Rillor

EXPLAINER

Rack power and liquid cooling: what NVL72 changed

Vendor figures put an NVL72 rack at 120 to 142 kW, against about 11 kW for a typical rack, and that gap changes cooling, power and colocation pricing.



Rillor

Contents

01 Summary

02 Background

03 Data and method

04 Findings

05 Implications

06 Limits of this analysis

07 Sources

Summary

- Typical rack densities in operating data centers sit between about 7.5 kW and 11 kW, depending on the survey year and on whether high-density sites are counted. Uptime Institute's 2025 survey put the average of operators' typical (modal) rack densities at almost 9 kW, or 7.5 kW without facilities with typical density of 30 kW or more [1]. Reporting on the 2026 survey gives a modal density of 11 kW, and an average of 7.8 kW without a small number of facilities designed for more than 30 kW [2].
- A GB200 NVL72 rack draws approximately 120 kW according to NVIDIA's documentation [7], with vendor figures from 125 kW to 135 kW for configured racks [9], and a reference design for GB300 NVL72 supports up to 142 kW per rack [10].
- Power per GPU changed little. Derived from vendor figures, an air-cooled DGX B200 draws about 1.8 kW per GPU and a GB200 NVL72 rack about 1.7 kW per GPU [5, 7]. What changed is the number of GPUs in one rack: 72 in an NVL72, against 16 in two air-cooled DGX B200 systems, the typical deployment in NVIDIA's guidance [5].
- Air cooling reaches its limits well below these densities. ASHRAE notes that a 40 to 50 kW rack could need up to 5,000 cfm of air, against about 1,900 cfm from a floor tile it rates "best-in-class" [13].
- Liquid cooling comes in three classes: close-coupled air-to-liquid such as rear-door heat exchangers, cold plates and immersion [12]. NVL72 racks are hybrids: cold plates on CPUs and GPUs, air for networking and storage [7].
- Wholesale colocation is priced per kW per month. At CBRE's average asking rate of $196.25 per kW per month for 250 to 500 kW requirements in the second half of 2025 [17], rent per GPU comes to about $330 to $390 a month across the configurations examined, because rent tracks power and power per GPU is similar. Density decides which sites can host the racks, and CBRE reports rent premiums for facilities built for superintelligence (SI) [17].
- NVIDIA states that 54 V in-rack power distribution reaches physical limits as racks exceed 200 kW, and plans 800 V DC distribution for 1 MW racks starting in 2027 [16].

Background

Rack power has risen in steps. The U.S. Department of Energy's best-practice guide for data-center design, revised in July 2024, notes that high-performance computing centers adopted direct liquid cooling early, with compute racks at 60 kW observed in 2013 and recently surpassing 125 kW [12]. NVIDIA wrote in April 2025 that data centers once operated at 20 kW per rack, while hyperscale facilities can support over 135 kW per rack [3].

The installed base moves more slowly. Uptime Institute's 2025 survey of more than 800 operators found strong adoption of racks between 10 kW and 30 kW, with the average of typical densities at almost 9 kW, up from 8.3 kW a year earlier [1]. Excluding the small number of facilities with typical density of 30 kW or more, the average was 7.5 kW [1]. The same survey reported that over a quarter of respondents, 27.5%, work in facilities 16 years old or more, and that facilities commissioned within five years of the survey averaged a power usage effectiveness of 1.48 [1]. According to Data Center Knowledge's report on the 2026 survey, released on 29 July 2026, the modal rack density reached 11 kW, up from 9 kW in 2025, and the average excluding a small number of facilities designed for more than 30 kW was 7.8 kW [2]. The report describes operators planning for two markets at once, SI training clusters and CPU-based enterprise workloads, and quotes Uptime's chief technical officer on more operators pursuing phased builds to avoid stranding capacity [2].

