silicon Data did not read the move as a reversal of the longer decline in H100 prices, and expected stabilization with modest further declines [\[2\]](#). SemiAnalysis's later contract data showed prices continuing to rise into March 2026 [\[1\]](#). The two sources measure different markets: posted hourly rates and surveyed term contracts.

2. Accelerator supply: GB300 at cluster scale

Date	Event	Source
9 October 2025	Microsoft Azure announces NDv6 GB300 VMs: a production cluster of 4,608 Blackwell Ultra GPUs for OpenAI workloads	[4]
3 November 2025	AWS and OpenAI: clusters of GB200 and GB300 GPUs on EC2 UltraServers, all capacity targeted before the end of 2026	[13]
3 November 2025	IREN and Microsoft: access to GB300 GPUs at Childress, Texas, deployed in phases through 2026 in data centers supporting 200 MW of critical IT load	[15]
19 November 2025	NVIDIA: data center revenue of \$51.2 billion for the quarter ended 26 October; "cloud GPUs are sold out"	[5]
25 February 2026	NVIDIA: data center revenue of \$62.3 billion for the quarter ended 25 January 2026, up 75% from a year earlier	[6]

Azure called it a supercomputing-scale production cluster and set a goal of deploying hundreds of thousands of Blackwell Ultra GPUs [\[4\]](#). By November, GB300 appeared by name in contracts. AWS's agreement with OpenAI described "hundreds of thousands of state-of-the-art NVIDIA GPUs", both GB200s and GB300s, clustered on EC2 UltraServers [\[13\]](#). IREN agreed to buy the GPUs and ancillary equipment for its Microsoft contract from Dell for about \$5.8 billion [\[15\]](#).

NVIDIA's results for the quarter ended 26 October showed data center revenue up 25% from the prior quarter and 66% from a year earlier, and its outlook for the next quarter was \$65.0 billion, plus or minus 2% [\[5\]](#). The quarter that followed, ended 25 January 2026, came in at revenue of \$68.1 billion and data center revenue of \$62.3 billion [\[6\]](#).

Buyers also committed to other accelerators at scale:

Date	Parties	Terms as stated	Source
5 to 6 October 2025	AMD and OpenAI	6 GW of AMD Instinct GPUs across multiple generations; first 1 GW of MI450 in the second half of 2026; AMD expects tens of billions of dollars of revenue; warrant for up to 160 million AMD shares at \$0.01, vesting with purchases and share-price targets up to \$600, expiring 5 October 2030	[7, 40]

Date	Parties	Terms as stated	Source
13 October 2025	Broadcom and OpenAI	10 GW of OpenAI-designed accelerators with Broadcom Ethernet and connectivity, starting in the second half of 2026 and running through 2029	[8]
23 October 2025	Google and Anthropic	Up to one million TPUs, worth tens of billions of dollars, with well over a gigawatt of capacity coming online in 2026	[11]
30 October 2025	Amazon	Trainium2 "fully subscribed"; Project Rainier contains nearly 500,000 Trainium2 chips	[31]

The AMD warrant gives a buyer equity in its supplier in return for volume: OpenAI's first tranche vests after delivery of the first gigawatt of MI450, and full vesting requires 6 GW of purchases [\[7\]](#). Anthropic described its approach as using "three chip platforms", Google's TPUs, Amazon's Trainium and NVIDIA's GPUs [\[11\]](#).

3. Capacity contracts and their counterparties

Date	Parties	Terms as stated	Source
21 October	Meta and funds managed by Blue Owl	Joint venture to finance the roughly \$27 billion development of the Hyperion campus in Richland Parish, Louisiana; Blue Owl 80%, Meta 20%	[9, 10]
30 October	Core Scientific stockholders	Reject the CoreWeave merger; agreement terminated	[12]
3 November	AWS and OpenAI	\$38 billion, with growth over seven years	[13, 14]
3 November	IREN and Microsoft	About \$9.7 billion over five years, including a 20% prepayment	[15]
11 November	Nebius and Meta	About \$3 billion over five years	[18]
12 November	Anthropic with Fluidstack	\$50 billion investment in U.S. data centers in Texas and New York, with sites coming online through 2026	[19]

