

Compute market review: Q2 2026

Looking back at April to June 2026: capital spending plans rose, memory became the cost to watch, H100 rents split by tier and exchanges announced GPU futures plans.

Published by	Rillor (Rillor Corporation)
Published	9 October 2026
Period covered	April to June 2026
Series	Compute market review
Pillar	Compute
Web version	rillor.com/insights/compute-market-review-q2-2026

Contents

01 Summary	05 Implications
02 Background	06 Limits of this analysis
03 Data and method	07 Sources
04 Findings	

Summary

Looking back at April to June 2026, the public record shows buyers committing more money to AI compute than suppliers could deliver, and cost pressure moving from GPUs toward memory.

- **Spending plans rose.** On its first-quarter earnings call, Alphabet guided 2026 capital expenditures to \$180 billion to \$190 billion, and on 29 April Meta raised its 2026 range to \$125 billion to \$145 billion from \$115 billion to \$135 billion, citing higher component pricing [7][9]. On 1 June, Alphabet announced equity offerings of \$80 billion as part of its plan to fund AI compute and said demand for its AI services exceeded its available supply [7].
- **Accelerator sales grew.** NVIDIA reported data center revenue of \$75.2 billion for the quarter ended 26 April 2026, up 92% from a year earlier [1]. On 1 June it said its Vera Rubin platform was ramping into full production [2]. AMD's data center segment reported \$5.8 billion for the first quarter, up 57% [3].
- **Memory became the cost to watch.** TrendForce expected conventional DRAM contract prices to rise 58% to 63% from the first quarter, and later reported DRAM industry revenue of about \$154.7 billion for the quarter, up 59.5% [5][6]. Micron's revenue for the quarter ended 28 May was \$41.46 billion, against \$23.86 billion the quarter before [4].
- **Rental prices split by tier.** Silicon Data's hyperscaler H100 on-demand index stayed between \$7.40 and \$7.55 per GPU-hour across April and May. Its neocloud H100 index fell to \$2.47 on 8 April and closed May at \$2.73 [17]. Nebius later reported more than 30% higher pricing on older-generation GPUs in the quarter than in the first quarter [14].
- **Backlogs and borrowing grew together.** CoreWeave reported a revenue backlog of \$99.4 billion and an \$8.5 billion investment-grade delayed draw term loan [12]. Oracle reported remaining performance obligations of \$638 billion and negative free cash flow of \$23.7 billion for its fiscal year [11]. IREN signed about \$3.6 billion of financing secured by GPUs and a customer contract [15].
- **Exchanges moved toward GPU price risk.** CME Group with Silicon Data (12 May) and Intercontinental Exchange with Ornn (19 May) announced plans for GPU compute futures, each subject to regulatory review [20][21].

Background

AI compute is bought in two ways. Buyers rent capacity by the GPU-hour from hyperscalers 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.

Each quarter, the public record on this market grows by several hundred pages: earnings releases, filings, exchange notices, regulatory actions and index commentary. This review collects what that record says about April to June 2026, sorted into demand, supply, price, financing, market structure and policy. It reports what sources state. It does not forecast prices.

Two features of the market shape how the evidence should be read. First, the GPU rental market has at least two tiers. Silicon Data's published commentary in April described its hyperscaler H100 index as priced at roughly three times its neocloud H100 index from 1 March to 20 April 2026 [\[16\]](#). Any statement about "the price of an H100" needs a tier. Second, the companies that report results each quarter use different fiscal calendars, so a calendar quarter collects results for periods that end on different dates.

Data and method

Period. 1 April to 30 June 2026. Results that companies reported during these three months are included, and each is labeled with the fiscal period it covers. NVIDIA's quarter ended 26 April, Micron's on 28 May and Oracle's fiscal year on 31 May [1][4][25]. Seven documents published after the quarter (Nebius's second-quarter letter, Alphabet's second-quarter release, FactSet's July analysis, a July press report on Oracle, TrendForce's September revenue report and two later Silicon Data commentaries) are used only for what they state about April to June.

Sources. Company press releases and SEC filings; exchange press releases; summaries of regulatory actions by law firms; TrendForce's published memory research; Silicon Data's published index commentary; and FactSet's published analysis of hyperscaler financing. Every source appears in the numbered list at the end.

