



Introduction

Two-Phase, or Not Two-Phase. After all, that is the question. At least it is for data center cooling architects and planners seeking to maximize cooling efficiency and performance for their ever-growing fleets of servers with the use of liquid-cooling. The use of water-based cold plate and dielectric immersion liquid cooling solutions have seen rapid increases in adoption over the past few product generations with a growth rate that continues to accelerate. With the acceptance (or to borrow the cliché “democratization”) of liquid in the data center, would-be customers of liquid-cooled servers are now turning their attention towards two-phase (or latent) liquid cooling technologies based upon these fluids’ non-conductive properties and potential for improved performance and efficiency when compared to water-based cold plate and single-phase immersion solutions.

The intent of this paper is to objectively compare the strengths and weaknesses of two-phase liquid cooling against traditional, single-phase, liquid solutions for data center servers. Using objective data collected on real world servers, single and two-phase cooling technologies are compared for cooling performance, pressure drop, and energy efficiency. Additionally, the paper also includes a qualitative discussion on some of the potential challenges that coincide with two-phase cooling design to better frame the decision landscape.

Background

Without a long, drawn-out, thesis on the rising cooling challenges of the modern data center (this author doesn’t get paid by the word), will simply summarize the current state of data center cooling as follows:

Recent missteps and stalling in generation over generation silicon lithography shrinks coupled with the rise of function-specific, discrete, accelerators have driven unprecedented increases in component Thermal Design Power (TDP) within the data center. Coincidental with rising energy costs and environmental objectives to optimize efficiency, this trend in heat density is fueling the proliferation of data center liquid cooling adoption in virtually every flavor it comes in.

From 2001 until 2017, there was very little movement in the top-end CPU heat dissipation requirements for a server. This trend was heavily disrupted when lithography-shrink misses and competitive threats forced Intel to ratchet up CPU power (heat) in order to meet the insatiable demand for higher levels of performance from server customers. Figure 1 below shows the generation over generation changes in CPU TDP from Intel’s mainstream Xeon processors which has begun to look a lot like the proverbial “hockey stick” growth curve.

Though Intel will inevitably shoulder a significant portion of the blame for the recent explosion in component heat density, they are certainly not the only company responsible. The buzz word of all buzz words, “Deep Learning” a

specific form of artificial intelligence, has promised limitless data analytics insights, machine automation, advances in medicine, and, at some point in the near future, a path for transitioning lead into gold. While the last part of that statement was made in jest, deep learning, or machine learning as a whole, has become an extremely valuable segment within the compute portfolio of many of the world's largest corporations. The challenge that the new AI wave has brought about is that the workload, itself, is extremely computationally intensive. Purpose-built accelerators designed to improve the throughput of deep learning models are now being pushed to 800 Watts and beyond as ASIC vendors each compete for the crown of neural network performance supremacy.

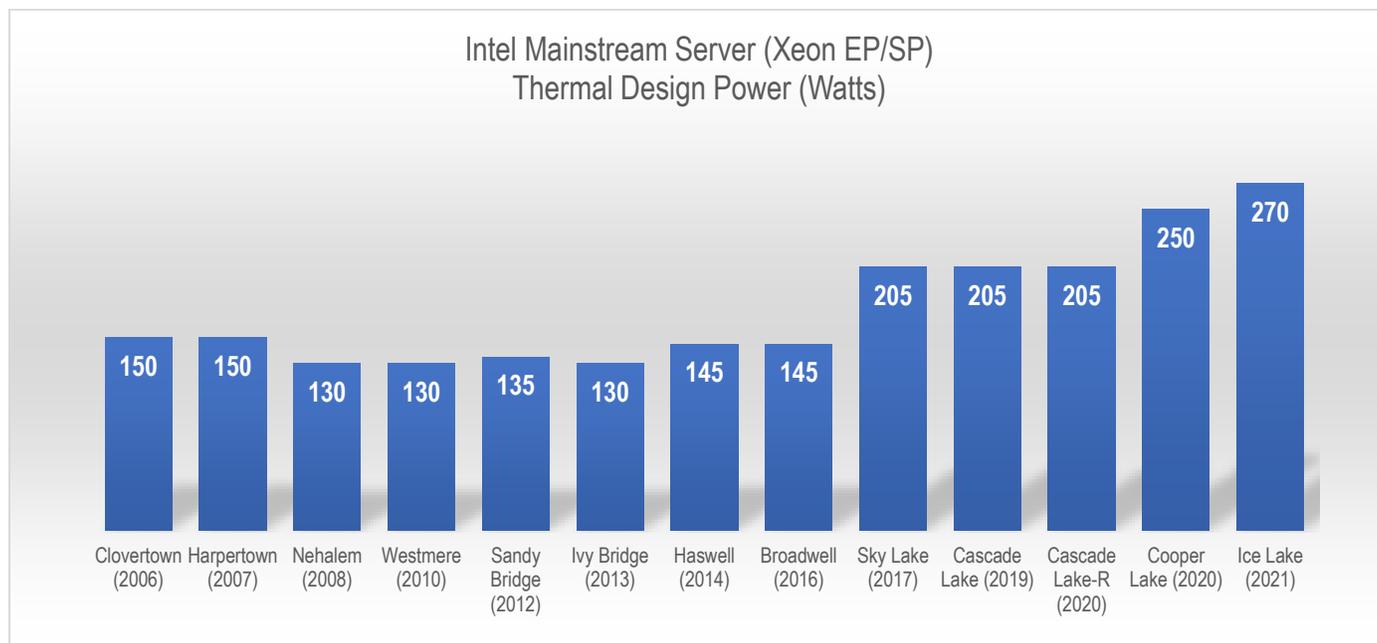


Figure 1. Intel mainstream server CPU Thermal Design Power shown from Clovertown in 2006 to Ice Lake in 2021 for top-end SKUs.

In order to support high heat flux CPUs and accelerators of various functions, many data center operators and server designers are turning towards liquid cooling in place of traditional air cooling for server components.

Why?

Liquid cooling affords substantially greater cooling performance and energy efficiency than what can be obtained using air. In many cases, cooling performance dictates limitations on server performance as a whole.

Does everyone need liquid-cooled servers?

Absolutely not. The capability of air-cooling technology has not remained stagnant over the years as component TDP has risen. Advances in heat transport devices such as 3D vapor chambers, loop heat pipes, ultra-high flow fans, and new, creative methods for expanding heat sink surface area have all significantly shifted the limitations of air-cooling over time. That said, even with recent advancements, there are still server components and physical configurations that exist today which cannot be cooled by air. Unfortunately, these advances in air-cooling performance have done very little to improve the energy efficiency limitations of air itself, brought about by air's miniscule volumetric heat capacity. Those organizations seeking top-end performance from their server components and maximization of energy efficiency will be drawn to liquid cooling. For many others, the end-to-end total cost of ownership is still better with air for at least the next few product cycles. Look for a future whitepaper to objectively breakdown the limits of air cooling and when a leap to liquid is going to make sense for most data center solution providers.

