How to Pick the Right Cooling Solution for HPC Server Systems? A Tech Guide by GIGABYTE

25/05/2022 by **GIGABYTE for ISC 2022**

The exponential increase in performance of key components in HPC servers, such as CPU, GPU, network adapters and so on, is a trend that will persist and rapidly evolve in upcoming years, if not months. In this tech guide based on GIGABYTE's decades of experience in HPC server design, we explain three popular options for system cooling – air, liquid, immersion – and demonstrate what can be done to reduce the energy footprint and costs while gaining extra computing power.

1. Air Cooling

Airflow in the data center must follow a specific path designed to disperse heat efficiently both in servers and throughout the data center. Sub-optimal airflow will negatively impact system performance from the individual components level up to clustered systems, as hot air builds up and creates a snowball effect in each step of heat transfer and cooling. Furthermore, measures such as prevention of airborne contaminants into system fans, optimization of the physical dimensions of chassis vents and design, and configurable system cooling (either via hardware design or software control) help to ensure airflow efficiency.

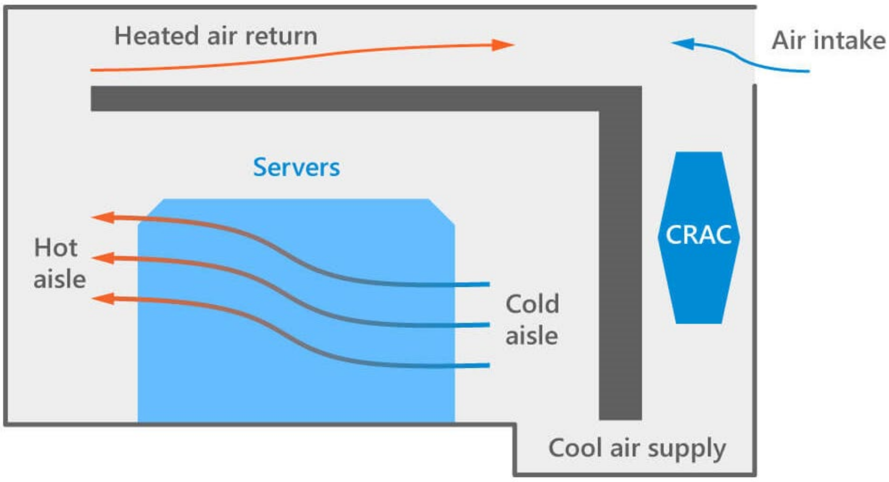

Traditional hot/cold aisle cooling: a cross section view demonstrating how heat is dissipated in an air-cooled data center. Servers must also be designed to facilitate airflow. They can be customized to provide greater protection against airborne contaminants that may cause equipment failure. (CRAC: Computer Room Air Conditioning)

An Airflow-Friendly Hardware Design

Powerful fans and high-performance heatsinks are installed to maximize contact between the cool air and the surface area of key components. Streamlined fins are a part of the heatsink design to further improve cooling and airflow. Furthermore, the components are made from heat-resistant materials to prevent heat transfer to the rest of the system, to ensure heat reduction when dealing with heavy workloads.

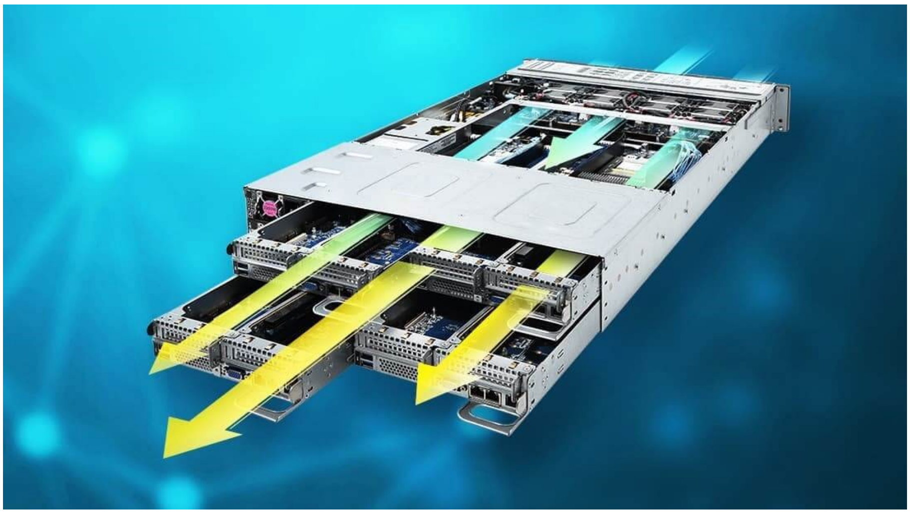
Example of heat transfer: the **H263 Series**, from the GIGABYTE high-density server series with 8 x CPUs (up to **400W TDP each)**, benefits from GIGABYTE's know-how in system cooling design. The overall airflow direction inside the system chassis is evaluated with simulation software and then fine-tuned to optimize ventilation.