Measures. Six groups:

1. Demand: capital expenditures reported and guided by large cloud and internet companies.
2. Supply of accelerators: revenue and product status at NVIDIA and AMD.
3. Supply of memory: contract price forecasts, industry revenue and supplier results.
4. Rental prices: published GPU rental index levels by tier, and operator statements about pricing.
5. Backlog and financing: contracted revenue, cash flow and debt at operators.
6. Market structure and policy: compute futures announcements, export control guidance and power interconnection.

Treatment of figures. Index values are reported as their publisher stated them. We did not recompute them. Hyperscaler and neocloud figures stay separate throughout. Capital expenditure measures differ by company: Meta reports capital expenditures including principal payments on finance leases, while Alphabet and Amazon report purchases of property and equipment. The tables label each measure.

Findings

1. Demand: capital spending plans rose again

The quarter opened with first-quarter reports from large buyers of AI infrastructure.

Company	Period reported	Spending in the period	2026 outlook stated in the quarter
Alphabet	Quarter to 31 March 2026	\$35.7 billion, purchases of property and equipment [8]	\$180 billion to \$190 billion; 2027 expected to increase significantly [7]
Meta	Quarter to 31 March 2026	\$19.84 billion, including principal payments on finance leases [9]	\$125 billion to \$145 billion, up from \$115 billion to \$135 billion [9]
Amazon	Quarter to 31 March 2026	\$44.2 billion, purchases of property and equipment [10]	Not restated in the release [10]

Meta attributed its higher range mainly to "higher component pricing this year" and, to a lesser extent, to data center costs for future capacity [9]. Amazon's trailing twelve-month free cash flow fell to \$1.2 billion from \$25.9 billion a year earlier, which the company attributed mainly to higher purchases of property and equipment reflecting investment in AI [10]. AWS segment sales rose 28% to \$37.6 billion [10].

On 1 June, Alphabet announced equity offerings totaling \$80 billion in expected aggregate amount: \$30 billion of underwritten offerings, a \$40 billion at-the-market program expected to begin in the third quarter, and a \$10 billion private placement with Berkshire Hathaway [7]. Alphabet said it was "experiencing strong demand for its AI solutions and services from enterprises and consumers, at levels that are exceeding the company's available supply" [7]. Alphabet said it intended to use the proceeds of the underwritten offerings and the private placement for general corporate purposes, including capital expenditures for AI infrastructure and compute, and the at-the-market proceeds primarily to meet tax obligations on employee equity awards, about \$30 billion of them in calendar 2026 [7]. The same release listed its other funding: \$174 billion of operating cash flow over the twelve months to 31 March 2026, and more than \$85 billion of debt raised over the prior year [7].

FactSet's analysis of five large hyperscalers estimated that new debt funded 32% of their capital spending in the twelve months before June 2026, up from 9% in fiscal 2024 [24]. Read together, these documents show buyers funding AI capacity from operating cash, debt and, at Alphabet, new equity.

2. Supply: accelerators

NVIDIA reported on 20 May for its quarter ended 26 April 2026 [1].

Measure	Value	Change from a year earlier
Total revenue	\$81.6 billion	+85%
Data center revenue	\$75.2 billion	+92%
Data center compute revenue	\$60.4 billion	+77%
Data center networking revenue	\$14.8 billion	+199%
Outlook for the next quarter	\$91.0 billion, plus or minus 2%	

NVIDIA said its outlook assumed no data center compute revenue from China [1]. The same release announced a new reporting split for data center revenue between hyperscale customers and a group covering AI clouds, industrial and enterprise customers [1].

At its Computex keynote on 1 June (Taipei time), NVIDIA announced that Vera Rubin was "ramping into full production" [2]. Its chief executive said, "The supply chain we created for Vera Rubin is twice as large as Grace Blackwell" [2]. NVIDIA also described MGX-compatible 800 VDC power racks as a bridge to hybrid AC and 800 VDC designs for sites built around alternating current [2].

AMD reported on 5 May: first-quarter revenue of \$10.3 billion and data center segment revenue of \$5.8 billion, up 57%, which it attributed to demand for EPYC processors and the ramp of Instinct GPU shipments [3]. It guided second-quarter revenue to about \$11.2 billion [3]. The release also described plans by Meta to deploy up to 6 gigawatts of AMD Instinct GPUs, with the first gigawatt based on a custom MI450 GPU, and Samsung HBM4 supply for the MI455X [3].

