The Advancement of Liquid Cooled Solutions: The Perfect Storm Brewing

A brief from the IT subcommittee of ASHRAE TC9.9

The data center industry is at a crossroads when it comes to cooling technologies and solutions. Though the most common and most recognized cooling solutions utilize air, new technologies are just on the horizon which requires liquid direct to rack and chip. The timeline for availability and adoption appears to be accelerating. It is the express interest of the ASHRAE TC9.9 technical committee, who is the industry authority on power & cooling trends and best practices, to author this technology brief to accelerate the preparedness and readiness of the industry. This article will not only highlight industry factors and influences anecdotally but will articulate the urgent need to prepare for liquid cooling technologies now.

Data center cooling has been evolving over the last 15-20 years. While water has always been the primary waste heat transport out of most buildings, it is once again creeping closer and closer to the electronics. From the perimeter, it moved to the row. From the row, it moved to the rack with rear door heat exchangers. And some have started cooling within the server at the chip. When looking at the industry over a longer period, it is not merely evolving; it is repeating.

Factors and Influences

Had this brief been written a year ago, the following factors would have been discussed as reasons to consider liquid cooling:

1. Disruption from workload type - driven by insatiable demand for more data, and the need for real-time data consumption and analytics. High performance analytics and AI crossing market segments from pure HPC play to more central, edge, and colo data center solutions
2. Data center infrastructure limitations - high heat density solutions utilizing air, driving footprint with capacity limited CRAHs. By using liquid direct to chip, or a hybrid variant, significant improvement in heat density can be attained extending the useful life of the data center
3. Energy reduction - pressure being applied to reduce the energy consumption of data center power. Water cooled solutions allow you to reduce PUE from greater than ~1.6 to less than ~1.1
4. Reduced capital- capital expense reduction for new data center builds
5. Waste heat reuse – predominately a European driven initiative to put the heat back into the grid

Now today, there are a couple of crucial additional points to discuss, both related to silicon performance.

1. Specialty SKUs- responding to the need for real-time data analytics, chip manufacturers are responding with higher frequency SKUs, raising thermal design power while lowering chip junction temperature requirements. Many of these (off roadmap) SKUs cannot be air-cooled. With multi-core parts topping out, the only option for increased performance is increased power. Expect a shift from off-roadmap SKUs to on-roadmap SKUs soon. Also expect the whole on-road SKU stack to creep upward, which relates to the next point.
2. Competitive landscape - NVidia, IBM Power, ARM, AMD, and Intel are all competing for a performance edge. The gloves are off to claim the #1 spot. (an accelerant to the fire…i.e., no

longer one chip manufacturer dominating this space trying to stave off the transition to non-air-cooled solutions).

Liquid cooling as driven by chip density

A visual look at the history of chip densities might help explain the challenge for air cooling and point out the inevitable transition to liquid cooling.