Types of Server Liquid Cooling

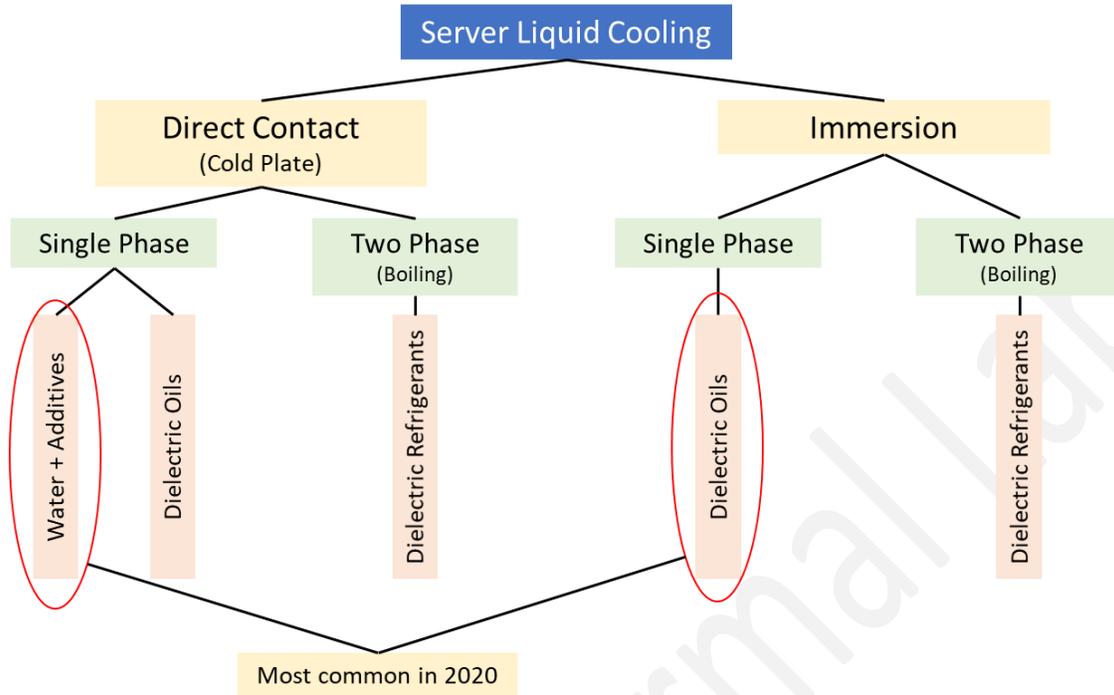


Figure 2. Diagram showing the most common categories of server liquid cooling found in contemporary data centers.

Figure 2 above is an attempt to bulk-categorize the different types of server liquid cooling commonly seen in datacenters. The notable absence of Rear Door Heat Exchanger (RDHx), easily the most prolific rackmount liquid cooling device in data centers today, is simply because this paper is focused upon liquid cooling within the server itself. The most popular implementation of server liquid-cooling uses cold plates attached directly to heat dissipating components within the server to direct the flow of liquid coolant which is typically water-based with additives to inhibit corrosion and biological growth. Additionally, single-phase immersion using low vapor-pressure oils is also becoming increasingly common in data center cooling, however, it does not really compete with single-phase water or two-phase liquid cooling in terms of heat transfer performance. That is not to say that single-phase immersion is without significant value; there are many scenarios where it may, actually, yield the best TCO. Therefore, single-phase immersion will be covered in a future publication.

Though single-phase cold plate cooling is the most common form of installed liquid cooling today, recently, significant attention is gravitating towards two-phase liquid cooling.

While two-phase cooling can technically describe everything from the internal physics of commodity heat pipes all the way to refrigerant-based DX cooling employed by various CRAC and row-cooling devices, this paper is focused, specifically, upon the use of low boiling point, dielectric, fluids that, either through direct contact or full immersion, transport heat away from server components through the transition of liquid to vapor phases of matter in latent heat absorption.

Why is Two-Phase Liquid Cooling gaining so much momentum?

There are three fundamental reasons why two-phase liquid cooling is garnering so much attention:

1. **Non-Conductive (dielectric) fluid:** The specialty fluids most commonly used in two-phase cooling possess high dielectric strengths; so much so that electronics can be directly immersed within the fluids without any risk of electric shock or component damage associated with the direct exposure to water. Liquid cooling with water-based coolants does present the risk of coolant leaks on expensive server components. Plenty of design strategies exist to mitigate the probability and risk of a water leak on these systems, however, none can 100% eliminate the risk of component damage like a dielectric coolant can.
2. **High Heat Transfer Performance Potential:** Boiling heat transfer coefficients can be as much as 10x larger than single-phase heat transfer coefficients under the right circumstances which creates a smaller temperature rise between components and the cooling fluid.
3. **High Efficiency Potential:** When substances change phases (freezing, melting, boiling, evaporating, condensing) they either absorb or release tremendous amounts of thermal energy (heat) depending upon which direction the phase change occurs. For many of the specialty fluids used in two-phase (latent) liquid cooling, they can absorb roughly 100x more heat per unit of volume when undergoing phase change compared to what they are able to absorb in single phase (sensible) heating. This means substantially less liquid is required to be pumped into a given server component in order to cool it which can reduce the amount of pumping energy required and drive-up efficiency.

In plain English: Two-phase liquid cooling eliminates the risk of component damage from leaks and presents the potential for incredible heat transfer performance and cooling efficiency. Changing phases means the required volume of liquid coolant is substantially less than that of single-phase cooling.

Table 1 below helps to illustrate these points. The Specific Heat of a fluid reflects how much heat it can absorb, per unit of mass, per degree of temperature increase in single-phase (sensible) heat transfer. The Latent Heat of Vaporization of a fluid reflects how much heat it can absorb, per unit of mass, as it transitions phase from liquid to vapor in two-phase (latent) heat transfer. As can be seen in the table, a fluid's latent heat capacity is typically several orders of magnitude larger than its sensible capacity.

Table 1. Thermophysical properties for pure water, water mixed with propylene glycol, four 3M engineer fluids, and common mineral oil.

	Unit	Water and Propylene Glycol		3M Novec				Mineral Oil
		Water	Water + 25% PG	HFE-7000	HFE-7100	Novec 649	FC-72	-
Boiling Point @ 1 atm	°C	100	122	34	61	49	56	349
Liquid Density	kg/m ³	980	995	1400	1510	1600	1680	870
Specific Heat	J/kg-K	4187	3787	1300	1183	1103	1100	1933
Sensible Volumetric Heat Capacity	kJ/m ³ -K	4103	3768	1820	1786	1765	1848	1682
Latent Heat of Vaporization	J/kg	2,256,000	1,870,000	142,000	112,000	88,000	88,000	977,000
Liquid Thermal Conductivity	W/m-K	0.63	0.498	0.065	0.065	0.065	0.057	0.143
Surface Tension	N/m	0.072	0.063	0.0124	0.0136	0.0108	0.01	0.035
Absolute Viscosity	cP	1	11.25	0.45	0.64	0.64	0.64	44
Dielectric Strength	kV/in	varies		>25	>25	>40	>35	30-40
Dielectric Constant @ 1kHz	-	78.4	-	7.4	7.4	1.8	1.75	2.1
Global Warming Potential	GWP	0.001	-	530	297	<1	>5000	1.07
Ozone Depletion Potential	ODP	0	-	0	0	0	0	
Single Phase Cooling Performance		Best	Very Good	Poor	Poor	Poor	Poor	Mid-Low
Is the fluid electrically conductive?		Yes	Yes	No	No	No	No	No

As the saying goes, a picture is worth a thousand words. Figure 3 below shows the relative volumetric heat capacity of air, water, and 3M's Fluorinert engineered fluid operating in single phase heat transfer. It is by this metric that many liquid-cooling solutions marketers will boast about water's 4000x cooling efficiency advantage over air, as one can see from the chart, water can absorb 4000 times more energy per unit of volume than air.