Programming: Automatic Fan Speed Control

Besides hardware design, the programming of automatic fan speed control is equally sophisticated. Thermal sensors are placed near key components to monitor their temperatures. If the baseboard management controller (BMC) detects a change in the temperature reported by the sensors, the speed of the corresponding fans near the key components in question will self-adjust automatically. More granular control of fan speeds can be achieved for different sections / key components of the server system by using fan speed profiles which can be created or edited via IPMI or a proprietary platform management software such as GIGABYTE Server Management (GSM) to accommodate specific requirements and balance temperature control and power efficiency.

Adaptive Customization to Protect Against Contaminants

By default, an air-cooled data center is susceptible to damage from elements such as moisture and dust. Modern data centers should have air purifiers and adhere to a strict guideline for operating environmental controls, but the risk of airborne contaminants is ever-present. Studies must be done by the hardware vendor on the operating environment in data centers to allow for a customized mechanical design (such as fan filters, dust covers over chassis I/O's, etc.) and to achieve an ingress protection rating (for dust resistance).

2. Direct Liquid Cooling

Direct liquid cooling (DLC or direct-to-chip liquid cooling, D2C) uses cold plates and sealed tubes to create a cooling loop which is filled with coolant; these elements can be placed directly in contact with key components such as CPU, GPU, RAM, and network cards to remove the generated heat, with little hardware modification. A heat exchanger is then used for removing heat from the coolant, allowing the coolant to circulate back into the servers and repeat the cycle.

Example of DLC: a liquid-cooled server node from **a 2U 4-node server,** the **H263/H273 Series** of high-density servers: cold plates and closed loops filled with coolant are placed on the CPUs to dissipate heat (up to 400W TDP per CPU, 8 x CPUs per server).

A liquid-cooled server is not only more effective in dissipating heat, but also removes unnecessary air-cooling parts like heatsinks and fans, which leads to a reduction in costs and energy consumption of active mechanical parts. For this reason, high-density and HPC servers often utilize some form of liquid cooling, whether via a hybrid "liquid-to-air" system or a pure "liquid-to-liquid" heat exchanger.

Liquid-to-Air Cooling

This hybrid approach is ideal for air-cooled data centers without having to completely overhaul the existing infrastructure. Liquid-cooled servers are installed in standard server racks, and closed cooling loops are connected to heat exchangers fitted in the same rack. Heat is evacuated into the hot aisle, similar to how an air-cooled server expels heat, which makes it possible to have a mix of air-cooled and liquid-cooled servers in the same facility. Though heat exchangers take up some rack space, liquid-cooled servers can house more dense high-power components like top-tier CPUs and GPUs; thus, they provide an overall increase in computing performance that offsets the extra space required.

