



ASHRAE W17 to W+ Liquid Cooling Classes for GPU Racks

Plan and build private AI compute with GPU Smith. We connect workload requirements with GPU hardware, networking, deployment and operating decisions so your infrastructure fits the work it must perform.

Published by GPU Smith · 2026-10-29 · ashrae w17 w27 w32 w40 w45 · ashrae liquid cooling classes · facility water temperature gpu rack · liquid cooling for gpu clusters · b300 liquid cooled · w45 facility water dry cooler gpu

Original article: <https://gpusmith.com/articles/en/ashrae-liquid-cooling-w-classes-facility-water-temperature>



In this article

1. The short answer
 2. Why facility water class matters more than rack kW
 3. From W1-W5 to W17-W+
 4. Class definitions table
 5. What compliance means
 6. Facility Water System vs Technology Cooling System
 7. What GPU vendors actually specify
 8. Chillers vs dry coolers vs hybrid
 9. Approach temperature and sizing the heat rejection
 10. Choosing W17 vs W27 vs W32+ in practice
 11. Commissioning and monitoring implications
 12. Specification checklist for RFPs
-

The short answer

ASHRAE's liquid cooling classes are named for the highest facility water supply temperature that IT equipment in that class must tolerate: W17 (17°C), W27 (27°C), W32 (32°C), W40 (40°C), W45 (45°C) and W+ (above 45°C). Every class has a lower limit of 2°C. The names were introduced with the fifth edition of *Thermal Guidelines for Data Processing Environments* in 2021, replacing the 2011 labels W1 to W5 and adding W40 as a new class ([ASHRAE TC 9.9, 2021](#)).

For a GPU build, the W class is the number that tells you whether the heat rejection plant needs chillers, can run on cooling towers or dry coolers alone, or can export heat. Two caveats matter more than the table. W classes describe the facility water system (FWS), not the coolant at the cold plates; a coolant distribution unit (CDU) between them adds an approach temperature. And GPU racks documented to run at 45°C supply need far more flow and pump head there than at 25°C or 30°C, so the class you can design to depends on that curve, not the headline maximum.

For the wider cooling picture (air, containment, chiller plant basics), see our [GPU data center cooling design guide](#).

Why facility water class matters more than rack kW

Rack power sets how much heat you remove; water temperature sets what each kilowatt costs and how many hours run without compressors.

Johnson Controls estimates that a typical chiller uses about 2 to 3 percent less power for every 1°C rise in chilled water temperature, and legacy chilled water set points sit at 42 to 45°F (6 to 7°C) ([DCD, 2024](#)). Warm enough water drops the chiller entirely: a 2026 National Laboratory of the Rockies study found 50°C/60°C inlet/outlet water eliminated chillers, cooling towers and heat recovery equipment in many modeled scenarios for a 1 MW inference facility, with up to 75 percent capital savings on that equipment ([Becker et al., ACEEE 2026](#)).

On the silicon side, hotter chips with lower allowable case temperatures push the other way. ASHRAE TC 9.9 warned in 2021 that direct liquid cooled products able to run on W45 "may be replaced by newer, higher-powered products only capable of W32 (W3) operation, and eventually cooling requirements may drive facility water temperatures down into the W27 (W2) range. Facilities should plan accordingly" ([ASHRAE TC 9.9, 2021](#)). The same paper notes that most cold plate solutions of the time could "live easily on W32 water with case temperatures in the 80°C–90°C range", but expected case temperature requirements to drop into the low 70s °C.

Schneider Electric's cooling general manager told DCD that at 700 W per accelerator, supply water "has to come down to about 80°F (27°C)", and that some customers ask for 20 to 24°C; Uptime Institute described facility water converging around 32°C ([DCD, 2024](#)). These are interviews, not measurements.

From W1-W5 to W17-W+

ASHRAE TC 9.9 created the liquid cooling classes in 2011, in collaboration with the HPC Energy Efficiency Working Group, to define facility water supply temperature ([Demetriou, 2025](#)). The original five classes, W1 to W5, had upper limits of 17°C, 27°C, 32°C, 45°C and over 45°C, each with a 2°C lower limit ([Upsite, 2024](#)).