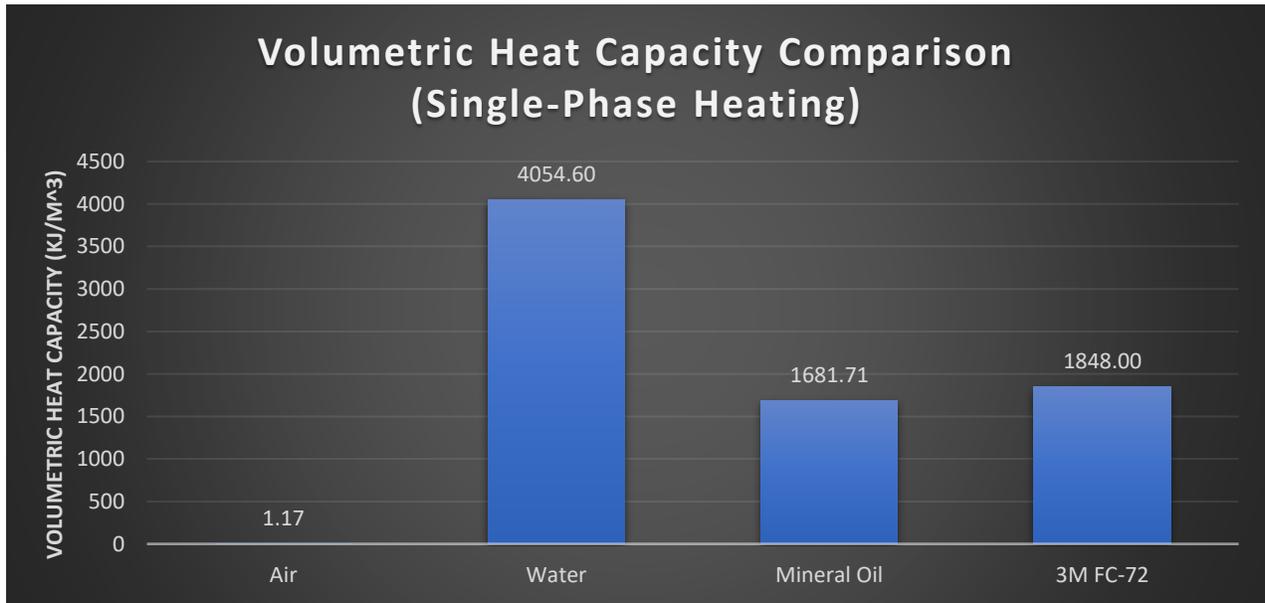


Figure 3. Sensible (single-phase) heat capacity comparison of common server cooling fluids

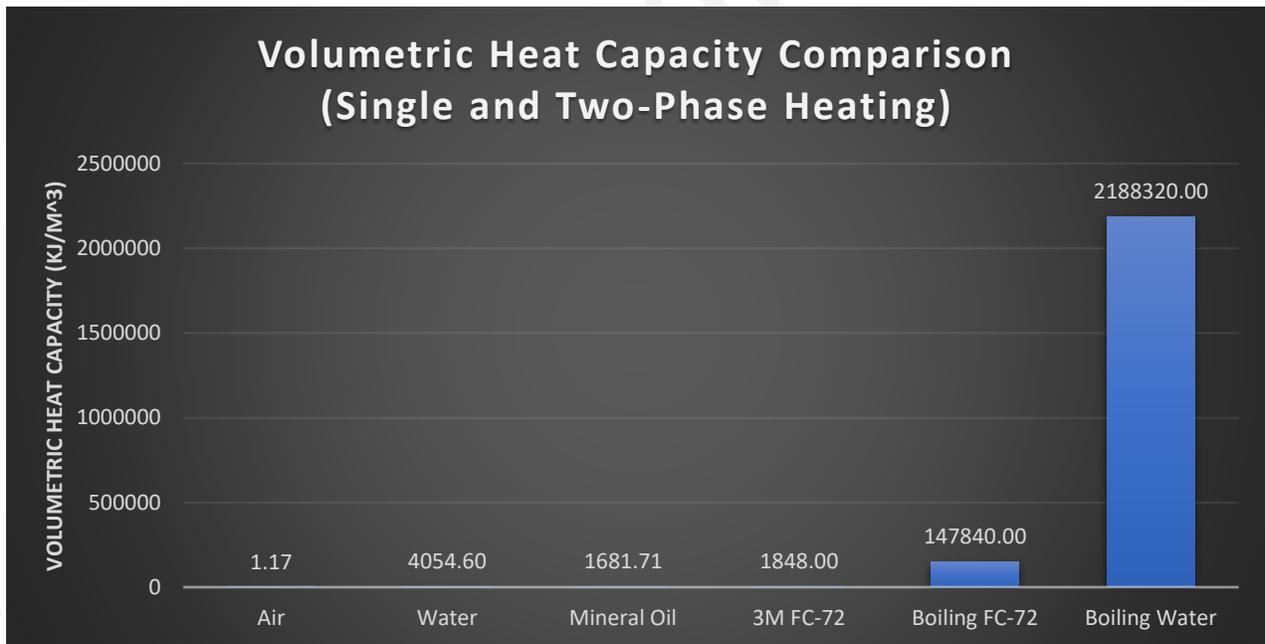


Figure 4. Latent (two-phase) heat capacity of water and 3M's Fluorinert compared to single-phase capacitance of common server cooling fluids.

Boiling liquids make this relative comparison even more interesting. Figure 4 shows the relative volumetric heat capacities of water and Fluorinert while operating in two-phase heat transfer. If comparing air cooling to single-phase water is like comparing the size of Earth with Jupiter, boiling either of the fluids and comparing them to air is more like comparing the size of the Earth with the size of our Sun.

Unfortunately, the story does not end there; as we will see, **Volumetric Heat Capacity Does Not Equal Cooling Efficiency!**

Simplifying the Cooling Loops for Comparison

With all the background out of the way, the next task is to discretize the primary components found in each of the three different cooling loop types being compared:

- **Single-Phase Cold Plate with water**
 - Server with cold plates and fans to cool internal components
 - Rackmount or Floor-Standing Coolant Distribution Unit (CDU)
 - Facility Water Supply
 - Facility Cooling Pumps
 - Facility Cooling Tower
- **Two-Phase Cold Plate with 3M's Novec HFE-7100**
 - Server with cold plates and fans to cool internal components
 - Rackmount or Floor-Standing Coolant Distribution Unit (CDU)
 - Facility Water Supply
 - Facility Cooling Pumps
 - Facility Cooling Tower
- **Single-Phase Immersion with Dielectric mineral oil variants**
 - Server immersed in fluid (no fans or cold plates)
 - Floor-standing Coolant Distribution Unit (CDU)
 - Facility Water Supply
 - Facility Cooling Pumps
 - Facility Cooling Tower
- **Two-Phase Immersion with 3M's Novec FC-72 (Fluorinert)**
 - Server immersed in fluid (no fans or cold plates)
 - Vapor to liquid condensers within the rack
 - Facility Water Supply
 - Facility Cooling Pumps
 - Facility Cooling Tower

The first thing to notice when comparing the elements that comprise each loop type is that the path from “chip to outside air” has a significant amount of common infrastructure between all three loops. ***To that end, this paper will make the stated assumption that facility cooling towers and pumps needed for all three loop types will more or less operate at similar efficiency levels*** to keep the focus on the server-liquid cooling performance and ensure apples to apples comparisons.