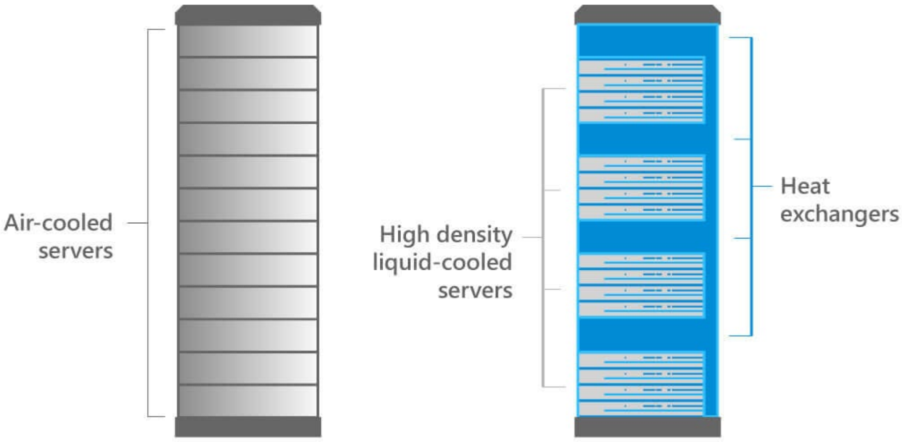

Comparison of air vs liquid cooled: a liquid-to-air cooling solution – heat exchangers are fitted in the same rack as the liquid-cooled servers, transferring heat from the coolant into the "hot aisle". This approach allows processors to be packed closely together, resulting in greater computing power.

Liquid-to-Liquid Cooling

The liquid-to-liquid method can rapidly remove a greater amount of heat compared to the liquid-to-air method, but it comes with a caveat: the data center must be equipped with a liquid cooling infrastructure. Rather than evacuating heat into the hot aisle, the rack-mounted heat exchanger transfers heat to the facility's liquid-based cooling loop, which is more efficient in dispersing heat than a standard CRAC. Therefore, server density can dramatically increase and result in even greater computing power. The liquid-to-liquid solution comes in different sizes: from localized liquid-cooling in just one server rack to an entire liquid-cooled data center.

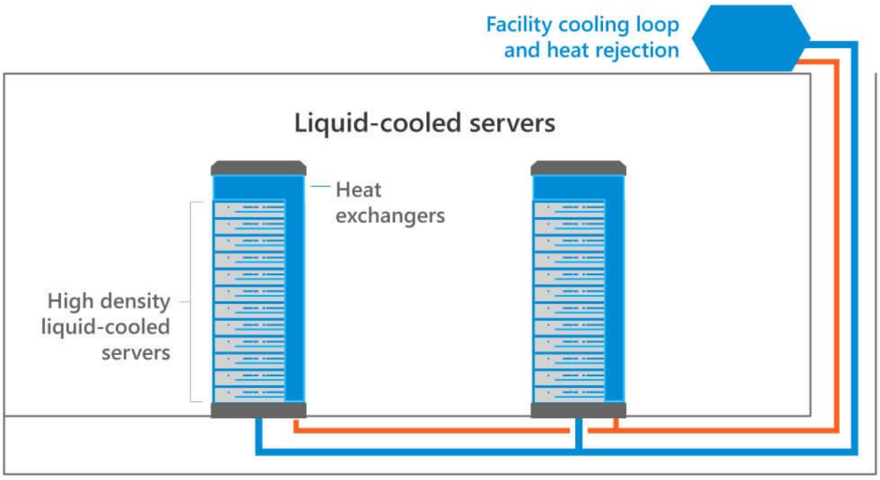

Liquid-to-liquid cooling system: transfers heat from the coolant to the data center's facility water, delivering even higher server density and computing power than can be achieved with the liquid-to-air approach.

All industry sectors can benefit from liquid cooling, for a wide range of scenarios. For instance, edge data centers can use liquid-cooled high-density servers to squeeze more performance into a limited space. In the financial sector where server systems are used to run programs for high frequency trading (HFT), fast and stable liquid-cooled servers are in high demand.

A quick case in point is the green computing by the German Aerospace Center (Deutsches Zentrum für Luft- und Raumfahrt, also known as DLR) for aerospace, energy, and transportation research. The DLR team built a green data center to support space exploration as well as other research programs by running HPC server systems that operate at an ambient temperature of 40°C with no air conditioning equipment. The team's objective was achieved by using liquid-cooled high-density servers from GIGABYTE (H261-Z60: for full success story). Paired with optimized power supplies, the liquid cooling HPC solution was able to reduce power consumption by 15% compared to the competition. And as more CPUs could be installed, the energy footprint was reduced by 50% while providing up to twice the maximum computing power of HPC solutions in contrast with the competition.