The fifth edition of *Thermal Guidelines for Data Processing Environments* (2021) made three changes, as described by TC 9.9 itself ([ASHRAE TC 9.9, 2021](#)):

1. The classes were renamed "with the upper temperature limits incorporated in the name": W17 (previously W1), W27 (W2), W32 (W3), W40 (new), W45 (W4) and W+ (W5).
2. A new 40°C class was added. Upsite's account of the revision says the need for a 40°C class "wasn't debated": several manufacturers were designing around 40°C entering facility water, while few could run consistently above 45°C ([Upsite, 2024](#)).
3. Compliance was tightened (next section).

DCD dates the renaming to 2022, but TC 9.9 says the fifth edition was "released in 2021" (Upsite: March 2021). Use 2021.

Class definitions table

Fahrenheit values and typical plant equipment below are from a 2025 ASHRAE Distinguished Lecture by the TC 9.9 IT subcommittee chair ([Demetriou, 2025](#)). The equipment columns are "typical infrastructure design", not requirements.

Class (5th ed.)	Former name	Facility supply water, upper limit	Typical main cooling equipment	Typical supplemental equipment
W17	W1	17°C (62.6°F)	Chiller / cooling tower	Water-side economizer (cooling tower)
W27	W2	27°C (80.6°F)	Chiller / cooling tower	Water-side economizer (cooling tower)
W32	W3	32°C (89.6°F)	Cooling tower	Chiller or district heating system

Class (5th ed.)	Former name	Facility supply water, upper limit	Typical main cooling equipment	Typical supplemental equipment
W40	(new)	40°C (104.0°F)	Cooling tower	Chiller or district heating system
W45	W4	45°C (113°F)	Cooling tower or dry cooler	District heating system
W+	W5	>45°C (>113°F)	Cooling tower or dry cooler	District heating system

Minimum water temperature for all classes is 2°C (35.6°F). The Green Grid's glossary gives the same values ([The Green Grid](#)).

S classes for the technology cooling system

Because the W classes say nothing about the fluid at the cold plate, TC 9.9 has also presented Technology Cooling System classes. The 2025 lecture shows S30 (30°C / 86°F), S35 (35°C / 95°F), S40 (40°C / 104°F), S45 (45°C / 113°F) and S50 (50°C / 122°F), with a minimum of 2°C (3.6°F) above room dew point. It pairs S30/S35 with chiller/cooling tower plants, S40/S45 with cooling tower or dry cooler, and S50 with dry coolers ([Demetriou, 2025](#)). A 2026 national laboratory paper lists them as S20, S25, S30, S35, S40, S45 and S50, with the numbers in degrees Celsius, so S20 and S25 mean 20°C and 25°C TCS supply ([Becker et al., ACEEE 2026](#)).

ASHRAE's formal S-class text now sits in its subscription-only TC 9.9 Datacom Encyclopedia; check it before writing S classes into a contract.

What compliance means

Before 2021, some manufacturers labeled ITE "W4" if it could operate anywhere between 32°C and 45°C, so the FWS had to be designed to W3 to be sure of supporting it ([Upsite, 2024](#)).

The fifth edition fixed this with a performance requirement: "operating within a particular environmental class requires full performance of the equipment over the entire environmental range of the specified class, based on nonfailure conditions" ([ASHRAE TC 9.9, 2021](#)). Upsite paraphrases it as "full, unthrottled operation of the ITE at all temperatures within the respective W-class", so that "a facility designed to support W32 ITE can support any ITE which is W32 compliant."

Consequences for GPU buyers:

- **A class claim is only as good as its conditions.** Ask for flow, pressure drop and fluid at the class limit, not just the class.
- **Component labels are not rack labels.** Lenovo's GB300 NVL72 guide lists the NVLink switch tray operating water temperature as "2°C to 50°C (35.6°F to 122°F) (ASHRAE W45 compliant)" ([Lenovo, 2026](#)). That is a tray-level statement; the rack is limited by its hungriest loop.
- **Throttling is non-compliance.** A GPU that clocks down at the top of the class does not meet it.

Facility Water System vs Technology Cooling System

The CDU "separate[s] the facility water system (FWS) from the technology cooling system (TCS)" with liquid-to-liquid heat exchangers, circulates the TCS coolant, and allows a TCS fluid other than water ([ASHRAE TC 9.9, 2019](#)). The same white paper says the CDU must keep the TCS supply at least 3.6°F (2°C) above room dew point to prevent condensation.

TC 9.9's 2025 material is explicit: W classes "define Facility Water Supply Temperature (FWS) and not the Technology Cooling System (TCS)", "large CDU approach temperature ranges result in ITE inlet temperatures 3-15°+ C higher than FWS", and "ITE thermal requirement are based the inlet fluid temperature to the liquid loop: Not the FWS" ([Demetriou, 2025](#)).

