Rillor

EXPLAINER

Data-center Tier levels: what Tier III and IV cover

Uptime Institute's four data center Tier levels set maintenance and fault-tolerance criteria with no uptime percentage attached, and NVIDIA says dual feeds give a DGX B200 rack only N.



Contents

01 Summary

02 Background

03 Method

04 Findings

05 Implications

06 Limits of this analysis

07 Sources

Summary

- Uptime Institute defines four data center Tier levels by capacity components and distribution paths: Tier I (Basic Capacity), Tier II (Redundant Capacity Components), Tier III (Concurrently Maintainable) and Tier IV (Fault Tolerant) [1, 3]. Component counts do not set the Tier, and Uptime says Tier IV is possible with N+1 components [23].
- No Tier carries an uptime percentage. The 1.6 hours and 99.98% quoted for Tier III come from field data on 16 sites in a 2008 table that Uptime titled as unrelated to Tier requirements, and Uptime removed expected downtime from the standard in 2009 [5, 22, 23].
- Utility feeds neither predict nor influence Tier level. The standard treats the engine-generator plant as the reliable power source [4, 23].
- Design awards issued after 1 January 2014 expire after two years, and constructed-facility awards last until the facility is modified [8]. On Uptime's public list, 869 of 1,562 Tier entries show a design award as their only Tier award [9].
- TIA-942 ratings and EN 50600 classes are separate schemes with different wording and routes to conformity. EN 50600 has no formal certification [10, 12, 13].
- NVIDIA says a DGX B200 rack fed from two discrete sources delivers only N [14]. Uptime's pooled 2025 and 2026 surveys associate 2N power and cooling with fewer reported outages, a correlation its authors caution against reading as cause [26].

Background

Uptime Institute's Tier Standard: Topology "describes criteria to differentiate four classifications of site infrastructure topology based on increasing levels of redundant capacity components and distribution paths" [3]. Uptime says it is the only organization that can formally certify data centers against the Tier Classification System [22]. It reports over 4,300 awards in more than 120 countries across its Tier, Management and Operations and Efficient IT programs [9].

Tier language now appears in broker rent benchmarks [35], SEC filings [20, 21] and NVIDIA's reference architectures for rack-scale GPU systems [18, 19]. The 2008 table that carries the old downtime figures lists typical densities of a few kilowatts per cabinet, over 3 kW at Tier III and over 4 kW at Tier IV [5]. A GB200 NVL72 rack draws approximately 120 kW [16], and GPU servers bring their own power supply redundancy, which interacts with the facility's [14]. This report sets out what a Tier covers, how the other schemes compare and what a buyer of capacity for GPU systems needs to confirm.

Method

What was examined. The Tier definitions, the origin of the availability figures attached to them, the treatment of utility power, the award types and their validity, TIA-942 and EN 50600, survey data on redundancy and outages, and NVIDIA's power guidance for DGX and rack-scale systems.

Sources. Standard owners come first: Uptime's own pages [1, 2, 3, 6, 7, 8, 9, 10], its Journal articles, survey report and press releases [22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 32, 33], TIA's certification pages [11, 12] and the CEN, CENELEC and ETSI coordination group's summary of EN 50600 [13]. The current Tier Standard: Topology sits behind a download agreement and was not used. The 2009 edition and the 2008 white paper it restated were read from third-party copies, as historic texts [4, 5]. NVIDIA's documentation [14, 15, 16, 17, 18, 19], two SEC filings [20, 21], CBRE [35] and a guest post on TIA's site [34] complete the set.

How. Definitions and figures are recorded as each source states them, with dates. Uptime's Tier Classification System page lists redundant equipment under Tier I, which conflicts with Tier I as Basic Capacity in Uptime's other texts, so the definitions here follow the Tier Certification overview, the 2009 standard and the Journal [1, 2, 4, 22]. Uptime's public awards list was read on 7 November 2026 and each entry counted by Tier level and listed award types [9]. Derived figures are marked as derived. The report uses no Rillor data.

Findings

1. The four Tiers as Uptime defines them

Tier	Name	Uptime's definition
I	Basic Capacity	Site-wide shutdowns are required for maintenance or repair work. Capacity or distribution failures will impact the site.
II	Redundant Capacity Components	Site-wide shutdowns for maintenance are still required. Capacity failures may impact the site. Distribution failures will impact the site.
III	Concurrently Maintainable	Each and every capacity component and distribution path in a site can be removed on a planned basis for maintenance or replacement without impacting operations. The site is still exposed to an equipment failure or operator error.
IV	Fault Tolerant	An individual equipment failure or distribution path interruption will not impact operations. A Fault Tolerant site is also Concurrently Maintainable.