3. Immersion Cooling

Immersion cooling is the pinnacle of liquid cooling. By submerging servers directly into a bath of non-conductive fluids inside a cooling tank, heat can be removed without using parts such as heatsinks or cold plates and no air conditioning system is needed. The thermal energy is transferred by the fluids in one of two ways: through a heat exchanger, as is the case with single-phase immersion cooling; or through a cycle of vaporization and condensation, as is the case with two-phase immersion cooling.

Single-Phase Immersion Cooling

Single-phase immersion cooling can be deployed based on heat dissipation needs. Servers are immersed in a thermally conductive dielectric coolant. The coolant's chemical properties, just like in two-phase immersion cooling, must be cautiously studied in regard to the target components that need to be cooled, to ensure the best performance and durability of the total solution. In single-phase immersion cooling, the coolant does not change its physical state and will always remain a liquid; no boiling or vaporization occurs throughout the operation (and therefore ideal for safe materials handling and compliance with environment regulations). The coolant in the cooling tank is circulated by a cooling distribution unit (CDU) that transfers the coolant liquid to a heat exchanger outside the tank where the heat is transferred to a secondary cooling circuit running with facility cold water.

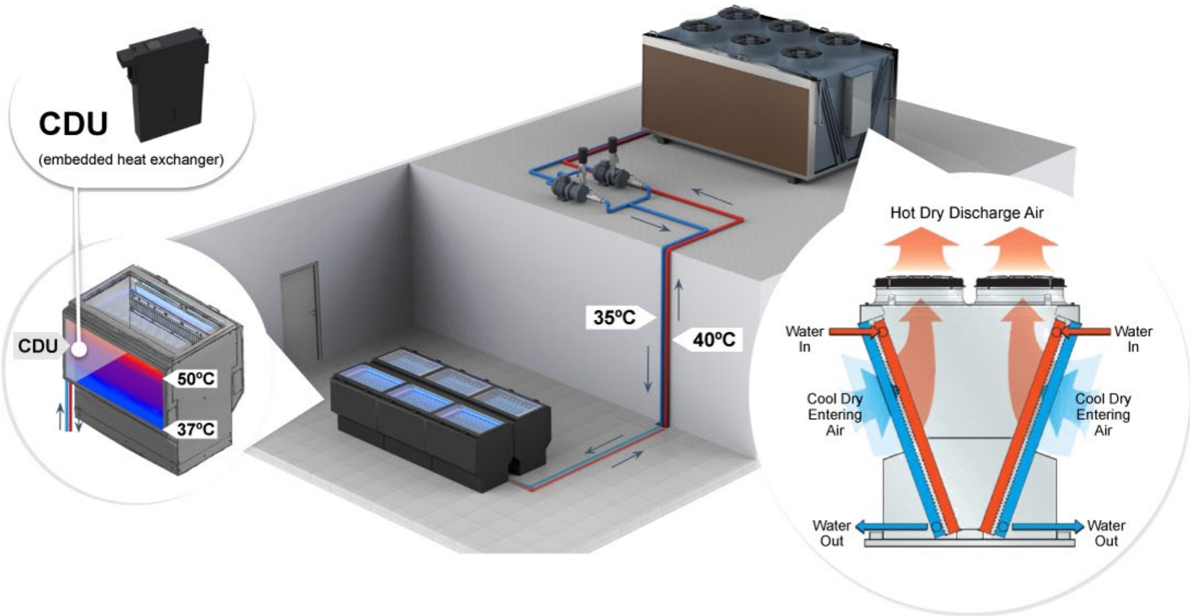

Single-phase system: the cooling cycle of a single-phase immersion total solution.