Date	Parties	Terms as stated	Source
18 November	Anthropic, Microsoft and NVIDIA	\$30 billion of Azure compute, plus up to 1 GW more; up to 1 GW of NVIDIA Grace Blackwell and Vera Rubin systems; NVIDIA to invest up to \$10 billion and Microsoft up to \$5 billion in Anthropic	[20]

Contracted backlogs grew. Oracle's remaining performance obligations rose by \$68 billion in its quarter to 30 November, to \$523 billion, up 438% from a year earlier, "highlighted by new commitments from Meta, NVIDIA, and others" [\[21\]](#). CoreWeave's revenue backlog was \$55.6 billion at 30 September and \$66.8 billion at 31 December [\[16, 27\]](#). Its contracted power rose from about 2.9 GW to about 3.1 GW over the same three months [\[16, 27\]](#).

Several of the quarter's agreements combined a compute purchase with an investment flowing the other way. NVIDIA and Microsoft committed to invest in Anthropic while Anthropic committed to buy Azure capacity running NVIDIA systems [\[20\]](#). AMD granted OpenAI a warrant tied to GPU purchases [\[7\]](#). These structures link a supplier's revenue to its customer's growth, and they make headline contract values harder to compare with plain capacity leases.

4. Financing took new forms

Date	Issuer	Instrument	Amount and terms	Source
21 October 2025	Meta and Blue Owl	Joint venture for the Hyperion campus	Blue Owl contributed about \$7 billion in cash at formation; Meta received a one-time \$3 billion distribution; debt issued to PIMCO and other bond investors; Meta leases the campus and provides a residual value guarantee	[9, 10]
3 November 2025	Meta	Senior notes in six tranches, due 2030 to 2065	\$30 billion, coupons 4.200% to 5.750%	[23]
6 November 2025	Alphabet	U.S. dollar and euro senior notes	\$17.5 billion and €6.5 billion; longest dollar tranche due 2075	[24]
11 November 2025	Nebius	At-the-market equity program	Up to 25 million Class A shares	[18]

Date	Issuer	Instrument	Amount and terms	Source
8 to 11 December 2025	CoreWeave	1.75% convertible senior notes due 2031	About \$2.59 billion; initial conversion price about \$107.80 per share	[25]
31 December 2025	CoreWeave	Amendment to its \$2.6 billion DDTL 3.0 term loan	Lower minimum liquidity for March and April 2026 payment dates; first debt service coverage test moved to 31 October 2027	[26]

The Hyperion structure moved a campus off one company's direct capital spending and into a vehicle owned mostly by outside capital, with Meta as tenant and guarantor of residual value [\[9, 10\]](#). Data Center Frontier reported an initial target of about 2 GW of capacity by 2030 [\[10\]](#). Bond markets supplied the rest of the large platforms' external funding in the quarter, at long maturities [\[23, 24\]](#).

For neoclouds, the quarter showed financing adjusting to delivery risk. CoreWeave said its DDTL 3.0 amendment aligned the facility with the delivery timing it described on its third-quarter call [\[26\]](#). It also expanded its revolving credit facility to \$2.5 billion during the quarter [\[27\]](#). IREN's contract with Microsoft includes a 20% prepayment, so the customer funds part of the build in advance [\[15\]](#).

Oracle reported trailing four-quarter free cash flow of negative \$13.2 billion at 30 November, with capital expenditures of \$20.5 billion in the first half of its fiscal year [\[21\]](#). Subtracting the \$8.5 billion it reported for its first quarter gives about \$12.0 billion for the quarter to November (our arithmetic) [\[21, 41\]](#). It raised its fiscal 2026 capital spending plan to about \$50 billion, from about \$35 billion in September [\[22\]](#).

5. Buildings and power set the pace

Date	Source	Statement
29 October 2025	Microsoft	Short "the space or the power" to put chips in; capacity constrained "through at least the end of our fiscal year"
30 October 2025	Amazon	Added more than 3.8 GW of power in 12 months; at least 1 GW more expected in the fourth quarter
3 November 2025 (reported)	Microsoft chief executive, in a podcast interview	"you may actually have a bunch of chips sitting in inventory that I can't plug in"
10 November 2025	CoreWeave	A third-party data center developer "is behind schedule"; 2025 capital spending guidance cut to \$12 billion to \$14 billion

Date	Source	Statement
17 December 2025	PJM	2027/2028 capacity auction clears at the \$333.44 per MW-day cap and falls 6,623 MW short of the reliability requirement

Sources: [\[30, 31, 36, 17, 37\]](#).