So the stack looks like this:

```
outdoor air / wet bulb
+ heat rejection approach (dry cooler, tower, chiller lift)
= FWS supply temperature      <- W class applies here
+ CDU approach (3 to 15+ C per TC 9.9)
= TCS supply temperature      <- S class and GPU vendor tables apply here
+ cold plate and coolant temperature rise
= chip case temperature       <- set by the silicon vendor
```

Piping facility water straight to the rack removes the CDU approach, but ASHRAE notes such cold plates need larger fin spacing for less controlled water, which raises thermal resistance ([ASHRAE TC 9.9, 2019](#)).

What GPU vendors actually specify

NVIDIA's public GB200 NVL72 page describes the system as 36 Grace CPUs and 72 Blackwell GPUs "in a rack-scale, liquid-cooled design" but gives no coolant temperature, flow or fluid ([NVIDIA](#)). The DGX GB200 user guide gives rack power consumption of "approximately 120kW", liquid-cooled compute trays and air-cooled networking and storage, but no temperatures ([NVIDIA DGX GB200 User Guide, updated 2026-09-17](#)). As of 2026-10-28 we found no public NVIDIA document stating a W class or TCS temperature for GB200 or GB300 NVL72, so the numbers below come from OEMs. Your OEM's site documents for your SKU govern.

Lenovo GB300 NVL72 (OEM user guide, August 2026)

Lenovo's GB300 NVL72 user guide (second edition, August 2026) publishes this water requirements table ([Lenovo, 2026](#)):

Fluid supply temperature	Required rack flow	Rack pressure drop
25°C	59 LPM	2.3 psi
30°C	71 LPM	3.2 psi
35°C	89 LPM	4.9 psi
40°C	119 LPM	8.5 psi

Fluid supply temperature	Required rack flow	Rack pressure drop
45°C	177 LPM	18.4 psi

The same guide lists approved coolants as PG25, NALCO TCS101 or HPCCL2000, which indicates a closed secondary loop. Read the "fluid supply temperature" as rack inlet (TCS side), not facility water. The air side is specified separately as ASHRAE Class A2 (10°C to 35°C).

Going from 25°C to 45°C supply triples the required flow (59 to 177 LPM) and multiplies pressure drop by about eight (2.3 to 18.4 psi). That is what "45°C capable" costs at the CDU: bigger pumps, pipework or more CDUs.

QCT GB200 NVL72 (OEM paper, 2025)

A QCT paper on its QoolRack product states that for GB200 NVL72 "the maximum liquid inlet temperature is 45° Celsius, while the maximum liquid return temperature is 65° Celsius", and that flow "goes up to 130 liters per minute (LPM) per rack" (QCT, 2025). It is OEM marketing from a custom publisher, and covers GB200, not GB300.

What the Lenovo table implies (illustrative calculation)

Lenovo does not publish the liquid heat load next to the table. Assume 100 kW rejected to liquid (an assumption, not a GB300 spec) and approximate PG25 properties of density 1.02 kg/L and specific heat 3.93 kJ/kg·K (also assumptions).

Temperature rise $\Delta T = Q / (\dot{m} \times c_p)$, where $\dot{m} = \text{LPM} / 60 \times 1.02$.

- 25°C at 59 LPM: $\dot{m} = 1.003 \text{ kg/s}$; $\dot{m} \times c_p = 3.94 \text{ kW/K}$; $\Delta T = 100 / 3.94 = 25.4 \text{ K}$; return $\approx 50.4^\circ\text{C}$
- 35°C at 89 LPM: $\dot{m} = 1.513 \text{ kg/s}$; 5.95 kW/K ; $\Delta T = 16.8 \text{ K}$; return $\approx 51.8^\circ\text{C}$
- 45°C at 177 LPM: $\dot{m} = 3.009 \text{ kg/s}$; 11.83 kW/K ; $\Delta T = 8.5 \text{ K}$; return $\approx 53.5^\circ\text{C}$

Under these assumptions return temperature stays near 50 to 54°C whatever the supply (the 30°C and 40°C rows fall in between): the design trades supply temperature for flow, so warmer supply buys little extra return temperature for heat reuse. Do not reuse these values as specs.