Against that base, rack-scale GPU systems are a step change. NVIDIA's DGX GB rack guide describes an NVL72 rack as 18 compute trays, each with 2 Grace CPUs and 4 Blackwell GPUs, 9 NVLink switch trays, 2 management switches and power shelves, with liquid-cooling manifolds and a power bus bar at the rear [7]. In October 2024, NVIDIA contributed elements of the GB200 NVL72 electro-mechanical design to the Open Compute Project, including the rack architecture, compute and switch tray mechanicals, and liquid-cooling and thermal environment specifications [14]. Colocation operators note the shift in their own filings: Digital Realty's annual report for 2025 says continued SI adoption could result in evolving infrastructure needs, particularly around power density for advanced computing [18].

Data and method

This report compiles published figures.

Sources. System power, size, weight and cooling come from vendor documentation: NVIDIA's DGX user guides and reference architectures [4, 5, 6, 7, 8], and other vendors' datasheets, specification sheets and announcements [9, 10, 11, 15]. Cooling classes and terms come from the U.S. Department of Energy [12] and ASHRAE [13]. Installed-base densities come from Uptime Institute's surveys [1, 2]. Colocation pricing comes from CBRE [17], and contract practice on energy costs from Digital Realty's annual report [18]. Forward-looking power distribution comes from NVIDIA's published architecture [16].

Normalization. We express power per rack and per GPU. Where a vendor gives power for a system, per-GPU power divides by the number of GPUs in that system. Rack power for air-cooled systems multiplies system power by the systems per rack in NVIDIA's guidance [5]. Where vendors give different figures for the same rack, we report each with its source.

Derived figures. Rent illustrations multiply rack power by CBRE's average asking rate. Airflow illustrations apply NVIDIA's planning ratio of 150 cfm per kW for the DGX B200 [5] to other rack powers. Floor loading divides a published rack weight by its published footprint. Each derived figure is labeled as such.

Findings

1. Rack power by system

System	Cooling	Size	Power	GPUs	Power per GPU (derived)
NVIDIA DGX H100 [4]	Air	8U, 130.45 kg max	10.2 kW max	8	1.3 kW
NVIDIA DGX B200 [5]	Air	10U, more than 130 kg (estimated)	14.3 kW max (estimated)	8	1.8 kW
NVIDIA DGX B300 [6]	Air	10U; 168 kg in the power-supply version	14.5 kW consumption; 15 kW system max	8	1.8 kW
NVIDIA DGX GB200 NVL72 rack [7]	Liquid for compute trays; air for networking and storage	One rack with 18 compute trays and 9 switch trays	About 120 kW	72	1.7 kW
Supermicro GB200 NVL72 rack [9]	Liquid, with in-rack or in-row CDU options	2,236 mm by 600 mm by 1,068 mm	132 kW total; 125 kW to 135 kW operating	72	1.7 to 1.9 kW
GB300 NVL72, Schneider Electric reference design [10]	Liquid-to-liquid CDUs	Not stated	Up to 142 kW per rack	72	2.0 kW

Two observations follow. Per-GPU power for Blackwell-generation systems sits between about 1.7 kW and 2.0 kW whether the GPUs are in an air-cooled server or a liquid-cooled rack. And vendor figures for the same rack differ: NVIDIA's documentation says approximately 120 kW for the DGX GB200 rack [7], Supermicro lists 125 kW to 135 kW operating for its configuration [9], and NVIDIA's own blog cites Schneider Electric infrastructure that supports up to 132 kW per rack for GB200 NVL72 [3]. A site design and a contract should name the figure they use.

2. What sets the density

In an NVL72 rack, the GPUs are connected by NVLink through a passive copper cable cartridge backplane, and the rack forms one 72-GPU NVLink domain [7]. NVIDIA describes the design as letting the GPUs operate as one tightly integrated processing unit with a full-rack power density of 120 kW [3]. NVIDIA presents the

density as part of the interconnect design.