CoreWeave's call showed delivery risk at a neocloud in detail. It said the delays in "powered shell delivery associated with the data center provider" would affect its fourth-quarter results, that the affected customer had agreed to adjust the delivery schedule and extend the expiration date, and that the total contract value was unchanged [\[17\]](#). It expected 2026 capital spending to be "well in excess of double that of 2025" [\[17\]](#). It ended the third quarter with about 590 MW of active power and ended the year with more than 850 MW [\[16, 27\]](#).

PJM's December auction secured 134,479 MW, plus 11,299 MW from regions under the fixed resource requirement, for a total of 145,777 MW [\[37\]](#). That left the region 6,623 MW short of its reliability requirement, with a reserve margin of 14.8% against a target of about 20% [\[37\]](#). Forecast peak load for the 2027/2028 delivery year was about 5,250 MW above the forecast used a year earlier, and PJM attributed nearly 5,100 MW of the increase to data centers [\[37\]](#). Its executive vice president for market services and strategy said the auction "leaves no doubt that data centers' demand for electricity continues to far outstrip new supply" [\[37\]](#).

6. Memory joined the cost stack

Date	TrendForce statement	Source
29 October 2025	Fourth-quarter conventional DRAM contract price forecast raised to 18% to 23%, from 8% to 13%; DDR5 contract prices expected to rise through 2026	[38]
26 February 2026	Fourth-quarter conventional DRAM contract prices rose 45% to 50% from the prior quarter; including HBM, 50% to 55%; DRAM revenue of \$53.58 billion, up 29.4%	[39]

TrendForce attributed the increase to data center expansion by large cloud providers, including general-purpose servers alongside AI servers, and to buyers struggling to secure supply, which strengthened suppliers' pricing power [\[38, 39\]](#). Memory sits in every server that hosts a GPU, so its price moves into the cost of AI capacity. On its January call, Microsoft said "rising memory prices would impact capital expenditures" [\[35\]](#).

7. Capital spending

Company	Measure	Quarter to September 2025	Quarter to December 2025	Full year 2025	Guidance in late October
Alphabet	Purchases of property and equipment	\$24.0 billion	\$27.9 billion	\$91.4 billion	\$91 to \$93 billion for 2025
Microsoft	Capital expenditures including finance leases	\$34.9 billion	\$37.5 billion	Fiscal year ends June	Fiscal 2026 growth rate above fiscal 2025
Meta	Capital expenditures including finance lease principal	\$19.4 billion	\$22.1 billion	\$72.2 billion	\$70 to \$72 billion for 2025; 2026 growth "notably larger"
Amazon	Cash capital expenditures (Q3, call); purchases of property and equipment (Q4 and year)	\$34.2 billion	\$39.5 billion	\$131.8 billion	About \$125 billion for 2025

Sources: Alphabet [28, 32]; Microsoft [30, 35]; Meta [29, 34]; Amazon [31, 33].

Microsoft said about two-thirds of its spending in the quarter to December went to short-lived assets, "primarily GPUs and CPUs" [35]. It reported commercial remaining performance obligations of \$625 billion at 31 December, about 45% of them from OpenAI [35]. Alphabet reported Google Cloud backlog of \$240 billion at the end of 2025, and its chief executive said he expected to go through 2026 "in a supply constrained way" [32].

Implications

For buyers of compute. The quarter separated two markets. Hopper capacity found a floor in contract pricing and began to firm in posted pricing [1, 2], while new GB300 capacity was committed in multi-year blocks before it was built [13, 15]. A buyer that needed capacity in 2026 faced a choice between available older-generation GPUs and newer systems with delivery dates tied to construction and power.

For builders and operators. Delivery risk moved from chips to powered buildings. A single developer's delay changed CoreWeave's quarterly revenue and capital spending, and the fix ran through both the customer contract (a new delivery schedule) and the loan agreement (covenant relief) [17, 26]. Grid capacity was a binding constraint in PJM, where the auction cleared short of the reliability standard [37].

For lenders. The quarter added structures that share risk across parties: joint ventures with residual value guarantees, convertible notes, customer prepayments and vendor equity tied to purchases [9, 25, 15, 7]. Each changes who bears a shortfall if delivery slips or demand weakens. Only the documents show which party that is.

