When and How Liquid Cooling Becomes Necessary in the Data Center

With Moore's Law having dead-ended, the ever-increasing power of processors has gone back to generating heat. Liquid cooling in its many forms is becoming a serious option once again. But is it really practical to mix flexible pipes and cooling plates with air ducts and blanking panels in a data center that's already very complex?

BY SCOTT M. FULTON, III

This much we know: Liquid is a more effective medium for transferring heat than air. This one isn't among the many tenets of science denied in recent years. As the components of data center hardware have resumed trending hotter, architects and engineers' attentions have once again shifted to the matter of leveraging the untapped thermal transfer power of liquid.

"Liquid cooling, it's going to happen," Joe Capes, global business development director for Schneider Electric, said. "It is already happening in HPC (high-performance computing). Ultimately, over the next five to ten years, you're going to see more and more companies with the ambition of doing away with air cooling in their data centers altogether, replacing that with liquid cooling — and remove 100 percent of the heat."

His theory goes something like this: The nature of applications running in enterprise data centers is changing, increasingly resembling high-performance computing, even supercomputing. Not only is there more analytics taking place, the rapid rise of orchestrators such as Kubernetes has illustrated the need for distribution and balance of processing power throughout a cluster of servers.

Putting Power in TDP

In data centers, as much as anywhere else, work expends energy, and energy is dissipated as heat. CPUs and other processors are designed to dissipate heat – now to the extent where the heat they generate is the premier measure of their performance. Today, thermal design point (TDP) is used synonymously with power so often, that engineers have revised the "P" to stand for "power." For example, when Intel was setting design goals for its mid-range Xeon E5 processors a few years ago, it capped TDP at around 150W.

Now, when chip makers appear to have reached the physical limits of Moore's Law, Intel has blown by that threshold, and it's looking back. The company's top-of-the-line Skylake-series Xeon chips are rated for 205W. Nvidia, as part of a marketing move last year aimed at pushing AI for the enterprise, rated its newest Volta accelerator GPUs for 300W TDP, characterizing that new power level as "unconstrained." More recently, Nvidia boosted TDP rating for its existing Tesla 32GB model to 350W.

In addition to translating to processing power, wattage also translates to heat. And if the processor makers can't make that heat go somewhere, someone else has to.

"The high level of interest and activity in the market today is not so much driven by wringing out a few more points of efficiency," Capes said. "It's driven by the fact that Intel, AMD, and Nvidia are coming out with much higher powered chips as part of their roadmaps. So, [liquid cooling in some form] is not really a choice anymore."

A server liquid-cooled by a CoolIT Systems solution

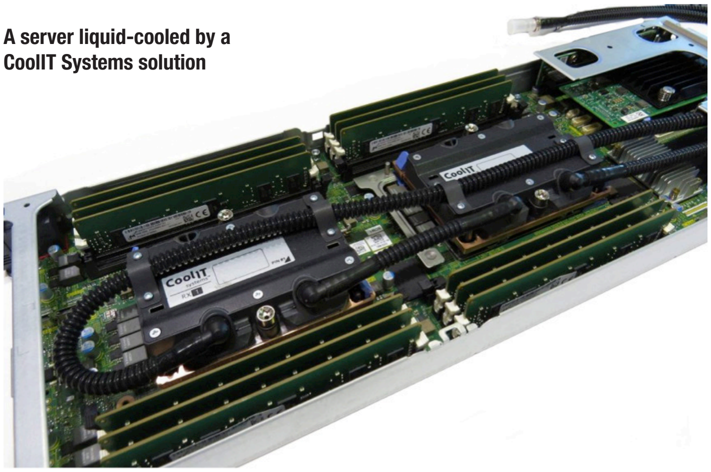

Photo Credit: CoolIt Systems

Very high-density solutions are coming to the market now that really call for liquid cooling, Lars Strong, senior engineer and chief science officer at Upsite Technologies, said. "Probably liquid is the only way to cool some of them. To not be considering that would be being blind."

Unless you count open air as technology, water cooling in the data center is actually older than air circulation and evaporative cooling – by a few decades. IBM's System 360 mainframe, introduced in 1964, and all the derivatives of that model up unto the late 1990s, were water-cooled. The whole industry of data center airflow management came about on account of modern servers being only slightly evolved versions of desktop PCs, with their exhaust fans in the rear and little or no air intake in the front.

