Everpure™ Storage in Liquid-Cooled Data Centers

Data storage in a liquid-cooled world

Heat is an inevitable byproduct of electronics at work. Increasingly power-hungry servers, switches, and storage systems are making heat dissipation a major factor in data center design and operation.

Moreover, rapid adoption of artificial intelligence (AI) systems with thousands of densely packaged *graphics processing units* (GPUs) is increasing heat generation manyfold. Conventional air cooling is becoming inadequate for these environments, so data centers are turning to liquid cooling for major components and systems.

Processors, especially GPUs, generate more heat than other data center equipment, making them the main drivers of the need for liquid cooling. Liquid cooling is available in multiple forms, but all require major changes in rack, and in some cases circuit board and chip design, as well as in data center layout and provisioning.

For storage systems, air cooling appears to be adequate for the foreseeable future, although at some point it is likely to become necessary. Today's challenge is incorporating air-cooled storage in *hybrid data centers* where other components use liquid cooling.

This brief summarizes the power and cooling needs of Everpure storage systems today and in the future and describes how they can coexist in data centers with liquid-cooled equipment.

Everpure

TB-251203-v01

THE HEAT IS ON

The more computing power a chip has, the more electrical power it draws. The top-of-the-line Intel Xeon 6 processor commonly found in servers at publication time consumes over 200 watts and generates nearly 700 BTU of heat. But that pales in comparison to the 1,200 watts needed to power NVIDIA's B200 (Blackwell) GPU. A full rack of servers containing these GPUs (e.g., Nvidia's NVL72) consumes 120-132KW (400,000-450,000 BTU), with newer products making even greater demands.

Artificial intelligence applications utilize hundreds of GPUs operating in parallel. When the server chassis that host GPUs are mounted in rows of 42U, 48U, or 52U equipment racks (73.5, 84, or 91 vertical inches of mounting space respectively), conventional air cooling is often inadequate. Liquid cooling makes higher thermal density possible, but it's not inexpensive or easy to implement.

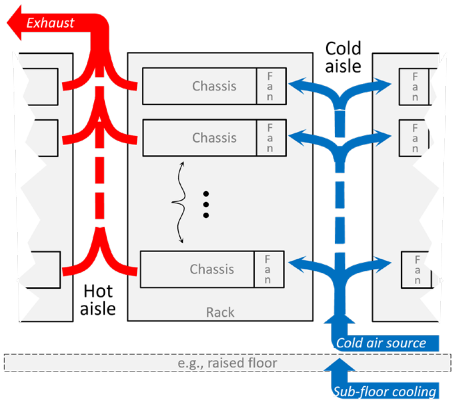

Figure 1: Typical Air Cooled Layout

In typical air-cooled data center layouts (Figure 1), server, switch, and storage chassis mounted in rows of racks each have their own fans that draw chilled air across their heat-producing components, usually from a "cold aisle" between two rows of racks into which chilled air is pumped. As air moves through equipment chassis, it absorbs the heat generated by their components before exiting, usually into a "hot aisle," where it is drawn out of the data center or channeled into heat exchangers that chill and recirculate it via the cold aisle. Air cooled data center designs attempt to manage large volumes of air flowing throughout the data center. Most do not focus airflow precisely on heat-generating components.

Liquid Cooling

Liquid cooling also circulates a fluid to move heat away from the chips that generate it. There are different versions of the technology, including rack doors with built-in heat exchangers, direct-to-chip plumbing, and immersion systems in which equipment is submerged in a non-conductive fluid.[1] Today, direct-to-chip liquid cooling appears to be the most commonly encountered technology.

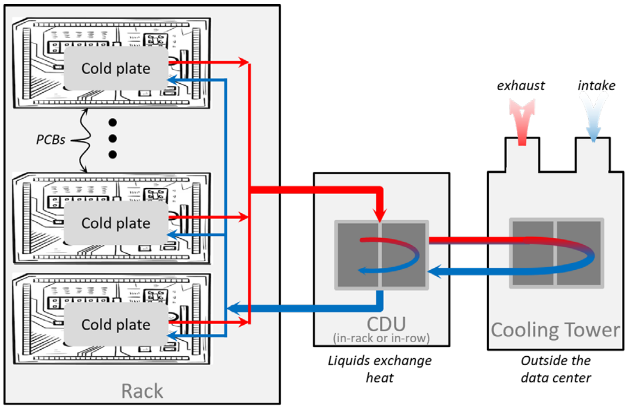

Figure 2: Direct-to-Chip Liquid Cooling

[1] https://www.vertiv.com/en-us/solutions/learn-about/liquid-cooling-options-for-data-centers/#/tab-heat-rejection

In direct-to-chip infrastructures, *cooling distribution units* (CDUs), located either in each rack or at the ends of aisles, pump chilled liquid to and from *cold plates* affixed to chips via elaborate networks of piping. Cold plates fulfill essentially the same function as the on-chip heat sinks found in air-cooled environments, but they do it more efficiently.