For researchers. Memory prices became a measurable input to compute costs in this quarter [38, 39]. Studies of GPU system costs and rental economics that hold component costs fixed across late 2025 misstate them.

Limits of this analysis

- This review uses public sources only. It contains no transaction-level price data and no data collected by Rillor.
- Index values are reported as their publishers stated them. Silicon Data's posted-price index and SemiAnalysis's contract survey describe different markets, and we did not recompute either.
- Several facts come from secondary reporting where no primary document was available to us: the Broadcom terms [8], the Meta and Blue Owl ownership split and cash contributions [10], Oracle's revised capital spending plan [22], the Microsoft chief executive's podcast remarks [36] and Amazon's call statements [31]. Transcripts and summaries can contain errors.
- Deal values are maximums or totals over long terms, and several depend on delivery. Values that combine investments and purchases are not comparable with plain leases.
- Oracle's quarterly capital expenditure figure is our arithmetic from its reported six-month and first-quarter totals.
- Fiscal periods do not match the calendar quarter. NVIDIA's, Oracle's and Microsoft's periods are labeled where used.
- This is a review of public information. It is not investment advice and makes no forecast of prices.

Sources

1. Daniel Nishball, Jordan Nanos, Cheang Kang Wen and others, The Great GPU Shortage: Rental Capacity, Launching our H100 1 Year Rental Price Index. SemiAnalysis, 2 April 2026. <https://newsletter.semianalysis.com/p/the-great-gpu-shortage-rental-capacity>
2. Carmen Li, H100 Price Spike: Understanding the 10% Surge in GPU Rental Costs. Silicon Data, 9 January 2026. <https://www.silicondata.com/blog/h100-price-spike>
3. Carmen Li, H100 Rental Price Over Time (2023-2025): A Complete Market Analysis. Silicon Data, 21 December 2025. <https://www.silicondata.com/blog/h100-rental-price-over-time>
4. Microsoft Azure Unveils World's First NVIDIA GB300 NVL72 Supercomputing Cluster for OpenAI. NVIDIA Corporation (blog), 9 October 2025. <https://blogs.nvidia.com/blog/microsoft-azure-worlds-first-gb300-nvl72-supercomputing-cluster-openai>
5. NVIDIA Announces Financial Results for Third Quarter Fiscal 2026. NVIDIA Corporation, 19 November 2025. <https://nvidianews.nvidia.com/news/nvidia-announces-financial-results-for-third-quarter-fiscal-2026>
6. NVIDIA Announces Financial Results for Fourth Quarter and Fiscal 2026. NVIDIA Corporation, 25 February 2026. <https://nvidianews.nvidia.com/news/nvidia-announces-financial-results-for-fourth-quarter-and-fiscal-2026>
7. Current Report on Form 8-K (warrant issued to OpenAI, event of 5 October 2025). Advanced Micro Devices, Inc., 6 October 2025. <https://www.sec.gov/Archives/edgar/data/2488/000119312525230895/d28189d8k.htm>
8. OpenAI, Broadcom outline custom AI accelerator, networking deal. Constellation Research, 13 October 2025. <https://www.constellationr.com/insights/news/openai-broadcom-outline-custom-ai-accelerator-networking-deal>
9. Kirkland Represents Blue Owl on Joint Venture with Meta to Develop Hyperion Data Center. Kirkland & Ellis LLP, 21 October 2025. <https://www.kirkland.com/news/press-release/2025/10/kirkland-represents-blue-owl-on-joint-venture-with-meta-to-develop-hyperion-data-center>
10. Meta's \$27B Hyperion Campus: A New Blueprint for AI Infrastructure Finance. Data Center Frontier, 5 November 2025. <https://www.datacenterfrontier.com/hyperscale/article/55325337/metas-27b-hyperion-campus-a-new-blueprint-for-ai-infrastructure-finance>
11. Expanding our use of Google Cloud TPUs and Services. Anthropic, 23 October 2025. <https://www.anthropic.com/news/expanding-our-use-of-google-cloud-tpus-and-services>
12. Core Scientific shareholders reject merger with CoreWeave, agreement terminated. Investing.com, 31 October 2025. <https://www.investing.com/news/sec-filings/core-scientific-shareholders-reject-merger-with-coreweave-agreement-terminated-93CH-4325472>
13. AWS and OpenAI announce multi-year strategic partnership. Amazon, 3 November 2025. <https://www.aboutamazon.com/news/aws/aws-open-ai-workloads-compute-infrastructure>