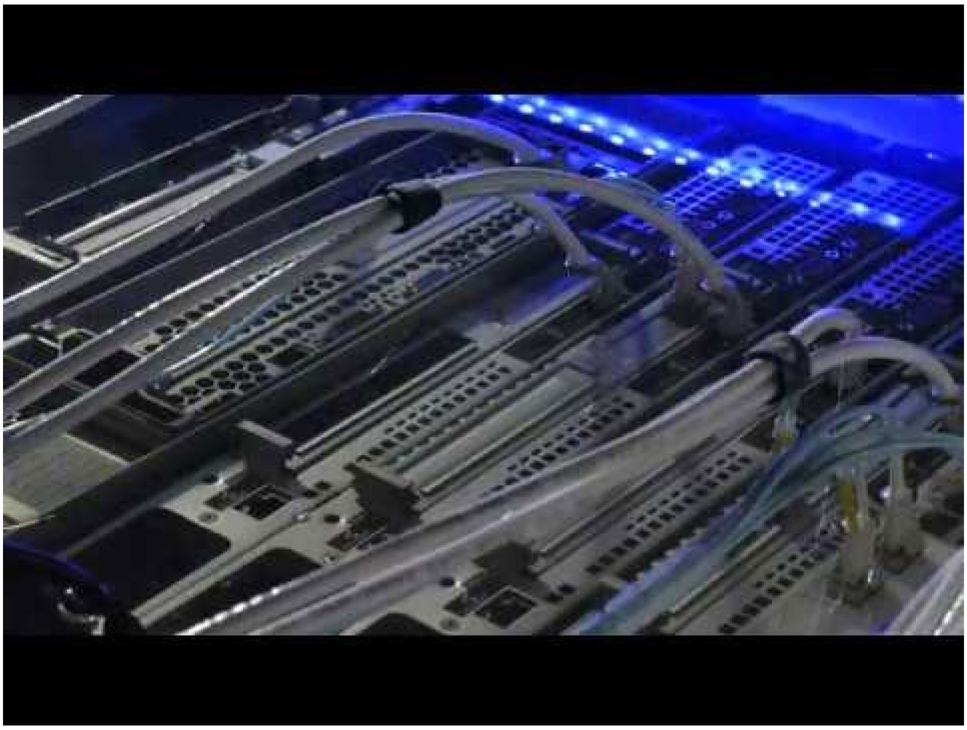

Single-phase demonstration: a single-phase immersion cooling solution from Submer, a GIGABYTE partner for immersion cooling total solutions, compatible with the **H262 Series** (2U 4-Node, 8 x CPU), the **G292 Series** (2U, 8 x GPU, the densest 2U GPU computing solution in the market), the **G152 Series** (1U, 6 x GPU, the densest 1U GPU computing solution in the market), and the **E152/E162 Series** (1U, 1 x GPU, 40-55cm in depth, for AIoT/Edge Computing and 5G NR use cases).

Two-Phase Immersion Cooling

Two-phase immersion cooling offers many advantages. Firstly, circulation of fluids inside the cooling tank occurs without using an additional mechanism to circulate; it works rather by the natural process of evaporation, and so power usage efficiency (PUE) improves. Servers can be placed inside a sealed tank filled with an open bath of transparent, dielectric fluids that are non-conductive and non-corrosive. Because the fluids have an ultra-low boiling point of 56°C, heat from key components naturally causes the fluids to undergo a low-temperature vaporization process, thus transferring heat out of the fluids. Vapor then rises to the top of the tank where it is cooled by a heat exchanger (in this case, a condenser coil) to transform back to a liquid state. Then the liquid drops back into the bath to be reused.