Source: [1].

The 2009 edition of Tier Standard: Topology summarizes the requirements in its Table 1 [4].

Requirement	Tier I	Tier II	Tier III	Tier IV
Active capacity components for the IT load	N	N+1	N+1	N after any failure
Distribution paths	1	1	1 active and 1 alternate	2 simultaneously active
Concurrently maintainable	No	No	Yes	Yes
Fault tolerance	No	No	No	Yes
Compartmentalization	No	No	No	Yes
Continuous cooling	Load density dependent	Load density dependent	Load density dependent	Class A

Source: [4].

Tier III and Tier IV. Tier III adds a redundant delivery path for power and cooling, so each component can be shut down and maintained without affecting IT [22]. Uptime requires critical power distribution, the UPS output and below, to be active/active at Tier III: "if a rack receives dual power from two separate power distributions, they must both normally be active" [23]. Mechanical systems carry no such requirement at Tier III [23]. Tier IV requires "multiple, independent, physically isolated systems", compartmentalization, continuous cooling and N capacity after any infrastructure failure [4], with all systems active/active [23]. Operating the fire alarm, fire suppression or emergency power off may still disrupt a Tier IV site [4]. Tier II, by contrast, allows concurrent maintenance of capacity components only [23].

Rules for every Tier. "The site's Tier rating is the lowest of the individual subsystem ratings", and there are no partial or fractional Tiers [4]. Uptime says terms such as "Tier III plus" have no basis in Tiers [23]. The Tiers prescribe no specific technology or design [22]. Uptime also says a Tier IV solution is not always better than a Tier II solution, because the business value of the application may be lower, and that costs and operational complexity rise with Tier level [22].

2. Why the standard carries no uptime percentage

The figures attached to Tiers come from Uptime's 2008 white paper. It reports that Uptime "measured actual performance of 16 data centers" over periods of up to 10 years, with uptime measured from the end user's point of view [5]. Its Table 2 is titled "Common Attributes Found in Data Centers that are Unrelated to Tier Requirements" [5].

Row in Table 2	Tier I	Tier II	Tier III	Tier IV
Annual site-caused end-user downtime (based on field data)	28.8 hours	22.0 hours	1.6 hours	0.8 hours
Resulting end-user availability	99.67%	99.75%	99.98%	99.99%
Ultimate kW per cabinet (typical)	Under 1 kW	1 to 2 kW	Over 3 kW	Over 4 kW
First deployed	1965	1970	1985	1995

Source: [5].

The paper says availability "calculations do not determine the Tier Classification", and that even a Tier IV site "will not satisfy an IT requirement of Five Nines (99.999 percent) uptime" [5]. The 2009 standard repeats the percentages only in a commentary marked "not part of the Data Center Site Infrastructure Tier Standard: Topology" [4]. Uptime says that in 2009 it removed all references to expected downtime per year, and that

the current standard assigns no availability predictions to Tier levels [22]. Its reason: "operational behaviors can have a huge impact on site availability regardless of the technical prowess of the design and build" [22]. A 2017 Uptime article adds that the figures "were never a part of the Tier definitions" [23].

The 99.982% figure is arithmetic: 1.6 hours out of 8,760 is 0.018% (derived from [5]). The figure remains in use. WhiteFiber's investor presentation, filed with the SEC on 14 September 2026, lists "99.982% Uptime and No More than 1.6 Hours of Downtime Annually" on a slide about Tier III data centers and names Uptime Institute as the source [20]. The same deck describes a "Tier III - compliant retrofit data center" without saying which Uptime award, if any, the site holds [20].

3. Utility feeds, generators and the Tier

Under Tier Standard: Topology, the engine-generator plant is "the only reliable source of power for a data center", because utility power is subject to unscheduled interruption [23]. "The number of utility feeds, substations, and power grids that provide public power to a data center neither predicts nor influences Tier level", and utility power is not required for Tiers [23]. The 2009 text calls the utility "an economic alternative" and says Tier III and IV generators "shall not have a limitation on consecutive hours of operation when loaded to 'N' demand" [4].