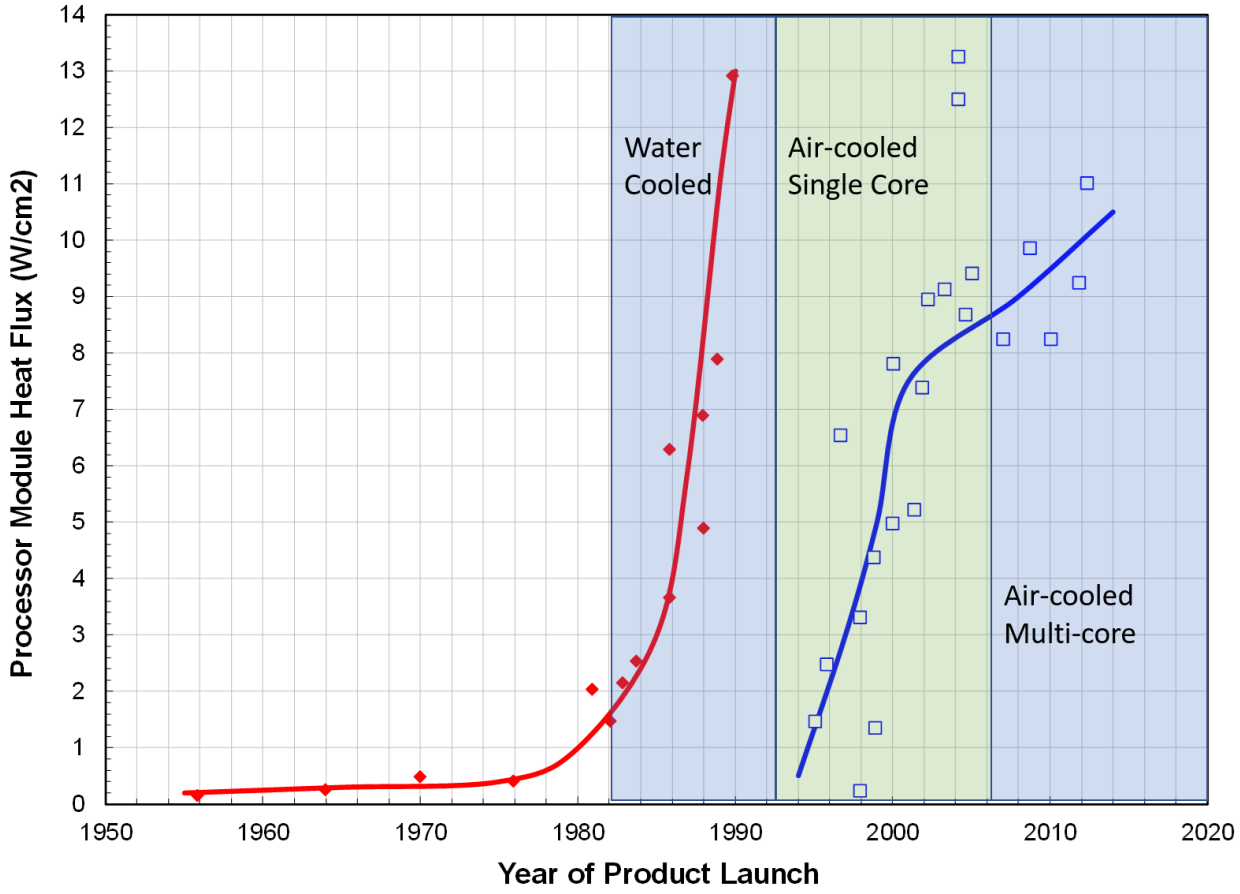

Each of the points represents a compute product, or more explicitly, the chip density in W/cm^2 in each product from various chip vendors. The products represented by solid diamonds were of earlier transistor technology called bipolar. In the early '90s, the industry transitioned to a different type of transistor technology, CMOS (Complementary Metal Oxide Semiconductor), represented by the squares. Bipolar devices, which had been air-cooled at once, fully transitioned to water cooling by the late '80s. CMOS brought about such a drastic reduction in chip power that the industry transitioned back to air cooling.

Fast-forward a decade or so, and IT vendors were, yet again, talking seriously about liquid cooling with the heat flux approaching 14 w/cm^2. All this talk subsided with the introduction of multi-core designs around 2005. All of the sudden, the same wattage chip became easier to cool once the primary heat source (core) within the chip was split and spread out. Performance gains were achievable by splitting cores without increasing overall chip wattage significantly. Running out of room now, the chips just

cannot get any larger and cannot continue to increase the core count. Performance gains will only be possible by increasing chip power, ushering in chips with mandatory liquid cooling.

Liquid cooling as driven by application performance requirements

For new workloads such as machine learning and artificial intelligence the need for specialized processors, such as graphical processing units (GPU), has driven significant power density to reach the performance level required. As an example, Table 1 provides an example over the last few generations of IT equipment and the server requirements to support an AI training framework. It is clear that in order to get generation-to-generation performance advantages, server designs have already exceeded the cooling capability of air-cooling. To support high-density server configurations such as a (2) CPU and (6) GPU 2U-server, water cooling is a requirement. In this example, designing an air-cooled server to meet the performance requirements would have resulted in a server that would have been highly inefficient from a space, power and cooling perspective.

Table 1 Evolution of server design on the way to water cooling

	2015	**2016**	**2017**	**2018**	
Server Size	2U	2U	2U	2U	2U
Cooling	Air	Air	Water	Air	Water
CPU	2x 325 W	2x 250 W	2x 250 W	2 x 250 W 16-core 2.7 GHz	2 x 250 W18-core 3.45 GHz
GPU	0	4x 300 W	4x 300 W	4x 300 W	6x 300 W
AI Training Framework Relative Performance	less than 1.0x			1.0x	1.4x

Summary

The data center industry has rejected the notion that water would ever again be a mainstream play. The history of Bi-Polar to CMOS, along with multi-core advances has created a false sense of security that technological advances would thwart the return to liquid. As experts and industry authorities in the field of electronics cooling, the return to liquid is imminent; it's back to the future. Engage with the IT manufacturers and equipment providers to make plans for water soon. Though a more comprehensive whitepaper will be released by TC9.9 later this year to further discuss this topic, the rate of evolution required this advance communication.