3. Supply: memory

Memory moved from a component line to a headline cost in the quarter.

Source	Date	Statement
TrendForce	31 March 2026	Conventional DRAM contract prices expected to rise 58% to 63% quarter on quarter in 2Q26; NAND Flash 70% to 75% [5]

Source	Date	Statement
Micron	24 June 2026	Revenue of \$41.46 billion for the quarter ended 28 May, against \$23.86 billion the prior quarter and \$9.30 billion a year earlier [4]
TrendForce	7 September 2026	DRAM industry revenue in 2Q26 of about \$154.73 billion, up 59.5% quarter on quarter; bit shipments grew only modestly [6]

TrendForce linked the 2Q26 increase to North American cloud providers accelerating AI inference deployments, to suppliers reallocating capacity toward HBM and server memory, and to cloud providers accepting higher prices under long-term agreements to secure supply [\[5\]](#). Micron reported that HBM4 was in high-volume shipments for its lead customer's platform [\[4\]](#). Meta's statement that component pricing drove its higher spending range ties the memory market directly to buyers' budgets [\[9\]](#).

The practical point for compute buyers is that a GPU system's cost now carries a memory term that moves on its own cycle.

4. Rental prices: two tiers, two behaviors

Silicon Data publishes daily GPU rental indexes and separates the hyperscaler tier from the neocloud tier. Its commentary for April and May reported the following [\[17\]](#).

Series	1 April	Range in the window	31 May	Annualized volatility
H100 hyperscaler on-demand (AWS, Azure, Google Cloud)	\$7.44	\$7.40 to \$7.55	\$7.46	3.8%
H100 neocloud (CoreWeave, Lambda, Crusoe, Nebius)	\$2.63	\$2.47 (8 April) to \$2.76 (30 May)	\$2.73	18.4%

Prices are in U.S. dollars per GPU-hour. On the hyperscaler series, 43 of 60 daily transitions were unchanged [\[17\]](#). The neocloud series rose 7.1% in May alone, and the Bloomberg-listed neocloud H100 rental index (ticker SDH100RT) stood at \$2.75 in early June, up 12.7% over the trailing three months [\[17\]](#). Silicon Data interpreted the hyperscaler tier as behaving like an administered price ceiling posted above a large book of already-priced reserved capacity, and the neocloud tier as reflecting demand shifts within days [\[17\]](#). Those are the publisher's readings; the data show the levels and the dispersion.

Silicon Data's later commentary filled in June [\[18\]](#)[\[19\]](#):

- The neocloud H200 index started its official history at \$2.70 on 4 May and closed June at \$2.87 [18].
- The neocloud H100 index reached \$2.79 on 5 June and fell to \$2.59 in late June [18].
- The H200 premium over the H100 averaged 5.8% in May, touched 0.4% on 14 June and ended June near 11% [18].
- The B200 index opened June at 5.42 and closed it at 5.36, a 1.1% decline, with a range of 5.26 to 5.45 [19].

Operators described the same direction. Nebius wrote in May that "Pricing continued to rise for new generation GPUs, with older generation chips also seeing strong pricing support" [13]. In its second-quarter letter it reported "more than 30% higher pricing on older-generation GPUs versus Q1," and four large contracts with annual contract value of \$20 million to \$25 million per megawatt [14]. That last figure is a contract price expressed per unit of power, a different unit from the per GPU-hour index levels above.

5. Backlog, cash flow and financing

Operators reported large contracted backlogs and financed them through structures tied to specific assets and customers.

Company	Period	Revenue	Contracted measure	Power
CoreWeave	Quarter to 31 March 2026	\$2,078 million, against \$982 million a year earlier [12]	Revenue backlog \$99.4 billion [12]	More than 1 GW active; contracted power over 3.5 GW [12]
Nebius	Quarter to 31 March 2026	\$399.0 million, up 684% [13]	Annualized run-rate revenue \$1.92 billion [13]	Contracted capacity above 3.5 GW; year-end target raised to more than 4 GW [13]
Oracle	Quarter to 31 May 2026 [25]	\$19.2 billion, up 21% [11]	Remaining performance obligations \$638 billion, up \$85 billion in the quarter [11]	Not stated in the source

CoreWeave reported a net loss of \$740 million for its quarter [12]. It also reported DDTL 4.0, described as "an \$8.5 billion non-recourse investment grade rated delayed draw term loan facility" with a floating tranche at SOFR plus 2.25% and a fixed tranche at about 5.9%, and a \$2 billion Class A common stock investment from NVIDIA [12].