In October 2025 Uptime described a GPU campus "built to be concurrently maintainable, powered by two utility feeds" [33]. Under the standard, the number of feeds does not change the Tier [23].

Uptime estimates GB200 NVL72 racks can climb from around 60 to 70 kW to more than 150 kW, against a maximum specification of 132 kW, and expects engine generator sets to be stressed more than other power equipment by such swings [32]. It suggests UPS configurations such as N+2 to allow maintenance without repeated overloads [32]. Power was the top cause of impactful outages in Uptime's 2025 survey, at 45% of incidents against 54% in 2024 [27], and its 2026 outage analysis says failures involving UPS systems, transfer switches and generators are dominant [28].

4. Design, constructed-facility and operational certification

Uptime issues Tier awards in three steps, each a prerequisite for the next [6, 7].

Award	What it covers	Validity
Design Documents (TCDD)	The functionality and capacity in the design specifications [6]	Two years from award, for awards issued after 1 January 2014 [6, 8]
Constructed Facility (TCCF)	The built site, including equipment installed at the site visit and the performance demonstration tests [8]	Until the facility is modified, including changes to capacity components or distribution paths [8]

Award	What it covers	Validity
Operational Sustainability (TCOS)	Management and operations, building characteristics and site location [7]	Three years for Gold, two for Silver, one for Bronze [8]

Uptime says more than 85% of data center designs are incorrectly executed in construction, so as-built performance "may be drastically different" from a certified design [22]. Its consultants find changes after design certification "in almost every site visit" [23]. The constructed-facility review ends with a "pull the plug" test [22], and certifications can be revoked if unreviewed changes compromise concurrent maintainability or fault tolerance [23]. The foil "specifies the relevant phase, room, pod, module, or other nomenclature" covered [8], and certification of a defined phase "is limited to that specific phase alone" [25].

Uptime's public awards list. On 7 November 2026 the list had 1,950 entries, of which 1,562 carry a Tier level [9]. The rest are Management and Operations or Efficient IT awards [9].

Tier	Entries	Design award only	Constructed-facility award	Operational award, no constructed-facility award
Tier II	30	22	8	0
Tier III	1,274	699	559	16
Tier IV	258	148	106	4
All	1,562	869	673	20

Source: tally of [9], 7 November 2026.

Of the 104 Tier entries located in the United States, 67 list a constructed-facility award and 35 a design award only (tally of [9]). The list shows no award dates, so lapsed design awards cannot be counted, and one facility can appear once per phase or room. The counts describe entries, which differ from the number of certified data centers.

Tier language in filings. Corvex's 2025 annual report describes "a Tier III data center facility" and "our current tier III-designed facility", with "N+N power paths, N+1 uninterruptible power supply and generators, N+1 cooling" [21]. It names no Uptime award or date.

5. TIA-942 ratings and EN 50600 classes

Separation. Uptime and TIA agreed in 2014 on "a clear separation between their respective benchmarking systems", and "any reference to the TIA rating of a data center may not include the word Tier" [23]. Uptime's announcement says TIA agreed to remove the word Tier from ANSI/TIA-942 [10].

TIA-942. Maintained by TIA Engineering Committee TR-42 and now at revision TIA-942-C, the standard covers site location, telecommunications infrastructure, the building, electrical and mechanical infrastructure, fire safety and physical security [11]. TIA-licensed certification bodies issue Design, Facilities and Ready certifications, each at four rating levels [11, 12].

Rating	TIA's definition (abridged)
Rated-1	Single capacity components and a single, non-redundant distribution path
Rated-2	Redundant capacity components and a single, non-redundant distribution path
Rated-3	Redundant components and multiple independent paths. "Typically, only one distribution path serves the computer equipment at any time." Concurrently maintainable
Rated-4	Redundant components and multiple independent paths "which all are active". Concurrently maintainable and tolerates one fault anywhere without downtime

Source: [12].

A Design certification lasts one year, and a Facilities certification three years with surveillance audits by the end of years one and two [12]. Rated-3 says that typically only one distribution path serves the equipment at any time [12], while Uptime requires both critical power paths to a rack to be normally active at Tier III [23].

EN 50600. The CEN, CENELEC and ETSI coordination group says the series defines four design availability levels each for power, environmental control and cabling, with a data centre's overall level set by the lowest of the three [13]. Class 3 power distribution and environmental control use multiple paths with concurrent repair and operation, and Class 4 multiple paths that are fault tolerant except during maintenance [13]. On conformity: "Self-declaration, rather than formal 'certification', is all that is available" [13].