Breaking Down Cooling Loop Efficiency

Next, I have adapted a once popular term from the world of thermodynamics for the purpose of simplifying the comparison of operational efficiency between each cooling loop. The term is “*Coefficient of Performance*” or *COP*. *Coefficient of Performance* is simply the ratio of usable work output divided by the amount of energy consumed in the process. While *COP* still exists, it has largely been replaced by more application-specific derivations of the formula. *EER* or *SEER* would be examples of *COP* transitioning into a more application specific formula for air conditioning based upon seasonal temperature fluctuations. The simplified equation for *COP* is shown below:

$$COP_{device} = \frac{\text{Heat Removed by Device}}{\text{Power Consumed by Device}}$$

Using this definition of *COP*, we can then derive a relationship between *COP* and the much more popular datacenter term *PUE* or *Power Usage Effectiveness*. Before jumping into that derivation, it is important to quickly normalize the meaning of *PUE* in the context of this paper. The official definition of *PUE* is the ratio of total facility power divided by the amount of power consumed by the IT equipment (predominately servers).

$$PUE = \frac{\text{Total Facility Energy}}{\text{IT Equipment Energy}}$$

There are probably no fewer than 5 distinct industry standards bodies, today, who continuously debate new formulations of *PUE* to address several weaknesses that this metric has for accomplishing its stated purpose of quantifying data center efficiency. Before anyone reading this gets too fired up about *PUE* and all of its highly contentious shortcomings, and there are several of them, let us quickly rearticulate the term for the sole purpose of this paper that it might be more useful in communicating efficiency performance pertinent to the subject at hand: liquid-cooled servers. Therefore, please accept the temporary tweak of cooling-specific *PUE* to mean the following:

$$PUE = \frac{\text{server IT power} + \text{cooling equipment power}}{\text{server IT Power}}$$

Where *cooling equipment power* will reflect the total power consumed by critical (named) devices responsible for maintaining compliant temperatures of components operating in a server that lives within a datacenter. In this case, server fans will be included in the *cooling equipment power* sum where they are often omitted from traditional *PUE*.

From this tweaked cooling-specific *PUE* and the aforementioned *COP*, the unifying relationship then becomes:

$$PUE_{device} = 1 + \left[\frac{1}{COP_{device}} \right]$$

And where multiple elements function together to create the complete facility cooling loop the expanded formula then becomes:

$$PUE_{Loop} = 1 + \left[\frac{1}{COP_1} + \frac{1}{COP_2} + \dots + \frac{1}{COP_n} \right]$$

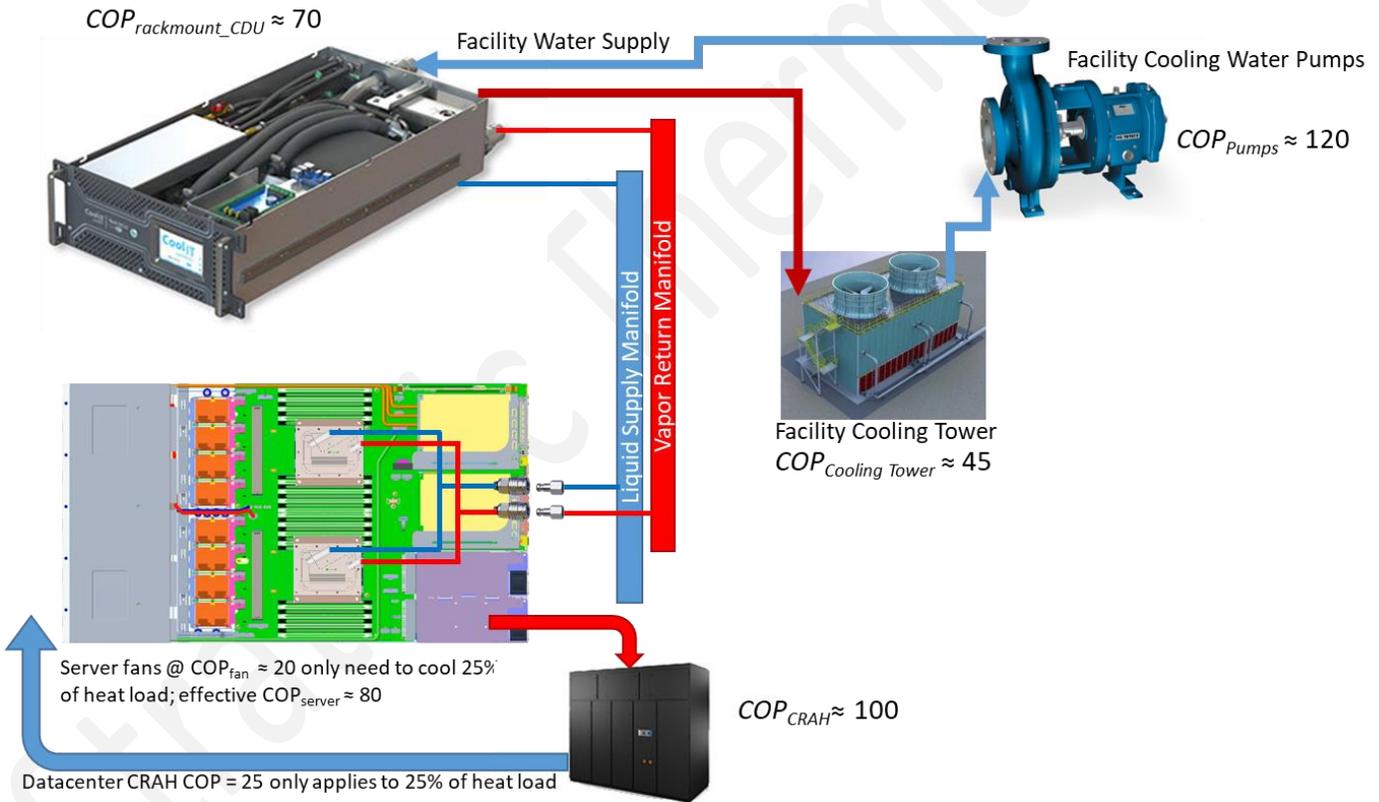
Where the subscripts 1, 2,... and *n* simply denote a given entity within the loop.

For a guy who doesn't get paid by the word, that's a whole lot of pretext. On to the data....

Single-Phase Cold Plate Liquid Cooling

Figure 5 below shows the basic elements of a cold-plate liquid cooling loop that utilizes a water-based coolant operating in single phase (sensible) heat transfer. Do not interpret the singular value shown for the $COPs$ of each device to suggest that these are static, unchanging values; they are not. COP depends heavily upon operating conditions like heat load and temperature gradients from components to external heat dumps. That said, these are real values and upon researching one will find these fit well into the expected *average* performance numbers for each device. Also, it should be noted that the same environmental perturbations that would affect these COP values would have a similar impact across all three loops, so the relative impact becomes moot.