- L4. Mark Brohan, OpenAI's \$38 billion AWS deal signals new AI era for ecommerce. Digital Commerce 360, 3 November 2025. <https://www.digitalcommerce360.com/2025/11/03/openai-38-billion-aws-deal-ai-ecommerce/>
- L5. IREN Secures \$9.7bn AI Cloud Contract with Microsoft (Exhibit 99.1 to Form 8-K). IREN Limited, 3 November 2025. https://www.sec.gov/Archives/edgar/data/1878848/000114036125040072/ef20058139_ex99-1.htm
- L6. CoreWeave Reports Strong Third Quarter 2025 Results (Exhibit 99.1 to Form 8-K). CoreWeave, Inc., 10 November 2025. <https://www.sec.gov/Archives/edgar/data/1769628/000176962825000059/coreweave3q25earningspress.htm>
- L7. CoreWeave, Inc. Q3 2025 Earnings Call, corrected transcript. CoreWeave, Inc., 10 November 2025. https://s205.q4cdn.com/133937190/files/doc_financials/2025/q3/CORRECTED-TRANSCRIPT_-CoreWeave-Inc-CRWW-US-Q3-2025-Earnings-Call-10-November-2025-5_00-PM-ET.pdf
- L8. Nebius reports third quarter 2025 financial results. Nebius Group N.V., 11 November 2025. <https://nebius.com/newsroom/nebius-reports-third-quarter-2025-financial-results>
- L9. Anthropic invests \$50 billion in American AI infrastructure. Anthropic, 12 November 2025. <https://www.anthropic.com/news/anthropic-invests-50-billion-in-american-ai-infrastructure>
- L10. Microsoft, NVIDIA and Anthropic announce strategic partnerships. Anthropic, 18 November 2025. <https://www.anthropic.com/news/microsoft-nvidia-anthropic-announce-strategic-partnerships>
- L11. Oracle Announces Fiscal Year 2026 Second Quarter Financial Results. Oracle Corporation, 10 December 2025. <https://investor.oracle.com/investor-news/news-details/2025/Oracle-Announces-Fiscal-Year-2026-Second-Quarter-Financial-Results/default.aspx>
- L12. Oracle defends infrastructure spending spree amid mounting AI demand. Channel Dive, 11 December 2025. <https://www.channeldive.com/news/oracle-capex-spike-cloud-ai-data-center/807716/>
- L13. Current Report on Form 8-K (offering of \$30 billion of senior notes). Meta Platforms, Inc., 3 November 2025. <https://www.sec.gov/Archives/edgar/data/1326801/000119312525262593/d75422d8k.htm>
- L14. Current Report on Form 8-K (U.S. dollar and euro senior notes offerings). Alphabet Inc., 6 November 2025. <https://www.sec.gov/Archives/edgar/data/1652044/000119312525269979/d56115d8k.htm>
- L15. Current Report on Form 8-K (1.75% convertible senior notes due 2031, event of 8 December 2025). CoreWeave, Inc., 11 December 2025. <https://www.sec.gov/Archives/edgar/data/1769628/000176962825000105/crww-20251208.htm>
- L16. Current Report on Form 8-K (first amendment to the DDTL 3.0 credit agreement, event of 31 December 2025). CoreWeave, Inc., 2 January 2026. <https://www.sec.gov/Archives/edgar/data/1769628/000176962826000003/crww-20251231.htm>
- L17. CoreWeave Reports Strong Fourth Quarter and Fiscal Year 2025 Results (Exhibit 99.1 to Form 8-K). CoreWeave, Inc., 26 February 2026. <https://www.sec.gov/Archives/edgar/data/1769628/000176962826000094/coreweave4q25earningspress.htm>