It's not surprising that among early leaders in the re-emerging liquid cooling are some of the very same companies that led the gaming-PC liquid cooling market in 2000s, including Stamford, Connecticut-based CoolIT Systems.

"Can people air-cool their current-to-next-generation enterprise-level data centers with air?" Patrick McGinn, CoolIT's VP of product marketing, asked. "They can. Is it the best way to do it? No."

CoolIT produces a kind of active heatsink for specific server processors, just as it did for gaming-desktop processors over a decade ago. But in the company's grandest vision, its optimum liquid cooling approach involves a stand-alone module that connects to the data center's existing air exchange system. Either unit is connected to a distribution manifold attached to the rear of the cabinet, using "dry-break quick disconnects" with built-in shut-off valves to direct coolant through the system using pipes reminiscent of coaxial cable.

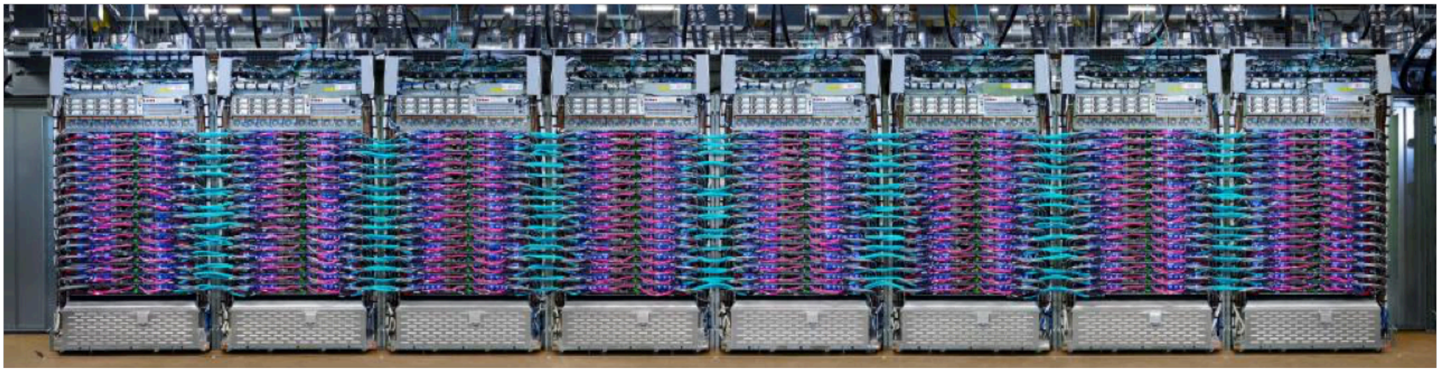

Photo Credit: Google

A row of TPU 3.0 server pods inside a Google data center, cooled with the company's custom direct-to-chip liquid cooling system

Such a system, CoolIT asserts, could eliminate the need for a separate water chiller attached to the airflow system, thus potentially reducing total cost of maintenance by about one-fifth.

Air's Limits

In a November 2016 company blog post, Upsite, which sells data center cooling solutions, suggested a different hybrid integration option with a similar purpose: rear-door heat exchangers for each rack, coupled with the facility's existing water chilling plant. The company acknowledged that this was by no means a new option, that in the early 2000s the concept was first floated as a heat-dissipation solution for racks exceeding what was then the generally accepted rack power density maximum of 8kW.

"Twenty years ago, the average density in a cabinet was somewhere around 3kW, or maybe even a little less," Upsite's Strong said. "And every few years there's somebody screaming about how it's going to start jumping to much higher densities. Yet, the average continues to just steadily, slowly, march up. It's still in the single-digits — still creeping up."

If we could plot the progress of this slow march on a chart, can we determine the absolute maximum power density that air-based solutions can handle? Strong answered yes, but with a caution: the equation won't be the same for everyone.

The relationship between power, delta-T, and airflow rate (measured in cubic feet per minute, or CFM) is fixed. "If you know any two of those, you can calculate the third," he explained. "Depending on the delta-T of the server and the IT load — the power it consumes — you're going to have a required CFM for that component."

With that CFM data in hand, can a data center operator calculate whether a liquid cooling solution is absolutely necessary? Schneider's Capes believes the answer is a definitive maybe.

While a traditional data center is designed to handle 4kW to 8kW per rack, a hyperscale data center (one that runs newer classes of distributed applications in an environment that's designed to scale out) may be designed around 12kW to 16 kW. At those density ranges, both types of facilities can be cooled with air. But, "I can tell you that the hyperscalers don't have an immediate answer to deal with heat loads that exceed 16kW per rack," Capes said.