The single-phase cold plate loop is relatively straight forward: quick disconnects provide the liquid coolant supply and return connections from a rackmount manifold into each server. Cold plates attach to the server CPUs and system fans are responsible for cooling the remaining server components (i.e. storage drives, memory, PCIE devices, etc.). With the CPUs consuming roughly 75% of the total server power and being cooled by liquid cold plates, the remaining system heat left for the fans to cool is typically around 25% of the total server power which is then cooled by the data center air handler (CRAH or CRAC). CPU heat is transported via water-based coolant from the server CPUs to either a rackmount or floor standing CDU and transferred to a facility water supply via liquid-to-liquid heat exchanger. The facility water supply then transports heat using centralized water pumps to a facility cooling tower where the heat is finally rejected into the outside air. The overall *cooling PUE* for this type of setup (inclusive of cooling fans within the server and CRAH units) is approximately 1.06.



$$PUE = 1 + \left[\frac{1}{COP_{Servers}} + \frac{1}{COP_{CRAH}} + \frac{1}{COP_{CDU}} + \frac{1}{COP_{Pumps}} + \frac{1}{COP_{Cooling\ Tower}} \right] = 1.06$$

Figure 5. Cooling loop and efficiency approximations for single-phase cold plate liquid cooling of servers using a water-based coolant.

Two-Phase Cold Plate Liquid Cooling

Figure 6 below shows the basic elements of a cold-plate liquid cooling loop that utilizes a dielectric coolant operating in two-phase (latent) heat transfer. It should look similar to single-phase cold plate as the diagram is, in fact, a copy and paste from the single-phase liquid cooling loop diagram directly above. Both cooling methods invoke the same fundamental building blocks including quick disconnects, cold plates mounted on top of heat dissipating components within the server, a brazed-plate liquid to liquid heat exchanger, a circulating pump to move the coolant throughout the loop, and server fans working with CRAH units to cool the remanent heat not captured by the cold plates. The overall *cooling PUE* for this type of setup (inclusive of cooling fans within the server and CRAH units) is, again, approximately 1.06 as shown in the figure below.

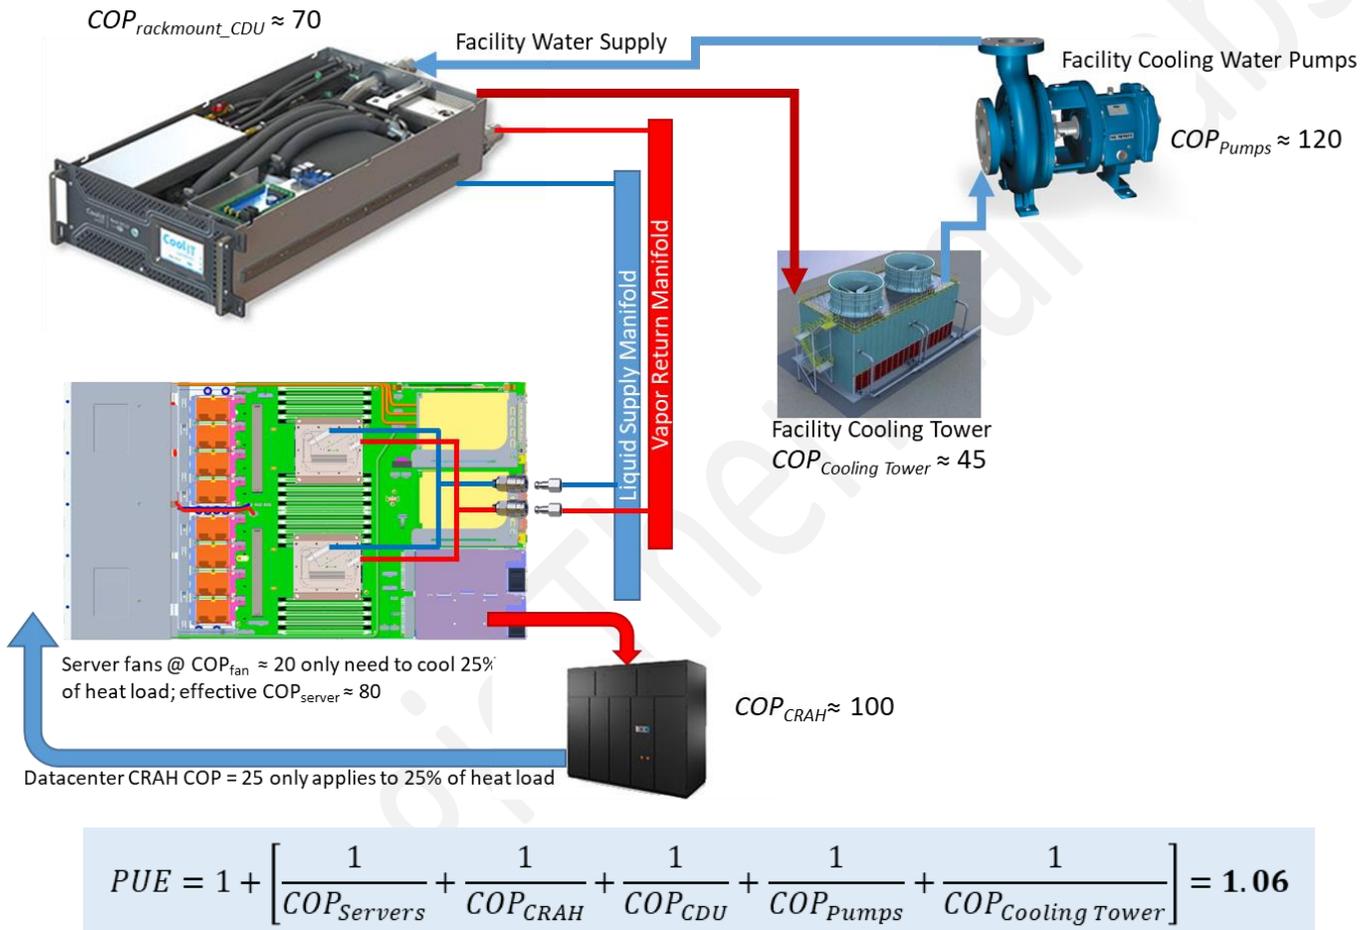


Figure 6. Cooling loop and efficiency approximations for two-phase cold plate liquid cooling of servers using a saturated, dielectric, coolant.

But wait, I thought that two-phase cooling loops were supposed to be drastically more efficient than single-phase cooling loops – what gives?

When you compare the key pieces of infrastructure between a single-phase and two-phase server cooling loop, it should become apparent that most of these pieces don't change at all. If we are trying to stay apples to apples, items like the facility cooling tower, the server fans, CRAH, etc. are still rejecting the same quantities of heat in both scenarios and doing so at what should be the same temperature gradients. The only component within the loop that will undergo significant change in operation is the technology coolant side of the CDU itself which will be pumping liquid supply to servers and returning a mixed liquid-vapor coolant back to the CDU heat exchanger where the vapor will condense back into a liquid. Hypothetically, if the CDU responsible for transferring heat from the two-phase cooling loop to the facility water supply consumed zero energy whatsoever, the overall cooling PUE might improve from ~1.06 to ~1.05. This would be a reasonable and desirable improvement in efficiency, however, the pumping power required from the CDU in a two-phase cooling loop isn't 0 Watts. In fact, the actual operating power draw of pumps in both single and two-phase loops may not be that dissimilar.