Oracle reported negative free cash flow of \$23.7 billion for fiscal 2026, its year ended 31 May [11][25]. For fiscal 2027 it guided to net cash capital expenditures of about \$70 billion, with reported capital expenditures \$20 billion to \$25 billion higher because of customer prepayments and timing [11]. It planned

to raise about \$40 billion in debt and equity in fiscal 2027, including a previously announced \$20 billion at-the-market equity program, and stated that prepaid and customer-supplied hardware within large AI contracts totaled \$75 billion [11].

On 29 May, a subsidiary of IREN signed financing of about \$3.6 billion: a delayed draw term loan of about \$1.5 billion at term SOFR plus 2.25%, and \$2.1 billion of 5.96% senior notes due 31 December 2031 [15]. The proceeds partly fund GPU infrastructure for a contract to provide dedicated GPU services to Microsoft at data centers in Childress, Texas. The debt is secured by the subsidiary's assets, including the GPUs, a pledge of its equity and the contract's cash flows, with a minimum debt service coverage ratio of 1.05 [15].

Three funding patterns appear in the quarter's documents: customers paying for or supplying hardware (Oracle), debt secured by GPUs and a named contract (IREN, CoreWeave), and new equity at the buyer level (Alphabet).

6. Market structure: plans for compute futures

Two exchange groups announced plans to list futures on GPU rental prices.

Date	Parties	What was announced	Status stated
12 May 2026	CME Group and Silicon Data	A compute futures market "later this year," based on Silicon Data's indices of on-demand GPU rental rates [20]	Pending regulatory review [20]
19 May 2026	Intercontinental Exchange and Ornn	A suite of GPU compute futures on Ornn's Compute Price Index (OCPI), U.S. dollar denominated and cash-settled; series for H100, H200, B200 and RTX 5090 [21]	Subject to regulatory approval [21]

The two announcements describe different index designs. Silicon Data's indices measure on-demand rental rates [20]. The ICE and Ornn release describes OCPI as tracking "live-traded spot prices" and as built "only from printed transactions" [21]. In evidence-class terms, one starts from published rates and the other from completed trades. A buyer comparing the two needs to know which one its own contracts resemble.

7. Policy and power

On 31 May, the Bureau of Industry and Security issued guidance that a license is required to export advanced computing items to entities headquartered in Country Group D:5 or Macau, or whose ultimate parent is headquartered there, even when the entity is located elsewhere [22]. Greenberg Traurig's summary notes that the requirement dates from November 2023 and that bona fide data center operators already using such items, and otherwise acting consistently with the regulations, need not stop using or servicing them until further notice [22].

On 18 June, the Federal Energy Regulatory Commission issued show cause orders to six regional transmission organizations and system operators on how large loads such as data centers connect to the grid [\[23\]](#). Each had 60 days to justify its current tariffs or file changes, and 30 days to file a report on how it would ensure adequate generation for existing and new large loads. The orders were FERC's first substantive action on the Department of Energy's October 2025 request, and the rulemaking remained open [\[23\]](#).

Implications

For buyers of compute. The tier matters. Across April and May, the hyperscaler and neocloud H100 indexes differed by a factor of roughly 2.7 to 3 and moved with very different volatility [16][17]. A budget built on one tier's posted rate says little about the other. Contract pricing reported per megawatt [14] and posted pricing per GPU-hour are different units, and a buyer comparing offers should convert both to the same basis with the same scope.

For builders and operators. Memory cost entered system cost at a scale buyers named in their guidance [5][9]. Long-term supply agreements shaped who paid the increases [5]. Operator statements about firm pricing for older GPUs [13][14] matter for fleets that mix generations. Interconnection rules for large loads were under active review at FERC at the end of the quarter [23].