GPU racks. In a guest post on TIA's site, Barry Elliott argues that the standard does not currently address GPU computing architectures. He notes that each liquid-cooled compute unit has "only one water feed pipe going in and only one coming out", and proposes treating such a rack as "one single computer" [34]. TIA says guest posts carry the authors' views [34].

NVIDIA. The GB200 reference architecture says: "Generally, the data center should meet or exceed Uptime Institute Tier 3 design standards, or alternatively the TIA942-B Rated 3 or EN50600 Availability Class 3 design standards, including concurrent maintainability and no single point of failure" [18]. The Vera Rubin NVL72 reference architecture repeats it [19]. NVIDIA treats the three as alternatives. Their definitions and conformity routes differ in detail [12, 13, 23].

6. N, N+1 and 2N in practice

Uptime's October 2026 analysis defines N as capacity sized to the expected peak load, N+1 and N+2 as one or two extra components, and 2N as the full capacity duplicated in two independent paths or systems [26]. Uptime calls N terminology "industry shorthand" and says added equipment "does not determine or guarantee achievement of any specific Tier level, because Tiers also includes evaluation of distribution pathways and other system elements" [23]. Tier IV is possible "with just N+1 components, depending on how they are configured and connected to redundant distribution pathways" [23]. A 2024 Uptime article describes fault tolerance from distributed N+1 UPS designs whose total capacity is "far less than twice the entire workload" [30].

In Uptime's 2025 survey, 42% of respondents named 2N as the usual level of redundant power equipment at their primary data center, and 41% named N+1 [27]. In 2020 Uptime reported a shift from N+1 to N+2, and said large cloud and internet operators commonly stipulate N+2 facilities from colocation partners [29]. In May 2026 it wrote that higher redundancy consumes a larger share of a site's provisioned capacity, and that for GPU deployments the permitted electrical capacity is often the limit [31].

7. GPU servers and redundancy at the rack

What the Tiers expect of IT equipment. Tier III and IV designs rest on dual-corded equipment and active-active dual paths [24]. Uptime's Fault Tolerant Power Compliance Specification asks equipment to keep running "from A and B power sources or from A alone or from B alone", and says equipment with an odd number of cords generally will not comply [24]. Rack-mounted transfer switches are allowed for odd-corded equipment, though their isolation or failure may cause an outage for the racks or devices behind them [23]. A single large static transfer switch solution does not meet Tier III or IV criteria [24].

DGX B200. Each 14.3 kW system has six power supplies, five must be energized, and the system "will not operate if more than one power supply unit is de-energized, regardless of upstream power redundancies" [14]. NVIDIA concludes that rack provisioning "should be understood as delivering ONLY N, even if the provisioning is comprised of dual feeds from discrete sources" [14]. Split three and three between two sources, a lost source leaves three supplies, two short of five (derived from [14]).

Layout	NVIDIA's description
Two discrete sources, each feeding half the systems (N)	A failure of either UPS impacts only half of the systems. Preferred when only two UPS sources reach the rack
Single source (N)	A UPS failure affects all of the systems
Two sources interleaved into each system (N)	A failure of either UPS impacts all of the systems. "This model is not recommended"
Dual sources with a static transfer switch (N+1)	The switch becomes a single point of failure. NVIDIA has not tested or validated any STS product with DGX B200 and does not currently recommend this solution
"6 to make 5" sources (N+1)	A failure of any single UPS source does not impact the systems. "This is the optimal configuration for high availability", though many sites are not equipped to provision six discrete UPS sources and it may limit racks to two systems

Source: [14].

Newer systems. NVIDIA's power designs change by generation. Each is given here in NVIDIA's words.

System	Power design as NVIDIA states it
DGX B200	Six PSUs, five must be energized [14]
DGX B300	12 PSUs "configured for N+N redundancy", minimum six required [15]
GB200 NVL72 rack	About 120 kW, eight power shelves of six 5.5 kW PSUs "that provide N+N redundancy" [16]. The reference architecture calls each shelf's PSUs "configured as N redundancy" [17]
Vera Rubin NVL72 rack	Four shelves of six 18.3 kW PSUs, "with only three required for full rack power, allowing for N+N redundancy in the data center on every rack" [19]

The B300 can lose six of its 12 supplies, against one of six for the B200 (derived from [14, 15]). The documents cited do not describe how supplies or shelves are wired to facility sources, so whether a rack rides through the loss of a whole source is a question for the vendor and the site.