Chillers vs dry coolers vs hybrid

The class you choose maps directly onto the plant:

- **W17 and W27** are typically chiller/cooling tower plants with a water-side economizer (Demetriou, 2025). Here the 2 to 3 percent chiller saving per 1°C applies.
- **W32 and W40** are cooling tower classes with a chiller as supplemental equipment. ASHRAE's 2021 paper warns that a facility built for economized operation with no mechanical cooling "should plan to add it in the future" (ASHRAE TC 9.9, 2021). Towers evaporate water; see our note on [data center water use](#).

- **W45 and W+** are where dry coolers become the main equipment and heat export (district heating) becomes realistic.

Schneider Electric told DCD that in a partner assessment, 95 percent of sites needed no mechanical assist for 300 W chips, but above 500 W only five percent could rely on free cooling ([DCD, 2024](#)). The study is not public; treat it as direction, not a number for your site.

Approach temperature and sizing the heat rejection

The question: on the design hour, what FWS temperature can the plant deliver without compressors, and is it below what the CDU needs?

Worked example (illustrative; every input below is an assumption)

Target: GB300 NVL72 racks, TCS supply at 35°C, using the Lenovo 89 LPM row.

1. **CDU approach.** Assume a CDU selected for a 3 K approach at design flow (TC 9.9 cites 3 to 15+ K across products). Required FWS supply = $35 - 3 = 32^{\circ}\text{C}$. The facility must guarantee 32°C or colder: a **W32** plant.
2. **Dry cooler only.** Assume a dry cooler approach of 8 K to outdoor dry-bulb. Maximum dry-bulb for compressor-free operation = $32 - 8 = 24^{\circ}\text{C}$. If the site's design dry-bulb is 35°C, the dry cooler alone delivers $35 + 8 = 43^{\circ}\text{C}$, which is 11 K too warm. You need adiabatic assist, a tower, or trim chillers sized for the hours above 24°C.
3. **Cooling tower.** Assume a 4 K tower approach to wet-bulb plus a 2 K plate heat exchanger approach. Maximum wet-bulb = $32 - 4 - 2 = 26^{\circ}\text{C}$. Check this against your site's design wet-bulb and the tower's water consumption.
4. **Go colder.** If the vendor requires 25°C TCS (59 LPM), the FWS must be 22°C with the same 3 K CDU approach. No class boundary sits at 22°C, so that equipment is effectively W17 and the facility must deliver 17°C or colder, which leaves up to 8 K for the CDU approach but puts chillers back in the plant.
5. **Go warmer.** At 45°C TCS (177 LPM, 18.4 psi), FWS can be 42°C with a dry cooler only up to 34°C dry-bulb ($42 - 8$). Compressors may disappear, but pumping and piping for three times the flow appear.

Each kelvin of CDU approach comes straight off your economizer limit.

Choosing W17 vs W27 vs W32+ in practice

- **Pick the class from the hardware, then add margin.** Take the coldest TCS temperature any GPU generation you expect to host will need, subtract the CDU approach, and round down to a class boundary. ASHRAE's 2021 forecast that supported water temperatures will "regress" argues against designing to the warmest class today's hardware tolerates ([ASHRAE TC 9.9, 2021](#)).
- **Colocation and multi-tenant halls** tend toward colder water (W17 to W27) because the operator does not control what tenants install; DCD heard anecdotally of W27 expectations ending at W17 ([DCD, 2024](#)).

- **Single-tenant AI factories** with one GPU platform can design to W32 or W40 if the vendor tables support it at acceptable flow, and keep trim chillers for hot hours.
- **Heat reuse** is the main reason to go to W40, W45 or W+. Per the arithmetic above, check the return temperature you would actually export before building a business case.

Liquid cooled GPU racks still reject some heat to air; coordinate the W class with the air class and our [rack power planning](#) guide.

Commissioning and monitoring implications

- **Set points:** keep FWS supply below the class limit by the control tolerance; set the TCS set point from the vendor table, not the W class.
- **Dew point:** CDU control must keep TCS supply at least 2°C above room dew point ([ASHRAE TC 9.9, 2019](#)).
- **Ride-through:** TC 9.9 notes that with liquid cooling "a loss of flow will result in ITE overheating within seconds", and has published a tech brief on redundant design of the TCS and FWS loops ([Demetriou, 2025](#)). Alarm on flow and differential pressure, not only temperature.
- **Excursions:** log FWS and TCS supply temperatures at high resolution; a heat-wave or chiller-staging excursion above the class limit is outside the ITE's specified envelope.
- **Ratings:** liquid-to-liquid CDUs now have a method of test: Addendum b to ASHRAE Standard 127-2020, approved 2026-02-27, rates them at 26°C facility water and 30°C TCS supply ([ASHRAE, 2026](#)). Most other liquid cooling equipment still has none; verify CDU approach at your flow during commissioning.