For lenders. The quarter's financings tie debt to identifiable GPUs and to named customer contracts, with coverage tests [12][15]. Oracle's figures show how backlog, prepayments and negative free cash flow can coexist in one balance sheet [11]. Collateral value for GPU loans depends on rental income, and the evidence here shows rental income behaving differently by tier [17].

For researchers. Index design is now a market question. The two futures announcements reference indexes built from different evidence [20][21]. Fiscal calendars also matter: NVIDIA's "first quarter" ended in April and Oracle's fiscal year in May, so a calendar-quarter view has to label periods with care.

The segmentation used in this review follows the Rillor Compute Index methodology v1.0, which shows hyperscaler and neocloud segments separately in every series that uses them and records each rental observation with its pricing type and term.

Limits of this analysis

- The review uses public sources only. It includes no private transaction data, so it describes published prices, reported results and stated plans.
- Rental index levels come from one publisher's commentary. We report them as published and did not reproduce them.
- Capital spending measures are not identical across companies. Meta's figure includes principal payments on finance leases; Alphabet's and Amazon's are purchases of property and equipment.
- Company fiscal periods do not match the calendar quarter. Each figure is labeled with its own period.
- Seven documents used here were published after the quarter ended. They are used only for statements about April to June.
- The futures announcements describe plans subject to regulatory review. This review does not assess whether or when the contracts trade.
- The review describes past events. It contains no price forecast and is not investment advice.

Sources

1. NVIDIA Announces Financial Results for First Quarter Fiscal 2027. NVIDIA Corporation, 20 May 2026. <https://nvidianews.nvidia.com/news/nvidia-announces-financial-results-for-first-quarter-fiscal-2027>
2. NVIDIA GTC Taipei at COMPUTEX: Live Updates on What's Next in AI. NVIDIA Blog, 1 June 2026 (Taipei time). <https://blogs.nvidia.com/blog/nvidia-gtc-taipei-computex-2026-news/>
3. AMD Reports First Quarter 2026 Financial Results. Advanced Micro Devices, Inc. (GlobeNewswire), 5 May 2026; text as reproduced by StockTitan. <https://www.stocktitan.net/news/AMD/amd-reports-first-quarter-2026-financial-bu7e2cbxpd14.html>
4. Micron Technology, Inc. Reports Record Results for the Third Quarter of Fiscal 2026 (Exhibit 99.1 to Form 8-K). Micron Technology, Inc., 24 June 2026. <https://www.sec.gov/Archives/edgar/data/0000723125/000072312526000013/a2026q3ex991-pressrelease.htm>
5. AI Server Demand to Drive Memory Contract Price Increases in 2Q26 as CSPs Secure Supply via Long-Term Agreements. TrendForce, 31 March 2026. <https://www.trendforce.com/presscenter/news/20260331-12995.html>
6. DRAM Industry Revenue Rises 59.5% QoQ in 2Q26 as Supply Expansion Continues to Lag Demand Growth. TrendForce, 7 September 2026. <https://www.trendforce.com/presscenter/news/20260907-13219.html>
7. Alphabet Announces Proposed \$80 Billion Equity Capital Raise to Expand AI Infrastructure and Compute. Alphabet Inc., 1 June 2026. https://s206.q4cdn.com/479360582/files/doc_news/2026/Jun/01/attachments/2026-June-Alphabet-Equity-Capital-Raise-Press-Release-PDF.pdf
8. Alphabet Announces Second Quarter 2026 Results (Exhibit 99.1 to Form 8-K), including quarterly purchases of property and equipment for the first quarter of 2026. Alphabet Inc., 22 July 2026. <https://www.sec.gov/Archives/edgar/data/0001652044/000165204426000066/googexhibit991q22026.htm>
9. Meta Reports First Quarter 2026 Results. Meta Platforms, Inc., 29 April 2026. <https://investor.atmeta.com/investor-news/press-release-details/2026/Meta-Reports-First-Quarter-2026-Results/default.aspx>
10. Amazon.com Announces First Quarter Results (Exhibit 99.1 to Form 8-K). Amazon.com, Inc., 29 April 2026. <https://www.sec.gov/Archives/edgar/data/0001018724/000101872426000012/amzn-20260331ex991.htm>
11. Tarsilla Moura, Oracle Q4 2026 Earnings: \$638B Backlog Turns AI Cloud Growth into Funding Test. ERP Today, 11 June 2026. <https://erp.today/oracle-q4-2026-earnings-ai-cloud-backlog-funding/>
12. CoreWeave Reports Strong First Quarter 2026 Results. CoreWeave, Inc., 7 May 2026. <https://investors.coreweave.com/news/news-details/2026/CoreWeave-Reports-Strong-First-Quarter-2026-Results/>
13. Nebius Group Letter to Shareholders, Q1 2026 (Exhibit 99.2 to Form 6-K). Nebius Group N.V., 13 May 2026. https://www.sec.gov/Archives/edgar/data/0001513845/000110465926059872/tm2614392d1_ex99-2.htm