Cooling. Elliott writes that a Rating 4 design "could be envisaged" only up to the rack's distribution manifold, with Rating 3 from that point [34].

8. Redundancy and outage rates in survey data

Uptime's October 2026 analysis combined its 2025 and 2026 surveys into a sample of over 1,200 respondents and compared the share reporting at least one impactful outage in three years by redundancy configuration [26].

Configuration	Cooling	Power
N	70%	68%
2N	39%	42%

Source: [26].

N+1 and N+2 fell between [26]. Serious outages, categories 3 to 5 on Uptime's scale, were reported by about 19% to 22% of respondents using N and about 8% to 11% of those using 2N [26].

Reduction in likelihood of at least one financially impactful (category 3 to 5) outage in three years, compared with N	N+1	N+2	2N
Cooling	About 36%	About 53%	About 63%
Power	About 23%	About 37%	About 44%

Source: [26].

Against an N+1 baseline, 2N cooling was 42% less likely to have a financially impactful outage, and N+2 26% less likely [26]. The authors caution that "Correlation is not causation", that the N results rest on small samples, and that redundancy cannot be separated from other traits of resilient facilities such as operational maturity [26]. They call the results "indicative reference points rather than predictions for an individual facility" [26]. Uptime also reports that per-site outage frequency has declined for the fifth consecutive year, at a slowing pace [28]. None of the sources reviewed reports outage rates by Tier level.

9. Where reference designs set the bar

NVIDIA's GB200 reference architecture sizes a scalable unit at 1.2 MW and asks for Tier 3 or an equivalent [18]. The Vera Rubin NVL72 version sizes a unit of 8 systems at about 2 MW and repeats the line [19]. The DGX B200 best-practices deck, with scalable units of up to 32 systems at 458 kW, sets no Tier level and handles redundancy at the rack [14].

Uptime reports that large cloud operators run mostly concurrently maintainable facilities [29], and former crypto miners building GPU sites with hyperscale engineers deliver facilities that "are concurrently maintainable" [33]. A concurrently maintainable site still exposes IT to an equipment failure or operator error [1], and under NVIDIA's guidance a dual-fed DGX B200 rack in it receives N [14].

10. Tier language in asking rents and filings

CBRE quotes North American rents as "asking rates for 250+ kW at N+1/Tier III requirements" [35]. On that basis, the average for a 250 to 500 kW requirement in primary wholesale markets rose 6.6% year over year to $196.25 per kW per month in the second half of 2025, while another section of the page gives $195.94, up 6.5% [35]. The footnote pairs an N+1 count with a Tier name to describe the requirement priced. It does not say the capacity holds an Uptime award. The filings in findings 2 and 4 make the same pairing without naming an award [20, 21].

Implications

For buyers and procurement teams. The sources point to the documents to request: the award type, its date and the scope on the foil [6, 8], and for a constructed-facility award, confirmation that the infrastructure has not changed since [8, 25]. A TIA-942 claim comes with a certification body and a term [12], and an EN 50600 claim rests on self-declaration or an assessment [13]. For GPU systems, the rack matters as much as the hall: how many power supplies each system needs, how they connect to the facility's sources, and whether both critical power paths are active to the rack [14, 15, 23]. Rillor's capacity agreements reserve power and space in 250 kW, 500 kW and 1 MW blocks, liquid-cooled or rear-door cooled, on a 24-month standard term priced per kW-month, sized to the GPU systems by power draw and cooling grade.

For finance teams and lenders. Tier claims age on fixed rules: design awards lapse after two years, constructed-facility awards end on modification and operational awards run one to three years [8]. Higher redundancy consumes more of a site's provisioned capacity [31], and its link to fewer outages is a correlation [26]. None of the sources reviewed quantifies a rent difference by Tier or award type.

For operators and site engineers. The server's power supply design is a facility design input: dual sources give a DGX B200 rack N, and "6 to make 5" reaches N+1 at a cost in sources and rack density [14]. Uptime expects generator sets, the reliable source under the Tier Standard, to bear more stress from the power swings of model training than other power equipment [23, 32].