If the reader has been tracking thus far, the next question they should be asking is:

You already told me that two-phase cooling requires substantially less liquid coolant volume per unit of heat rejection, how, then, is the pumping power similar with single-phase?

To answer this, we need to go a step deeper. The Law of Conservation of Mass states that, in a closed system, mass cannot be created or destroyed. In other words, the mass flowrate into a server must equal the mass flowrate out of a server. Mass flowrate (\dot{m}) is simply the volumetric flowrate (Q measured in GPM, CFM, LPM, m^3/s , etc.) multiplied by the density (ρ) of the substance.

$$\dot{m} = \rho Q$$

During two-phase heat transfer, the liquid coolant transitions into a vapor or gas which has a substantially lower density than the liquid itself. For example, the widely popular HFE-7100 series fluid from 3M has a liquid density of $1,300 \text{ kg/m}^3$ while the vapor density is roughly 9.6 kg/m^3 .

What this means is that if the liquid coolant flowing into a server undergoes 100% transition into vapor (full latent heat transfer), the exiting volumetric flowrate of the vapor will be 135 times larger than the entering liquid volumetric flowrate. The very high vapor volumetric flowrates on the return side of a two-phase cold plate loop will, in turn, drive very high pressure drops through tubing, manifolds, and quick disconnects which, in many cases, can eliminate the pumping power benefits obtained on the liquid-side of the flow loop.

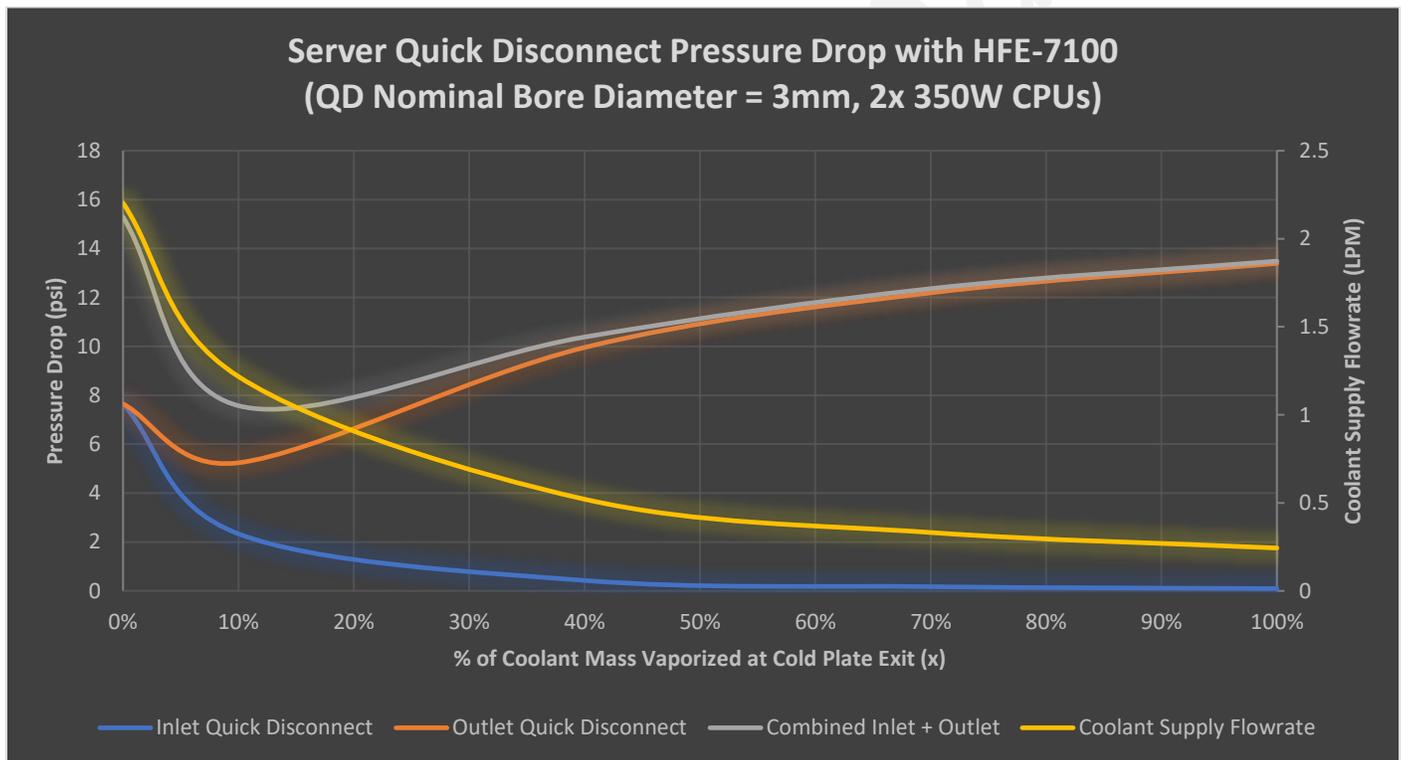


Figure 7. Estimated pressure drop across server quick disconnects based upon varying mass vaporization rates (exit vapor quality)

Figure 7 above depicts an example of quick disconnect pressure drop for a server with 2x 350 Watt CPUs being cooled at various percentages of two-phase vaporization using 3M's HFE-7100 fluid. In the graph, 0% indicates single-phase heat transfer and 100% indicates that all of the incoming liquid is vaporized. The graph illustrates an interesting trend: beyond the initial transition of a portion of the liquid undergoing phase change, increases in exiting vapor percentages also increase the combined pressure drop of the quick disconnects.

Figure 8 below shows a pressure drop and flowrate comparison between a single-phase water cold plate server loop and that of a two-phase cold plate server loop using HFE-7100 for the condition where 100% of the liquid undergoes phase change (lowest inlet liquid flowrate required). It is, again, interesting to note that while the required liquid

coolant flowrate is reduced by nearly 75% when switching from single-phase water to two-phase HFE-7100 (0.8 LPM down to 0.23 LPM), the pressure drop for the server loop more than doubles.