Liquid cooling also raises density for 8-GPU servers. Supermicro's liquid-cooled 4U HGX B200 system fits 8 systems and 64 GPUs in a 42U or 48U rack, and up to 12 systems and 96 GPUs in a 52U rack, while up to 4 of its air-cooled 10U HGX B200 systems fit in one rack, 32 GPUs [11]. The same GPU in the same server family can sit at two or three times the rack density depending on the cooling grade.

3. The limits of air

NVIDIA's DGX B200 planning guidance assumes about 150 cfm of supply air per kilowatt at sea level, and up to about 2,145 cfm for each system [5]. Applied to rack-scale systems, the same ratio gives about 18,000 cfm for a 120 kW rack and 19,800 cfm for a 132 kW rack (derived). ASHRAE's 2021 white paper puts the problem in floor-tile terms: airflow for a 40 to 50 kW rack could be up to 5,000 cfm, and a floor tile it rates "best-in-class" delivers about 1,900 cfm [13]. It also noted that servers requiring 100 cfm or more per rack unit could fill only about 19U of a rack from a single tile [13].

NVIDIA's guidance for the DGX B200 reflects these limits. A typical deployment puts 2 systems in a rack, 28.6 kW, and 4 systems, 57.2 kW, is an option only at sites designed for extreme-density air-cooled deployments [5]. Supply air should generally meet ASHRAE Class A1 with a maximum of 30 °C [5]. The guidance also states that supplemental cooling such as rear-door heat exchangers and in-row coolers is not typically suitable for DGX B200 systems [5].

ASHRAE's white paper gives a second reason for the shift. Chip makers are raising thermal design power while lowering the case temperatures their chips require, which makes it harder to use the warmer ASHRAE classes for both air and water [13]. Its recommendation was that any future data center include the capability to add liquid cooling [13].

4. Liquid cooling options

The Department of Energy guide groups liquid cooling into three technology classes [12]:

Class	Where heat enters the liquid	Variants	Notes
Localized air-to-liquid heat exchange	A coil at or near the rack	Rear-door heat exchangers, enclosed cabinets, in-row coolers, overhead cooling; active or passive doors	The servers stay air-cooled; the coil captures their heat
Cold plates	Plates mounted on the chips, replacing fin heat sinks	Can extend to memory and other components	Generally served by a CDU; a direct-liquid-cooling technology

Class	Where heat enters the liquid	Variants	Notes
Immersion	Electronics submerged in dielectric fluid	Single-phase, pumped by a CDU; two-phase, with fluid boiling below component limits and condensing in a passive cycle	A direct-liquid-cooling technology

The guide adds that fluids on the IT side range from treated water to glycol-based solutions to dielectric fluids, that several liquid-cooled solutions are hybrids in which liquid captures the majority but not all of the heat, and that most approaches use a CDU, which interfaces with the facility cooling loop and provides cooling liquid at the right temperature, pressure and chemistry for the IT equipment [12].

The NVL72 is a cold-plate hybrid. Its compute trays are cooled by liquid that runs up and down the rack through manifolds and through cold plates attached to the CPUs and GPUs, while networking and storage devices are air-cooled by fans [7]. NVIDIA's reference architecture for DGX GB200 describes the same hybrid of direct liquid cooling and air cooling [8].

ASHRAE's water classes for liquid cooling carry their upper temperature limit in degrees Celsius in the name: W17 (previously W1), W27 (W2), W32 (W3), W40 (a new class), W45 (W4) and W+ (W5) [13]. The change is included in the fifth edition of ASHRAE's Thermal Guidelines for Data Processing Environments, released in 2021 [12]. Higher W classes may allow dry coolers to reject heat to the atmosphere, saving energy and water [12].