See our [cold plate commissioning checklist](#) for flushing, leak testing and acceptance.

Specification checklist for RFPs

Ask the GPU OEM:

1. Rack inlet (TCS) temperature, flow and pressure drop table for your SKU, firmware and fluid.
2. The supply temperature and heat load at which performance is guaranteed without throttling.
3. Maximum return temperature, and the heat split between liquid and air.
4. Whether any W or S class claim refers to the rack, a tray or a component.

Ask the CDU vendor:

5. Approach temperature (TCS supply minus FWS supply) at your design flow and heat load, and test method.
6. Pump redundancy, maximum flow and available external head at the 45°C flow condition if you intend to use it.
7. Dew point control method and filtration on both sides.

Ask the MEP designer:

8. FWS class and design-day supply temperature, with the climate data used (dry-bulb and wet-bulb percentiles).
9. Hours per year of compressor-free operation at the chosen class.
10. Behavior on chiller or fan failure: how long until FWS exceeds the class limit.

GPU Smith helps teams turn vendor cooling tables into facility water specifications ([contact](#)).

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. https://www.ashrae.org/File%20Library/Technical%20Resources/Bookstore/Emergence-and-Expansion-of-Liquid-Cooling-in-Mainstream-Data-Centers_WP.pdf
2. <https://www.datacenterdynamics.com/en/analysis/hot-water-cold-water>
3. <https://www.aceee.org/wp-content/uploads/2026/08/data-center-high-temperature-liquid-cooling-and-heat.pdf>
4. https://dallas-ashrae.org/images/meeting/041625/the_ashrae_thermal_guidelines_for_data_centers____past__present_and_future.pdf
5. <https://www.upsite.com/blog/major-changes-to-ashraes-fifth-edition-of-thermal-guidelines-part-3-liquid-cooling-chapter-updates>
6. <https://www.thegreengrid.org/resources/glossary/W>
7. https://pubs.lenovo.com/gb300-nvl72/gb300-nvl72_user_guide.pdf
8. https://www.ashrae.org/File%20Library/Technical%20Resources/Bookstore/WhitePaper_TC099-WaterCooledServers.pdf
9. <https://www.nvidia.com/en-us/data-center/gb200-nvl72>
10. <https://docs.nvidia.com/dgx/dgxb200-user-guide/hardware.html>
11. https://blog.qct.io/wp-content/uploads/2025/04/QCT-Qoolrack-Stand-Alone_Advanced-Liquid-Cooling-for-NVIDIA-GB200-NVL72-Systems.pdf
12. https://www.ashrae.org/file%20library/technical%20resources/standards%20and%20guidelines/standards%20addenda/127_2020_b_20260227.pdf

THE PUBLISHER

About GPU Smith

GPU Smith provides independent engineering and research for private AI infrastructure. We help technical buyers and operators reason about GPU workload sizing, hardware selection, procurement, deployment and inference operations, from the compute system to the networking, power and cooling requirements around it.

Start with the workload

Useful infrastructure decisions begin with the models, data, throughput, latency and operating constraints a team actually has. GPU Smith helps connect those requirements with system design and deployment choices. The goal is a privately controlled AI environment whose capacity and operational demands are understood before a purchasing decision.

Hardware and supplier research

Our public reference library covers GPUs, complete systems and networking components. The server-vendor directory and manufacturing research help teams investigate suppliers, compare options and follow sources behind technical claims. We distinguish manufacturer specifications from measured benchmarks and advertised capabilities from tested configurations. Reference pages are research resources, not a live inventory listing or a binding equipment quote.

Deployment and operations

GPU Smith's engineering scope includes infrastructure integration, networking, power, cooling and the practical operation of inference workloads. Our published articles explain the tradeoffs behind deployment, procurement and ongoing operations so teams can ask better questions and document decisions.

Work with GPU Smith

Explore the [hardware library](#), [GPU references](#), [networking references](#) and [vendor research](#). [Contact GPU Smith](#) with your workloads and deployment constraints to discuss a project and confirm current scope, pricing and availability.

Inclusion of a manufacturer or product in our research does not imply a partnership, certification or endorsement.

Website: <https://gpusmith.com>

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.