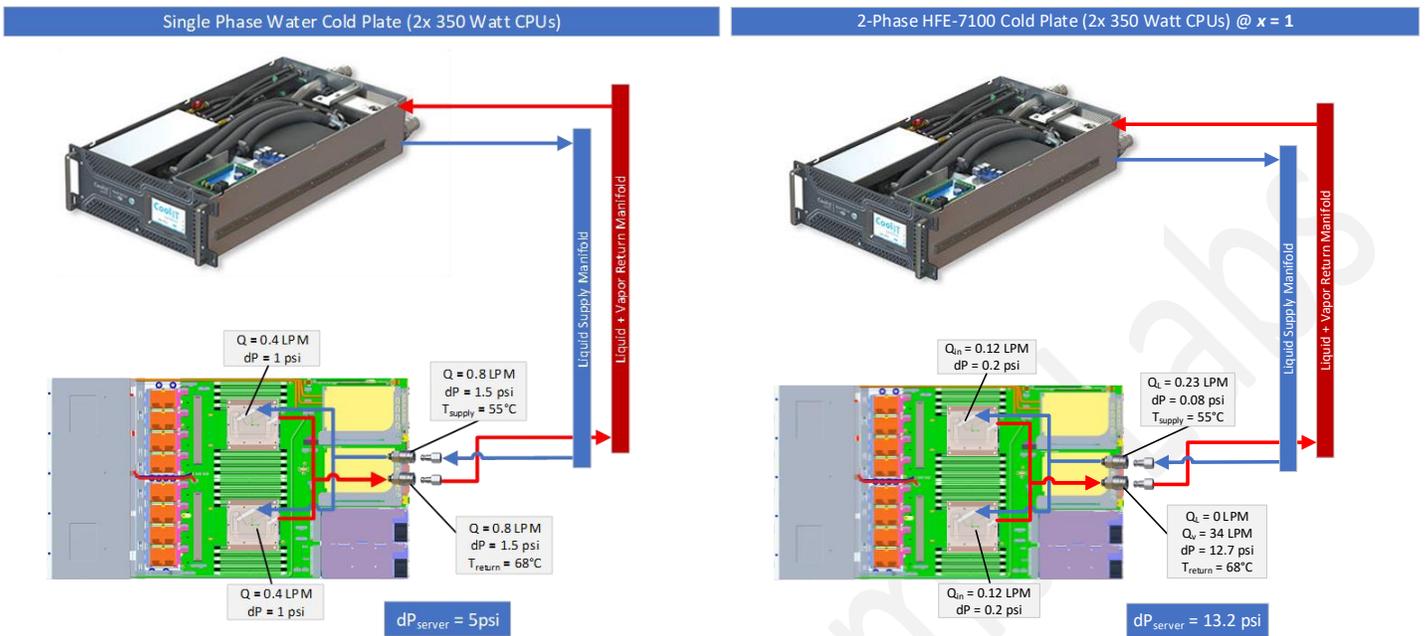


Figure 8. Comparison of flowrate and pressure drop for single and two-phase server cooling loops

So we see then that, while single-phase water requires substantially higher flowrates to cool the same 700 Watt heat load, two-phase HFE-7100 requires substantially higher pressure drop.

Pressure drop vs volumetric flowrate: which of these has the greater impact on required pumping power for the liquid cooling loop? Both of these quantities are, actually, both equally weighted in determining the pumping power required to cool a device. The equation below calculates the *mechanical work* or *hydraulic power* for a pumping system:

$$\text{Hydraulic Power (Watts)} = Q(\text{LPM}) * dP(\text{psi}) * 0.113$$

The actual power consumed by the pump in the cooling loop is then equal to the hydraulic power divided by the efficiency of the pump:

$$\text{Actual Pump Power (Watts)} = \frac{\text{Hydraulic Power (Watts)}}{\eta_{\text{pump}}}$$

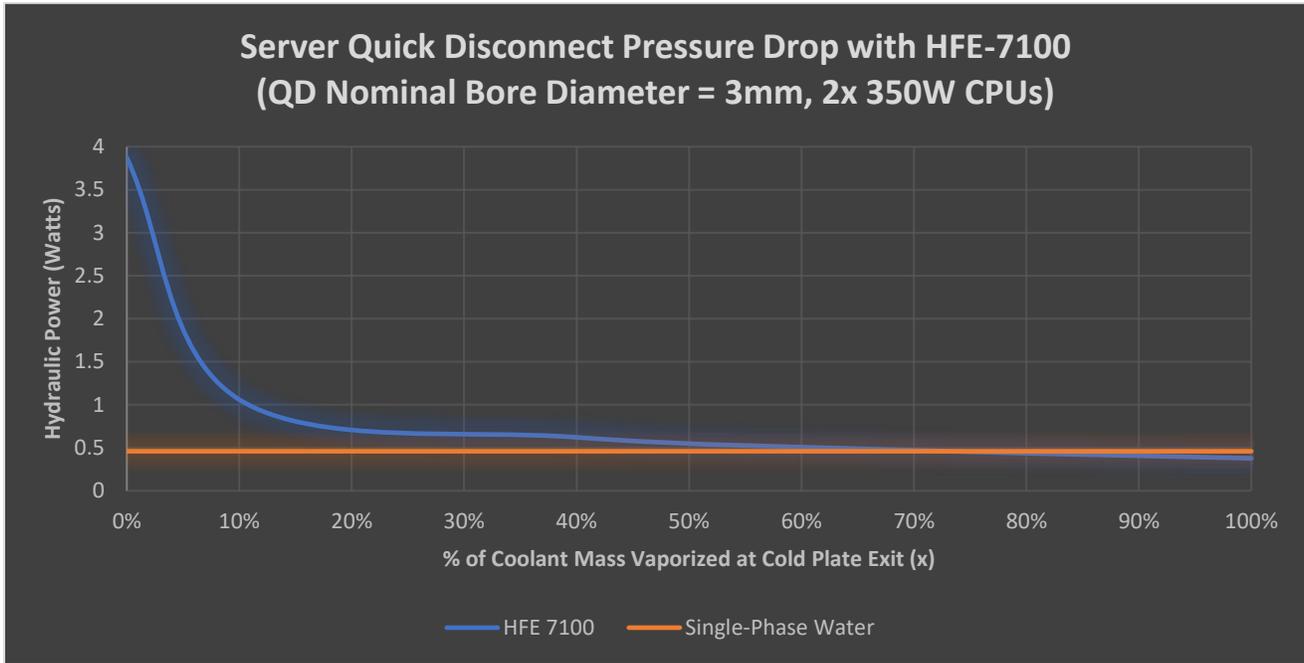


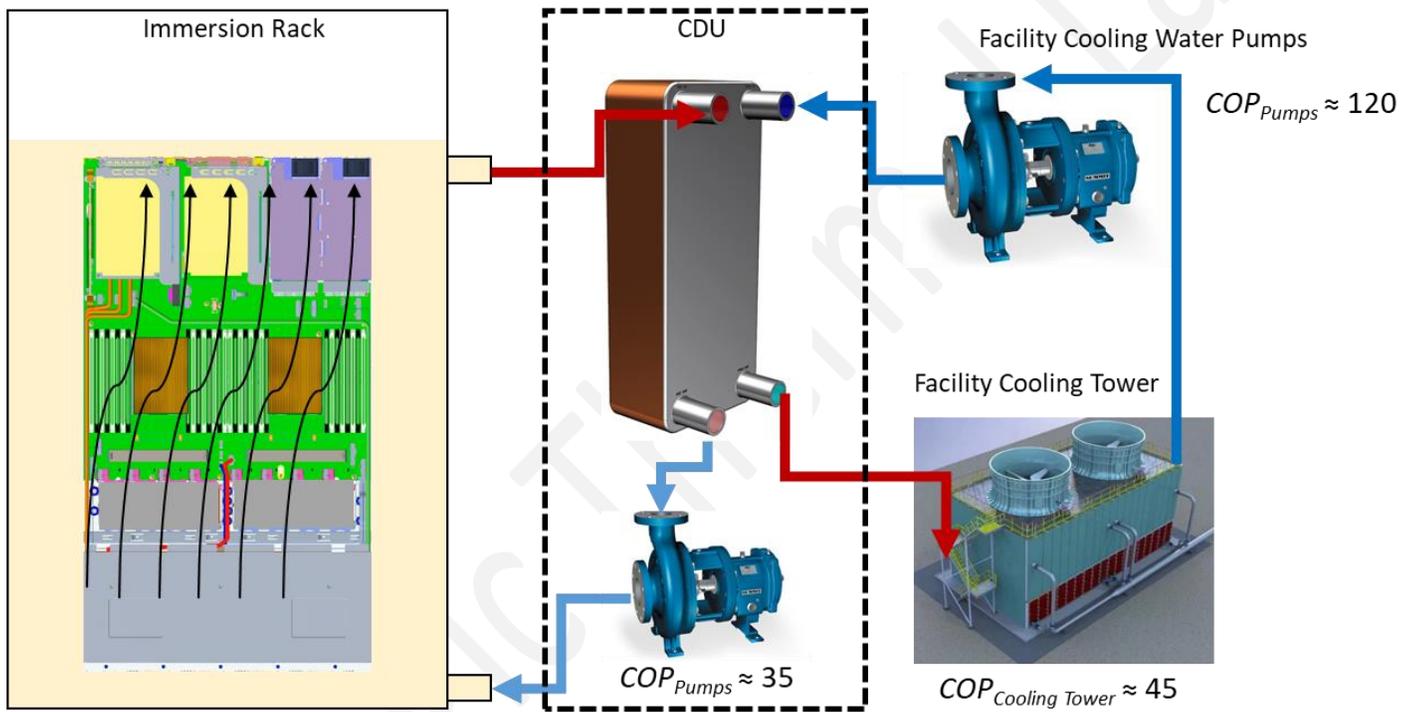
Figure 9. Hydraulic power (“pumping power”) required to cool a server with 2x 350 Watt CPUs at varying vaporization rates.