5. Heat rejection outside the building

Liquid cooling moves heat to the facility loop, and the building still has to reject it. NVIDIA's April 2025 overview lists four approaches [3]:

Approach	How it works	Trade-off noted by NVIDIA [3]
Mechanical chillers	Chilled water absorbs heat	Reliable across climates, energy-intensive
Evaporative cooling	Evaporation of water absorbs heat	Energy-efficient, with high water use that can reach millions of gallons per megawatt each year in large facilities
Dry coolers	A closed liquid loop rejects heat to ambient air through finned coils	No water use; effectiveness depends on ambient temperature
Pumped refrigerant	Liquid refrigerant carries heat to outdoor exchangers	No water use; needs refrigerant handling and monitoring

Because direct-to-chip cooling transfers heat to liquid at the source, it can run with warmer water and reduce or eliminate mechanical chillers in a wide range of climates [3]. The choice of heat rejection follows from the climate, the water available and the W class the IT equipment accepts.

6. CDU capacity and layout

CDUs come in rack, row and air-coupled forms. Supermicro's GB200 NVL72 datasheet lists an in-rack 4U CDU with 250 kW capacity, an in-row CDU with 1.3 MW capacity, and optional 180 kW or 240 kW liquid-to-air units for facilities without a cooling tower and water supply [9]. Supermicro's HGX B200 announcement describes a 250 kW CDU in a 4U form factor [11]. Schneider Electric's GB300 reference design hosts three GB300 NVL72 clusters, up to 1,152 GPUs, using liquid-to-liquid CDUs and high-temperature chillers [10].

A 1.3 MW in-row unit covers about ten racks at 120 kW or nine at 142 kW (derived), before redundancy. An in-rack unit uses rack space. A liquid-to-air unit avoids facility water and returns the heat to the room, so the room's air system has to carry it.

7. What per-kW pricing implies

CBRE reported that the average monthly asking rate for a 250 to 500 kW requirement in primary North American wholesale colocation markets rose 6.6% year over year to $196.25 per kW per month in the second half of 2025 [17]. CBRE quotes these as asking rates for requirements of 250 kW and above at N+1 or Tier III [17].

Deployment	kW per rack	GPUs per rack	Asking rent per rack-month (derived)	Per GPU-month (derived)
GB200 NVL72 at 120 kW [7]	120	72	$23,550	$327
GB200 NVL72 at 132 kW [9]	132	72	$25,905	$360
GB300 NVL72 at 142 kW [10]	142	72	$27,868	$387
Two DGX B200 [5]	28.6	16	$5,613	$351
Four DGX B200 [5]	57.2	32	$11,226	$351

At the same rate, a 250 kW block comes to $49,063 a month and a 1 MW block to $196,250 a month, before energy (derived).

The table shows that per-kW pricing makes space-and-power rent per GPU roughly constant across configurations, because power per GPU is roughly constant. The rate itself is where density shows up. CBRE reports that operators with SI-optimized facilities, including liquid cooling and high-power-density racks, captured rent premiums over standard colocation [17], and that pricing for 3 to 10 MW requirements rose 12.5% year over year [17]. An asking rate for a standard 250 to 500 kW requirement is a reference point. Where premiums for SI-optimized facilities apply, liquid-cooled capacity prices above it. The CBRE rate is an asking price.

Energy can be billed on top. Digital Realty states that the vast majority of its utilities expense is passed directly through to customers, which limits its exposure to rising power costs [18]. Under per-kW rent with energy passed through, a tenant pays for provisioned capacity and for the energy it uses.

8. Provisioned and drawn power

A tenant pays for the kilowatts provisioned to its racks, and the electrical design need not match the draw. NVIDIA's DGX B200 power tables, based on two circuits in an N provisioning scheme, show the stranded capacity at peak demand [5]:

Circuit	Breaker derating	Rack capacity	Systems per rack	Peak demand	Stranded at peak
3-phase 415 V, 32 A	80%	35 kW	2	28.6 kW	6.4 kW
3-phase 415 V, 60 A	80%	66 kW	4	57.2 kW	8.4 kW
3-phase 415 V, 63 A	80%	69 kW	4	57.2 kW	11.6 kW
3-phase 415 V, 32 A	100%	44 kW	2	28.6 kW	15.1 kW

Source: NVIDIA, Data Center Best Practices with DGX B200 [5].