28. Alphabet Announces Third Quarter 2025 Results (Exhibit 99.1 to Form 8-K). Alphabet Inc., 29 October 2025. <https://www.sec.gov/Archives/edgar/data/1652044/000165204425000087/googexhibit991q32025.htm>
29. Meta Reports Third Quarter 2025 Results (Exhibit 99.1 to Form 8-K). Meta Platforms, Inc., 29 October 2025. <https://www.sec.gov/Archives/edgar/data/1326801/000162828025047114/meta-09302025xexhibit991.htm>
30. Microsoft Fiscal Year 2026 First Quarter Earnings Conference Call. Microsoft Corporation, 29 October 2025. <https://www.microsoft.com/en-us/investor/events/fy-2026/earnings-fy-2026-q1>
31. Amazon (AMZN) Q3 2025 Earnings Call Transcript (call of 30 October 2025). The Motley Fool, 31 October 2025. <https://www.fool.com/earnings/call-transcripts/2025/10/31/amazon-amzn-q3-2025-earnings-call-transcript/>
32. Amanda Gerut, Alphabet plans to double capex spending to a possible \$185 billion, but it's keeping CEO Sundar Pichai up at night. Fortune, 4 February 2026. <https://fortune.com/2026/02/04/alphabet-google-ai-spending-supply-constraints/>
33. Amazon.com Announces Fourth Quarter Results. Amazon.com, Inc., 5 February 2026. <https://ir.aboutamazon.com/news-release/news-release-details/2026/Amazon-com-Announces-Fourth-Quarter-Results/default.aspx>
34. Meta Reports Fourth Quarter and Full Year 2025 Results. Meta Platforms, Inc., 28 January 2026. <https://investor.atmeta.com/investor-news/press-release-details/2026/Meta-Reports-Fourth-Quarter-and-Full-Year-2025-Results/default.aspx>
35. Microsoft Fiscal Year 2026 Second Quarter Earnings Conference Call. Microsoft Corporation, 28 January 2026. <https://www.microsoft.com/en-us/investor/events/fy-2026/earnings-fy-2026-q2>
36. Berry Zwets, Satya Nadella: Power shortage, not GPU inventory, is the problem. Techzine Global, 3 November 2025. <https://www.techzine.eu/news/infrastructure/135966/satya-nadella-power-shortage-not-gpu-inventory-is-the-problem/>
37. PJM Auction Procures 134,479 MW of Generation Resources. PJM Interconnection, 17 December 2025. <https://www.pjm.com/-/media/DotCom/about-pjm/newsroom/2025-releases/20251217-pjm-auction-procures-134479-mw-of-generation-resources.pdf>
38. Tight DRAM Supply to Boost DDR5 Contract Prices; Profitability in 2026 Expected to Surpass HBM3e, Says TrendForce. TrendForce, 29 October 2025. <https://www.trendforce.com/presscenter/news/20251029-12758.html>
39. Price Rally Drives 4Q25 DRAM Revenue Up 29.4%; Samsung Regains No. 1 Market Share, Says TrendForce. TrendForce, 26 February 2026. <https://www.trendforce.com/presscenter/news/20260226-12937.html>
40. AMD and OpenAI announce strategic partnership to deploy 6 gigawatts of AMD GPUs (Exhibit 99.1 to Form 8-K). Advanced Micro Devices, Inc., 6 October 2025. <https://www.sec.gov/Archives/edgar/data/2488/000119312525230895/d28189dex991.htm>

11. Oracle Announces Fiscal Year 2026 First Quarter Financial Results. Oracle Corporation, 9 September 2025. <https://investor.oracle.com/investor-news/news-details/2025/Oracle-Announces-Fiscal-Year-2026-First-Quarter-Financial-Results/default.aspx>

About this report

Published by Rillor on 9 October 2026. AI tools assist research and drafting, and every figure cites its source. Rillor builds agentic AI systems, datasets, market research and compute services.

Cite as: Rillor. Compute market review: Q4 2025. 9 October 2026.
<https://rillor.com/insights/compute-market-review-q4-2025>

Notices

Rillor is not a registered investment adviser or commodity trading advisor and does not provide investment or trading advice. Research, forecasts, data and software described on this site are for research and engineering use. Nothing here is an offer or recommendation to buy or sell any security, commodity interest or digital asset. Past or simulated results do not indicate future results.

Rillor Compute Index series are for research and planning use. They are not offered for settlement of any contract.

NVIDIA, AMD and related product names are trademarks of their owners and are used only to identify products. No affiliation or endorsement is implied.