For researchers. Outage data exist by redundancy level [26], and the awards list supports counts by Tier and award type, without dates [9].

Limits of this analysis

- The current edition of Tier Standard: Topology was not read. The 2009 edition and the 2008 white paper are historic texts, read from third-party copies [4, 5].
- No source reviewed says how Uptime certifies halls of GPU systems with N+1 internal power supplies, or how its current standard treats liquid-cooled racks and in-rack power shelves.
- The awards list counts are entries on one day, with no dates, and one facility can appear several times [9].
- TIA's list of certified data centers was not tallied, and EN 50600 was described from a public summary [13].
- The survey results are self-reported and correlational, with small N samples [26].
- NVIDIA's N+N descriptions are quoted as written. The documents do not state rack feed topology, and this report does not infer it [15, 16, 17, 19].
- The report uses no Rillor data. It is research, not engineering, legal or investment advice.

Sources

1. Tier Certification. Uptime Institute, viewed 24 October 2026. https://uptimeinstitute.com/tier-certification

2. Tier Classification System. Uptime Institute, viewed 24 October 2026. https://uptimeinstitute.com/tiers

3. Tier Standard: Topology (publication abstract). Uptime Institute, viewed 24 October 2026. https://uptimeinstitute.com/publications/asset/tier-standard-topology

4. Data Center Site Infrastructure Tier Standard: Topology (2009 edition, copy hosted by a third party). Uptime Institute Professional Services, 2009. https://35904.fs1.hubspotusercontent-na2.net/hub/35904/docs/tierstandard.pdf

5. Tier Classifications Define Site Infrastructure Performance (white paper, copy hosted by a third party). W. Pitt Turner IV, John H. Seader, Vince Renaud and Kenneth G. Brill, Uptime Institute, 2008. https://fsta.dk/wp-content/uploads/2023/07/Tier-Classifications-Define-Site-Infrastructure-Performance.pdf

6. Data Center Design Document Certification. Uptime Institute, viewed 24 October 2026. https://uptimeinstitute.com/tier-certification/design

7. Tier Certification Operational Sustainability. Uptime Institute, viewed 24 October 2026. https://uptimeinstitute.com/tier-certification/operations

8. Awards Terms and Conditions. Uptime Institute, viewed 24 October 2026. https://uptimeinstitute.com/awards-terms-and-conditions

9. Uptime Institute Issued Awards (public awards list). Uptime Institute, viewed 7 November 2026. https://uptimeinstitute.com/uptime-institute-awards/list

10. Uptime Institute and TIA agree on clear separation (press release). Uptime Institute, viewed 24 October 2026. https://uptimeinstitute.com/about-ui/press-releases/uptime-tia

11. TIA-942 Certification. Telecommunications Industry Association, viewed 24 October 2026. https://tiaonline.org/products-and-services/tia942certification/

12. TIA 942 Certifications & Ratings. Telecommunications Industry Association, viewed 24 October 2026. https://tiaonline.org/products-and-services/tia942certification/tia-942-certifications-ratings/

13. Review of standardisation activities: Energy Management and Environmental Viability of Data Centres, Edition 11. CEN/CENELEC/ETSI Coordination Group on Green Data Centres, November 2024. https://www.cencenelec.eu/media/CEN-CENELEC/AreasOfWork/CEN%20sectors/Digital%20Society/Green%20Data%20Centres/2024/brochuredatacentre-_standardizationedition11_2024.pdf

14. Data Center Best Practices with DGX B200. NVIDIA, 2024. https://docs.nvidia.com/dgx-pdf/nvidia-dgx-superpod-data-center-best-practices-with-dgx-b200.pdf

15. Introduction to NVIDIA DGX B300 Systems (DGX B300 User Guide). NVIDIA, updated 21 October 2026. https://docs.nvidia.com/dgx/dgxb300-user-guide/introduction-to-dgxb300.html

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

17. Key Components of the DGX SuperPOD (DGX GB200 reference architecture). NVIDIA, updated 19 November 2025. https://docs.nvidia.com/dgx-superpod/reference-architecture-scalable-infrastructure-gb200/latest/dgx-superpod-components.html

18. NVIDIA DGX SuperPOD: Next Generation Scalable Infrastructure for AI Leadership, Reference Architecture Featuring NVIDIA DGX GB200 (RA-11338-001). NVIDIA, 16 June 2025. https://docs.nvidia.com/dgx-superpod/reference-architecture-scalable-infrastructure-gb200/latest/_downloads/238ab33ee345a7241fc2f15eedc2ff6c/RA11338001-DSPGB200-ReferenceArch.pdf