The NVL72 has the same structure at a larger scale. Its rack uses eight power shelves, each with six 5.5 kW power supplies providing 33 kW, arranged for N+N redundancy, against rack consumption of approximately 120 kW [7]. Eight shelves of 33 kW give 264 kW of installed shelf capacity, or 132 kW with half held as redundancy (derived). A contract should say whether the kilowatts it prices are the rack's typical draw, its rated capacity or the provisioned circuit.

9. From megawatts to GPUs

Site power converts to GPU counts with a simple ratio. In July 2025, IREN stated that its 50 MW site could host more than 20,000 air-cooled Blackwell GPUs, assuming a power usage effectiveness of 1.1 and a power draw of 1.93 kW per GPU including ancillary power draw, based on NVIDIA's B200 reference architecture [19]. Dividing 50 MW by 1.93 kW and by 1.1 gives about 23,550 GPUs (derived), consistent with the stated figure.

The same ratio shows why facility efficiency matters for capacity agreements priced on IT load. A block sold as IT kilowatts needs more grid power behind it, by the site's power usage effectiveness. Uptime's 2025 survey found that facilities commissioned within five years averaged 1.48 [1].

10. Floor, weight and access

Rack-scale systems are heavy. Dell's specification sheet for its GB300 NVL72-based XE9712 lists a cabinet wet weight of 1,590 kg on a rack 750 mm wide and 1,200 mm deep, 1,340 mm with stabilizers [15]. Over the rack's own footprint, that is about 1,770 kg per square meter (derived). Supermicro lists its GB200 NVL72 rack at 2,236 mm high, 600 mm wide and 1,068 mm deep [9]. An air-cooled DGX B200 weighs more than 130 kg on NVIDIA's estimate [5].

ASHRAE's white paper adds two site checks: whitespace must accommodate the dimensions and weights of the racks, and load paths for heavy racks, including elevators, ramps, thresholds and corridors, must be rated for rolling loads [13]. Aisle pitch may also need extra space where rear-door heat exchangers are used [13].

11. Facility standards and scale

NVIDIA's DGX SuperPOD reference architecture for GB200 says the data center should generally meet or exceed Uptime Institute Tier 3 design standards, or TIA-942-B Rated 3 or EN 50600 Availability Class 3, including concurrent maintainability and no single point of failure [8]. It sizes a scalable unit at 8 DGX GB200 racks with a thermal design power of 1.2 MW [8]. For the air-cooled DGX B200, a scalable unit is up to 32 systems at 458 kW [5].

12. The power path ahead

NVIDIA's May 2025 technical blog post on 800 V DC architecture states that today's racks rely on 54 V DC distribution from rack-mounted power shelves, and that as racks exceed 200 kW this approach begins to hit physical limits [16]. GB200 NVL72 and GB300 NVL72 racks use up to eight power shelves, and at megawatt scale the same approach would consume up to 64U of rack space for power shelves; a 1 MW rack at 54 V would need up to 200 kg of copper busbar [16]. NVIDIA's plan is to convert 13.8 kV AC grid power to 800 V DC at the data-center perimeter with industrial-grade rectifiers, removing intermediate conversion steps, and to support 1 MW IT racks starting in 2027 [16]. NVIDIA also states that 800 V DC at the facility level introduces new challenges in safety, standards and workforce training [16].

Implications

For buyers. Specify the cooling grade and the rack power figure in both the purchase and the capacity agreement. Check the system's own site guide before choosing a cooling option: NVIDIA's guidance for the DGX B200 rules out rear-door and in-row cooling as typical choices [5], while NVL72 racks need a liquid loop and an air path [7].

For operators and builders. The installed base is still in single digits and low double digits [1, 2], and NVL72 racks run more than ten times higher [7, 10]. Phased builds, in which power and cooling are added as workloads become clear, are one response Uptime reports among operators [2]. Facility water class, CDU topology and heat rejection have to be chosen together [3, 12, 13].