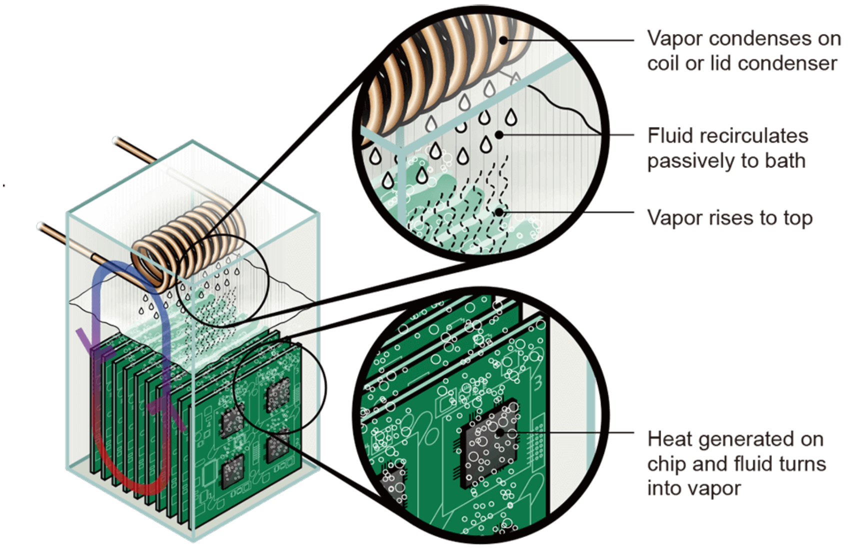

The Passive 2-Phase Immersion Cooling Cycle

Two-phase cooling: evaporation, condensation- repeat.

By estimation, replacing air cooling with two-phase immersion cooling can lead to energy savings of up to 90% and a reduction in cooling OPEX by up to 95% in data centers when compared to air cooling – including server power consumption and air conditioning costs. The energy and economic costs are even more evident in HPC projects using high-power/high-TDP computing parts such as GPUs. Furthermore, the removal of parts like pumps and jets means less maintenance and a lower chance of equipment failure, and system stability is improved while operating costs go down. As a result of space optimization, key computing components such as CPU and GPU can be packed closer together. The greater density of high-power components packed within a smaller space can increase computing power by up to 10 times, which becomes appealing, for example, to telecom companies who need to establish edge data centers for high-speed processing capabilities and for reducing data transfer latency. Such edge data centers are often located in crowded urban areas with exorbitant property prices, so greater server density is essential.

Two-phase demonstration: an improved PUE solution with two-phase liquid immersion cooling by GIGABYTE, running the **G292 Series** server (2U, 8 x GPU, the densest GPU server solution in the market).

Summary: Get Expert Advice and Proof-of-Concept

Roughly speaking, to use a standard server rack as a basis of comparison, air cooling offers up to 20Kw – 30kW of heat dissipation; liquid cooling offers 80kW of heat dissipation, and immersion cooling offers 200kW of heat dissipation (with the cooling water temperature below 13°C). The appropriate total solution design varies as every project has specific requirements which require expert consultation. AT GIGABYTE, we help data centers and customers analyze their projects by looking into energy consumption, heat dissipation, space optimization and PUE/ water use efficiency (WUE), among many other technical topics, at each step of solution design. Taking a step further, GIGABYE also offers installation/deployment services working with data center infrastructure companies, to ensure that customers receive smooth project delivery and short turn-around time for operational readiness. Most importantly, GIGABYTE strongly advises its customers to benefit from its Proof-of-Concept (PoC) resources for validating every solution design and project parameters for making the best decision, as many environmental factors might alter the expected performance and system stability. GIGABYTE has PoC units (in both single-phase and two-phase immersion cooling) for testing and validation of immersion cooled servers. The server model options come in 1U/2U/4U form factors and can be modified on request to suit different use cases and workloads. GIGABYTE works with all major liquid-cooling and immersion-cooling technology partners in the market, so that customers can count on design-in compatibility of GIGABYTE total solutions with their infrastructure.

GIGABYTE PoC: GIGABYTE's two-phase immersion cooling PoC unit, compatible with 2U server systems such as the **G292 Series** (2U, 8 x GPU per server, the densest GPU computing system in the market) and can be modified on request to be used with 1U or 2U servers (such as the **G152 Series** (1U 6 x GPU)**,** and the **H262 Series**, (2U 4-Node, 8 x CPU).

Ultimately, whether air cooling, liquid cooling, or immersion cooling, GIGABYTE has the solutions developed with its world-class engineers from all domains. We encourage you to reach out to our sales representatives at server.grp@GIGABYTE.com for consultation.