

Warm-Water Ready Mission Critical Cooling




Apples-to-apples testing between **single-phase direct-to-chip (1P D2C)** and **two-phase direct-to-chip (2P D2C)** on an 8-way NVIDIA B200 server revealed these key thermal differences:

Compared to single-phase, two-phase D2C runs:

 **9-14°C cooler at the full-system level**

50°C Facility Water Temperature



With 50°C FWS, 1P D2C forces GPU downtime—while 2P D2C has headroom to spare.

The Nine-Degree Difference

Mechanical Cooling Needed (Hours Per Year)

	At 45°C FWS	At 54°C FWS
Northern Virginia	5 hours/year	0 hours/year
Austin-San Antonio	73 hours/year	0 hours/year
Dallas-Fort Worth	129 hours/year	0 hours/year
Phoenix	665 hours/year	0 hours/year


NVIDIA says its liquid cooling systems can run with 45°C facility water (FW). **Ours can run at 54°C, or up to 59°C at higher flow rates.**

For most of the world, that difference **unlocks year-round free cooling** and **eliminates the need to build a mechanical plant** for assisted cooling—saving you up to **\$1M per 1MW** in CapEx alone.

2P D2C Is The Enabling Technology For AI Factories



Save 36-56%
In Annual
Cooling Energy



Deploy 5% More
GPUs In Same
Power Envelope



Be Ready For
Rising TDPs &
Rack Densities



Click Logo
To Download
The Full Results