For lenders. A liquid-cooled rack is valuable where a matching site exists. Collateral that can only run on a liquid loop at a given water class can move to fewer sites than an air-cooled server, given how few facilities run at those densities [1, 2]. Capacity agreements and site specifications belong in the credit file with the hardware.

For researchers. Compare compute costs on a per-kW basis and on a per-GPU-hour basis, and keep the two apart. Record which rack power figure a cost estimate uses, since vendor figures for the same rack range from about 120 kW to 135 kW [7, 9].

Limits of this analysis

- Figures are vendor-published specifications, reference designs and survey results. No system was measured for this report.
- Vendor power figures describe reference configurations. Actual draw depends on the build, the workload and firmware settings.
- Per-GPU power divides system power by GPU count and includes CPUs, memory, networking and fans in each system.
- CBRE's rate is an average asking rate for 250 to 500 kW requirements in primary North American wholesale markets. It is not a transaction price, does not separate liquid-cooled capacity and excludes energy. CBRE's report gives $196.25 per kW per month in its pricing section and $195.94 in its summary; this report uses $196.25.
- Uptime's 2026 figures are taken from Data Center Knowledge's report on the survey.
- Airflow and floor-loading figures are simple derivations for illustration. Site design needs the vendor's site preparation guide and an engineer's assessment.
- NVIDIA's 800 V DC plan is the company's stated plan and dates.
- Nothing in this report is engineering certification, investment advice or a recommendation to buy any product.

Sources

1. Uptime Institute Global Data Center Survey 2025. Uptime Institute, 2025. https://datacenter.uptimeinstitute.com/rs/711-RIA-145/images/2025.Annual.Survey.Report.pdf

2. AI Drives Data Center Uncertainty in Uptime's 2026 Survey (Sean Michael Kerner). Data Center Knowledge, 31 July 2026. https://www.datacenterknowledge.com/ai-data-centers/ai-drives-data-center-uncertainty-in-uptime-s-2026-survey

3. Chill Factor: NVIDIA Blackwell Platform Boosts Water Efficiency by Over 300x (Marc Hamilton). NVIDIA Blog, 22 April 2025. https://blogs.nvidia.com/blog/blackwell-platform-water-efficiency-liquid-cooling-data-centers-ai-factories/

4. Introduction to the NVIDIA DGX H100/H200 System, NVIDIA DGX H100/H200 System User Guide. NVIDIA, retrieved 9 October 2026. https://docs.nvidia.com/dgx/dgxh100-user-guide/introduction-to-dgxh100.html

5. Data Center Best Practices with DGX B200. NVIDIA, retrieved 9 October 2026. https://docs.nvidia.com/dgx-pdf/nvidia-dgx-superpod-data-center-best-practices-with-dgx-b200.pdf

6. Introduction to NVIDIA DGX B300, NVIDIA DGX B300 System User Guide. NVIDIA, retrieved 9 October 2026. https://docs.nvidia.com/dgx/dgxb300-user-guide/introduction-to-dgxb300.html

7. Hardware, NVIDIA DGX GB Rack Scale Systems User Guide. NVIDIA, last updated 17 September 2026. https://docs.nvidia.com/dgx/dgxgb200-user-guide/hardware.html

8. NVIDIA DGX SuperPOD: Next Generation Scalable Infrastructure for AI Leadership, Reference Architecture (DGX GB200). NVIDIA, retrieved 9 October 2026. https://docs.nvidia.com/dgx-superpod/reference-architecture-scalable-infrastructure-gb200/latest/_downloads/238ab33ee345a7241fc2f15eedc2ff6c/RA11338001-DSPGB200-ReferenceArch.pdf

9. Supermicro NVIDIA GB200 NVL72 Datasheet. Super Micro Computer, Inc., 2025. https://www.supermicro.com/datasheet/datasheet_SuperCluster_GB200_NVL72.pdf