And things are continuing to heat up in hyperscale data centers. Schneider has already seen densities north of 100kW per rack in some instances, he said. A 36U rack at 3kW per rack unit brings it to 108kW.

Liquid-based rear-door heat exchangers like Upsite's are one option hyperscalers have gone with. Google recently revealed that its latest-generation TPU 3.0 processors for machine learning pushed data center power density beyond what an air-based solution could handle. As a result, the company has been retrofitting racks across its data center footprint with a custom direct-to-chip cooling solution. (Google declined to specify power densities inside its data centers.)

In less extreme cases – ones traditional enterprise data center operators are more likely to encounter – a lot can be done before air's cooling limits are reached.

"One of the drivers for liquid cooling is ignorance about what can be done with air," Upsite's Strong said. "A lot of people think the upper limit for air is 10, 20kW in a cabinet. But there are numerous people doing 40, 45kW in a cabinet, just with air."

Tom Cabral, product application specialist with Chatsworth Products, a maker of cabinets and enclosures for high-performance and hyperscale data centers, agreed. "With a greater strategy of airflow management, you can solve a lot of problems, and you don't have to go to direct liquid cooling," he said.

A typical enterprise customer is under the impression that their data center is starved for airflow. In many cases, however, it takes little more than rearranging equipment to change that perspective. It could be something as simple as changing the direction its cabinets are facing, for example.

Any full-scale solution, Cabral contended, must take into account two simultaneous snapshots of the current scenario: "micro" and "macro." A reorganization solution that takes macro into account can completely resolve the issues the customer believes to be micro. And this is where liquid cooling can exacerbate the debate, he said, by playing on customers' beliefs that a micro problem can only be addressed with a micro solution.

High Density at What Cost?

Cooling equipment typically comprises 35 percent to 42 percent of a data center's initial required capital investment, Steven Bornfield, a senior data center consultant at Chatsworth, said. Energy delivery, management, and redundancy can command 30 to 40 percent of CapEx.

"Whenever you look at the server racks, you're at one or two percent of the total cost of the data center," he said. "So, we're going to spend all this time, money, and effort on advanced cooling systems that are well above the one-to-two-percent value. That's where I really run question marks, and say, folks, are we really focusing on the right things? Your key elements — power and cooling — are your biggest buckets."

The next largest cost center after power is labor, Bornfield continued. With direct-to-chip liquid cooling systems, which introduce non-standard equipment form factors and the foreign element of coolant, labor cost may go up, which remedial planners often fail to consider.

"You're adding costs," he said. "That's why I question, does it make sense? In certain environments, it could, but overall? No."

Opinions vary on how much cost and operational complexity liquid cooling adds. Also, different technologies add different levels of complexity.

Immersion cooling requires an entirely different infrastructure design, since vertical racks are replaces with tanks of coolant, and air-cooling components are taken out of servers. There are also operational complexities associated with immersion.

Rear-door heat exchangers are relatively non-disruptive, since they mount on the backs of existing racks, while direct-to-chip liquid cooling requires replacing air-cooled heat sinks with liquid-cooled cold plates or heat sinks and running fluid connections into IT boxes and throughout the racks.

While Schneider's Capes said he thinks the added complexity of direct-to-chip liquid cooling is negligible, Chatsworth's Bornfield thinks it's enough to make an operator think long and hard before pulling the trigger on liquid cooling.

Changes On the Horizon

Capes foresees a future where liquid cooling is the principal option for data centers, and where colocation facilities offer immersion-cooled racks as premium options. In that future, the principles of data center design as we know them today no longer apply.

"Think of a data center that has no air movement with air handling units, no requirements for hot-aisle or cold-aisle containment," he said. "Even the whole IT rack structure changes: it no longer becomes a method to distribute cold air or exhaust air; it's simply a mounting structure for your IT equipment. There's so much cost that can be stripped away from the infrastructure and so much benefit that can be gained by operating at these higher water temperatures, more flexibility in site selection, better overall operating efficiency, and potentially the elimination of evaporative cooling."

Upsite's Strong warned that adopting liquid cooling to any degree changes the formula for implementing and managing air cooling. And many operators will have both technologies in their facilities.

Very few data centers will be "all-liquid," Strong said. "There's going to be a mixture. And that partitioning, or lack of partitioning, requires awareness of how to do this well — how to keep managing the air cooling in an efficient way."