Liquid flowing through cold plates absorbs heat generated by the chips and returns to the CDUs, where it is cooled and recirculated by a separate thermal circuit connected to cooling towers outside the data center. Because liquid transfers heat far more efficiently than air, direct-to-chip cooling uses relatively small volumes of coolant which it moves more precisely to and from sources of heat.

While the fundamental job of liquid cooling is to dissipate heat generated by high-power components, it also reduces data center electricity consumption, noise and space.[2] It removes most of the heat generated by densely packed components, but it is usually augmented by air-cooled infrastructures that function like that shown in Figure 1.[3] Thus, data centers that use direct-on-chip cooling are in fact hybrids that use both forms of thermal management.

Rack Design Considerations

Air cooled server, network, and storage chassis place few constraints on rack design. Generic industrial quality racks usually cost less than $10,000 (including power distribution units).

Racks for direct-on-chip liquid cooling can be more than five times as expensive.[4] In addition to greater power distribution (100-150KW vs. about 40KW for typical air-cooled racks), they include plumbing for carrying coolant to and away from heat sources and CDUs. External heat exchangers and leak prevention and recovery measures add to the overall cost.

Data center liquid cooling is complex and expensive to implement, but as power requirements continue to grow, it is becoming inevitable, both at hyperscale and in enterprise data clouds. Direct-on-chip liquid cooling and data center design and retrofit to use it are rapidly evolving technologies. Hyperscale operators have both the most pressing need for it and the resources to design custom infrastructures for their own use.[5] Equipment vendors, notably Nvidia and Supermicro, are developing easier-to-deploy standardized infrastructure products[6] that make liquid-cooled data centers accessible to large enterprises and AI service providers.

[2] For example, https://www.supermicro.com/white_paper/white_paper_LC.PowerSavings.pdf

[3] https://www.airedale.com/2025/05/13/how-hybrid-cooling-is-future-proofing-data-centers-for-the-next-generation-of-ai

[4] For example, https://www.reddit.com/r/RigBuild/comments/1p5sig3/cooling_system_for_a_single_nvidia_blackwell

[5] For example, https://www.aboutamazon.com/news/aws/aws-liquid-cooling-data-centers

[6] For example, https://www.youtube.com/watch?v=dr9n2cjx0Bs

STORAGE IN LIQUID-COOLED DATA CENTERS

Today, storage systems have considerably lower power density than high-powered processor complexes. For example, at publication time the largest FlashArray model, XL170R5 fully populated with DirectFlash™ modules peaks at 3.16KW. A FlashBlade//S500 chassis with 10 blades and 40 DFMs tops out at 2.75KW. A full rack of these systems is well within the typical 15-20KW per rack power budget common in air-cooled data centers.

But FlashBlade systems frequently provide storage for artificial intelligence applications, as do FlashArrays for performance-sensitive databases. As these applications increasingly implement liquid cooling to meet processing requirements, their data centers will necessarily be hybrids, supporting both liquid and air-cooled equipment.

While direct-on-chip cooled racks are actually hybrids of liquid and air cooling, it is good data center design practice to locate liquid and air-cooled equipment in separate rooms or on separate levels because the coolant flow requirements of the two differ. For example, in air-cooled data centers, raised flooring is often a sub-floor conduit for chilled air, whereas racks for direct-on-chip cooled equipment are usually located directly on concrete floors. Moreover, the racks themselves are very different—a $50,000 liquid-cooled rack would be an expensive mounting option for air-cooled equipment, even if one ignores the need for adaptation.

LIQUID-COOLED STORAGE?

Designers of artificial intelligence servers are exploring the feasibility of adding local storage to their server chassis to front-end the huge data sets typically hosted by FlashBlade and similar systems. This usage would require multi-terabyte devices capable of tens of thousands of IOPS. Beefing up processing, buffer, and external communication resources would enable more concurrent I/O, but it would increase power draw and heat generation beyond what is common in today's DFMs and SSDs. Moreover, for mounting in liquid-cooled server chassis, storage devices would likely have to be liquid-cooled as well, even if they did not require it per se.

As flash density increases, the trend in Everpure's own systems is toward configurations with fewer, higher-capacity DFMs. To maintain and improve system performance, future DFMs will have to be capable of more concurrent I/O, with corresponding increases in power consumption and heat generation. More powerful DFMs and higher-speed host/client communications will need more powerful blades and controllers. It is therefore likely that within two or three product generations, the blades in FlashBlade systems, FlashArray controller mainboards, and the DFMs used by both product families will require liquid cooling. The company's hardware engineering teams are already analyzing the ramifications of storage systems with much higher thermal density and are exploring the various possible paths for transitioning from air-cooled to liquid-cooled storage devices and systems.