19. NVIDIA DGX SuperPOD: Next Generation Scalable Infrastructure for AI Factories, Reference Architecture Featuring NVIDIA DGX Vera Rubin NVL72 Systems. NVIDIA, September 2026. https://docs.nvidia.com/dgx-superpod-reference-architecture-with-dgx-vera-rubin-nvl72.pdf

20. Investor Presentation of WhiteFiber, Inc. as of September 2026 (Exhibit 99.1 to Form 8-K). U.S. Securities and Exchange Commission, filed 14 September 2026. https://www.sec.gov/Archives/edgar/data/2042022/000121390026099424/ea030525701ex99-1.htm

21. Corvex, Inc. Annual Report on Form 10-K for the year ended 31 December 2025. U.S. Securities and Exchange Commission, filed 31 March 2026. https://www.sec.gov/Archives/edgar/data/1734750/000121390026036609/ea0278320-10k_corvex.htm

22. Explaining the Uptime Institute's Tier Classification System (April 2021 Update). Matt Stansberry, Uptime Institute Journal, 30 September 2014, updated April 2021. https://journal.uptimeinstitute.com/explaining-uptime-institutes-tier-classification-system/

23. Myths and Misconceptions Regarding the Uptime Institute's Tier Certification System. Kevin Heslin, Uptime Institute Journal, 20 July 2017. https://journal.uptimeinstitute.com/myths-and-misconceptions-regarding-the-uptime-institutes-tier-certification-system/

24. Dual-Corded Power and Fault Tolerance: Past, Present, and Future. Uptime Institute Journal, 18 August 2014. https://journal.uptimeinstitute.com/dual-corded-power-details-change-theme-remains/

25. Tier Certification for Modular and Phased Construction. Chris Brown, Uptime Institute Journal, 14 December 2015. https://journal.uptimeinstitute.com/modular-and-phased-construction/

26. Survey findings link higher redundancy to lower outage risk. Owen Rogers and Douglas Donnellan, Uptime Institute Journal, 14 October 2026. https://journal.uptimeinstitute.com/survey-findings-link-higher-redundancy-to-lower-outage-risk/

27. Uptime Institute Global Data Center Survey 2025. Uptime Institute, 2025. https://datacenter.uptimeinstitute.com/rs/711-

RIA-145/images/2025.Annual.Survey.Report.pdf

28. Uptime Announces Annual Outage Analysis Report 2026 (press release). Uptime Institute, 13 May 2026. https://uptimeinstitute.com/about-ui/press-releases/uptime-announces-annual-outage-analysis-report-2026

29. Why data center operators are investing in more redundancy. Andy Lawrence, Uptime Institute Journal, 26 October 2020. https://journal.uptimeinstitute.com/datacenter-operators-redundancy/

30. Resiliency v low PUE: regulators a catalyst for innovation? Andy Lawrence, Uptime Institute Journal, 4 September 2024. https://journal.uptimeinstitute.com/resiliency-v-low-pue-regulators-a-catalyst-for-innovation/

31. Capacity allocation and the next generation of AI-era KPIs. Rand Talib, Uptime Institute Journal, 20 May 2026. https://journal.uptimeinstitute.com/capacity-allocation-and-the-next-generation-of-ai-era-kpis/

32. Electrical considerations with large AI compute. Daniel Bizo, Uptime Institute Journal, 19 November 2025. https://journal.uptimeinstitute.com/electrical-considerations-with-large-ai-compute/

33. Crypto mines are turning into AI factories. Max Smolaks, Uptime Institute Journal, 22 October 2025. https://journal.uptimeinstitute.com/crypto-mines-are-turning-into-ai-factories/

34. Aligning the AI High Performance Computing infrastructure with the ANSI/TIA-942 Ratings resilience model. Barry Elliott, Telecommunications Industry Association (guest blog), 10 December 2025. https://tiaonline.org/aligning-the-ai-high-performance-computing-infrastructure-with-the-ansi-tia-942-ratings-resilience-model/

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

About this report

Published by Rillor on 7 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. Data-center Tier levels: what Tier III and IV cover. 7 November 2026. https://rillor.com/insights/data-center-tier-levels-explained

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, Kubernetes 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