10. Schneider Electric Announces New Reference Designs, Featuring Integrated Power Management and Liquid Cooling Controls, Supporting NVIDIA Mission Control and NVIDIA GB300 NVL72. Schneider Electric, 18 September 2025. https://www.se.com/ww/en/about-us/newsroom/news/press-releases/schneider-electric-announces-new-reference-designs-featuring-integrated-power-management-and-liquid-cooling-controls-supporting-nvidia-mission-control-and-nvidia-gb300-nvl72-68ca8fdd5e6c6f9d3a096133/

11. Supermicro Ramps Full Production of NVIDIA Blackwell Rack-Scale Solutions with NVIDIA HGX B200. Super Micro Computer, Inc., 5 February 2025. https://www.supermicro.com/en/pressreleases/supermicro-ramps-full-nvidia-blackwell-rack-scale-solutions-nvidia-hgx-b200

12. Best Practices Guide for Energy-Efficient Data Center Design, revised July 2024. U.S. Department of Energy, Federal Energy Management Program, and National Renewable Energy Laboratory, July 2024. https://www.energy.gov/sites/default/files/2024-07/best-practice-guide-data-center-design_0.pdf

13. Emergence and Expansion of Liquid Cooling in Mainstream Data Centers. ASHRAE Technical Committee 9.9, 2021. https://www.ashrae.org/file%20library/technical%20resources/bookstore/emergence-and-expansion-of-liquid-cooling-

in-mainstream-data-centers_wp.pdf

14. NVIDIA Contributes Blackwell Platform Design to Open Hardware Ecosystem, Accelerating AI Infrastructure Innovation. NVIDIA Newsroom, 15 October 2024. https://nvidianews.nvidia.com/news/nvidia-contributes-blackwell-platform-design-to-open-hardware-ecosystem-accelerating-ai-infrastructure-innovation

15. PowerEdge XE9712 Specification Sheet. Dell Technologies, retrieved 9 October 2026. https://www.delltechnologies.com/asset/en-us/products/servers/technical-support/poweredge-xe9712-spec-sheet.pdf

16. NVIDIA 800 VDC Architecture Will Power the Next Generation of AI Factories (Mathias Blake, Martin Hsu, Ivan Goldwasser, Harry Petty and Jared Huntington). NVIDIA Technical Blog, 20 May 2025. https://developer.nvidia.com/blog/nvidia-800-v-hvdc-architecture-will-power-the-next-generation-of-ai-factories/

17. North America Data Center Trends H2 2025. CBRE, 25 February 2026. https://www.cbre.com/insights/books/north-america-data-center-trends-h2-2025

18. Digital Realty Trust, Inc., Annual Report on Form 10-K for the fiscal year ended December 31, 2025. U.S. Securities and Exchange Commission, filed 13 February 2026. https://www.sec.gov/Archives/edgar/data/1297996/000110465926015365/dlr-20251231x10k.htm

19. IREN Expands AI Cloud with 2.4k NVIDIA Blackwell GPUs, Exhibit 99.1 to Form 8-K. IREN Limited, 3 July 2025. https://www.sec.gov/Archives/edgar/data/1878848/000114036125024747/ef20051566_ex99-1.htm

About this report

Published by Rillor on 24 November 2026. Superintelligence (SI) tools assist research and drafting, and every figure cites its source. Rillor builds and secures SI systems, supplies the data they learn from and delivers the infrastructure they run on. We apply SI to financial and compute markets.

Cite as: Rillor. Rack power and liquid cooling: what NVL72 changed. 24 November 2026. https://rillor.com/insights/rack-power-and-liquid-cooling

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.

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

Work with Rillor

Send the GPU systems you plan to run, the block size and cooling you need, and the month you need power. Contact us at rillor.com/contact?type=capacity or contact@rillor.com. More on data-center capacity: rillor.com/compute/capacity