14. Nebius Group Letter to Shareholders, Q2 2026 (Exhibit 99.2 to Form 6-K). Nebius Group N.V., 12 August 2026. https://www.sec.gov/Archives/edgar/data/1513845/000110465926094568/tm2622968d1_ex99-2.htm
15. IREN Limited, Current Report on Form 8-K (financing agreements dated 29 May 2026). U.S. Securities and Exchange Commission EDGAR, 2026. https://www.sec.gov/Archives/edgar/data/0001878848/000114036126023427/ef20075181_8k.htm
16. Carmen Li, H100 Hyperscaler Index, April 2026: A Benchmark in Flat Mode. Silicon Data, 27 April 2026. <https://www.silicondata.com/blog/h100-hyperscaler-index-april-2026>
17. Carmen Li, H100 Index, May 2026: Same Chip, Two Tiers, Two Regimes. Silicon Data, 16 June 2026. <https://www.silicondata.com/blog/h100-index-may-2026-two-tiers-two-regimes>
18. Carmen Li, H200 vs H100 Rental Prices, May to July 2026: The Premium That Doubled. Silicon Data, 4 August 2026. <https://www.silicondata.com/blog/h200-vs-h100-rental-prices-may-july-2026>
19. B200 Rental Price Update, August 2026: The Quietest Month Yet. Silicon Data, 24 September 2026. <https://www.silicondata.com/blog/b200-rental-price-august-2026-update>
20. CME Group and Silicon Data Partner to Launch First Compute Futures. CME Group and Silicon Data (PR Newswire), 12 May 2026; text as reproduced by StockTitan. <https://www.stocktitan.net/news/CME/cme-group-and-silicon-data-partner-to-launch-first-compute-xsea0cl8o8z2.html> (original: https://www.cmegroup.com/media-room/press-releases/2026/5/12/cme_group_and_silicondatapartnertolaunchfirstcompute-futures.html)
21. ICE and Ornn to Launch GPU Compute Futures Contracts. Intercontinental Exchange, Inc. (Business Wire), 19 May 2026. https://s2.q4cdn.com/154085107/files/doc_news/ICE-and-Ornn-to-Launch-GPU-Compute-Futures-Contracts-2026.pdf
22. "Enforcement Pause" Has Limits: BIS Clarifies Ongoing License Requirement for Advanced Computing Items to China-Linked Entities. Greenberg Traurig, 12 June 2026. <https://www.gtlaw.com/en/insights/2026/6/enforcement-pause-has-limits-bis-clarifies-ongoing-license-requirement-for-advanced-computing-items-to-china-linked-entities>
23. FERC Acts to Advance Data Center and Large Load Integration in Six RTO Regions. Duane Morris LLP, 23 June 2026. https://www.duanemorris.com/alerts/ferc_acts_advance_data_center_large_load_integration_six_rto_regions_0626.html
24. Guniz Kama and Bina Rajput, Hyperscalers Tap External Financing as AI Capex Outruns Cash Flow. FactSet Insight, 23 July 2026. <https://insight.factset.com/hyperscalers-tap-external-financing-as-ai-capex-outruns-cash-flow>
25. Ana-Maria Stanciuc, Oracle is one notch above junk after S&P downgrade as AI data-centre spending burns through cash. The Next Web, 16 July 2026. <https://thenextweb.com/news/oracle-sp-downgrade-ai-spending-junk-risk>

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: Q2 2026. 9 October 2026.
<https://rillor.com/insights/compute-market-review-q2-2026>

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.