Figure 9 above uses the relationship of flowrate and pressure drop to calculate the hydraulic power required for the two-phase cooling loop using the same data presented in Figure 5. The orange line provides the single-phase water hydraulic power for comparison. As can be seen in the chart, at very high vaporization rates (above 75%), the hydraulic power for two-phase cooling does drop below that of single-phase water, however, only slightly so. It should also be noted that the pressure drop from rack components like the liquid-vapor manifold, quick disconnects into the CDU, and the liquid to liquid heat exchanger are completely omitted from this study and, while flowrate will remain constant, the pressure drop of each of those devices will add in series to the complete loop pressure drop.

Based upon the preceding discussion on flowrate and vapor-side pressure drop, it is likely that, in most cases, the two-phase cooling loop will have a pumping power requirement that is greater than or equal to a single-phase water-cooling loop.

Single-Phase Immersion Liquid Cooling

Next to single-phase water cold plate cooling, direct immersion utilizing dielectric oil variations in single phase heat transfer is the 2nd most commonly deployed server liquid cooling technology in the data center today. In single-phase immersion cooling, servers are fully submerged inside a tank filled with a non-conductive liquid coolant that typically has a very low vapor pressure (low evaporation rates, highly stable). Unlike cold plate cooling, single-phase immersion doesn't use a pump to create forced convection through an encapsulated heat sink, but rather, leverages natural (or free) convection to dissipate component heat into the fluid. As the fluid near each heated component increases in temperature, its density decreases allowing the hot liquid to rise and the cold fluid to sink. A Coolant Distribution Unit (CDU) pumps the warm oil through a liquid-to-liquid heat exchanger where the Facility Water Supply (FWS) absorbs the heat from the servers and then returns the cool fluid to the immersion tank as shown in Figure 10 below.



Servers directly immersed in Dielectric Oil

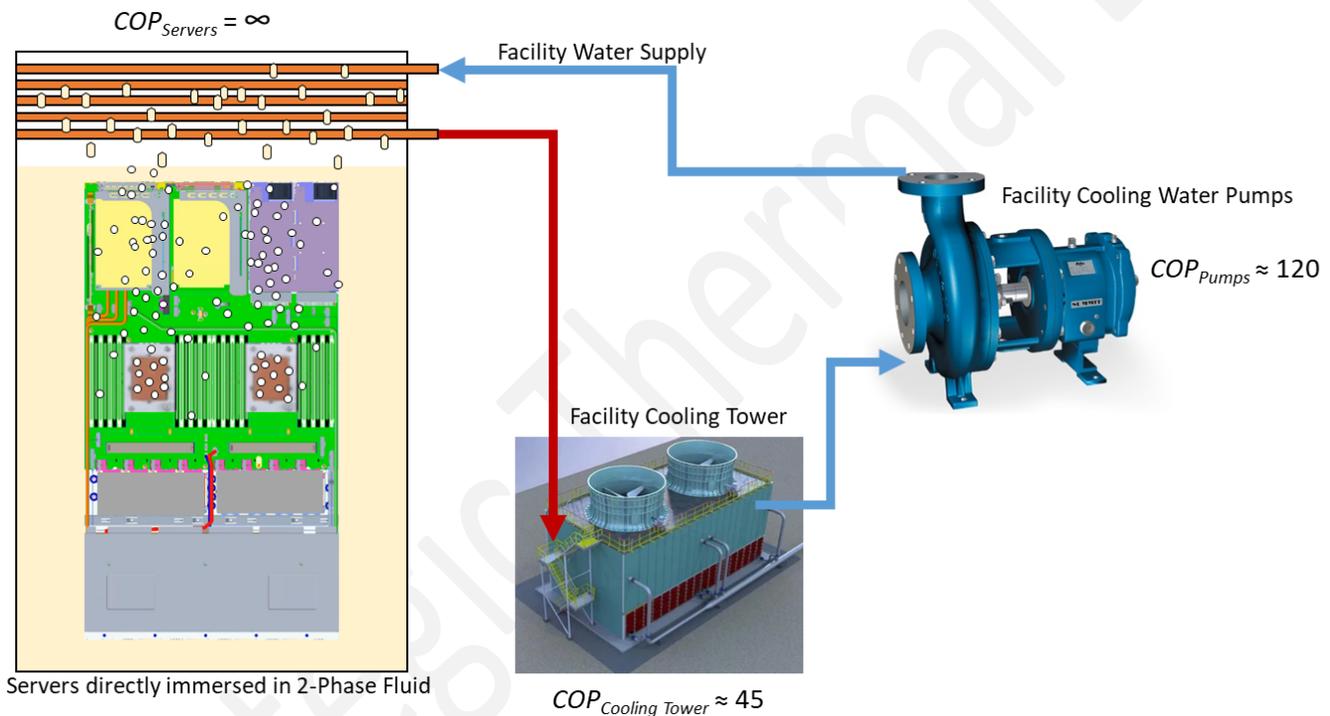
$$PUE = 1 + \left[\frac{1}{COP_{CDU}} + \frac{1}{COP_{Pumps}} + \frac{1}{COP_{Cooling Tower}} \right] = 1.06$$

Figure 10. Cooling loop and efficiency approximations for single-phase immersion liquid-cooling of servers in a dielectric oil.

Removal of all server fans and the necessary air movers throughout the data center (at least those necessary for server cooling) does have significant cooling energy reduction-*potential*, as the list of contributing devices to the overall cooling-specific PUE has shrunk considerably. There is a catch, however, in that single-phase oils, having roughly 40x higher absolute viscosity than water on average, require significantly more energy to pump than water while carrying with them far less energy than water with less than half the volumetric heat capacity. The result is a cooling-specific PUE that looks much like single and two-phase cold plate solutions even though they still depend upon fans for a portion of their cooling.

Two-Phase Immersion Liquid Cooling

Unlike two-phase cold plate liquid cooling, two-phase immersion liquid cooling does reflect a significant divergence in physical infrastructure from the traditional single-phase water liquid cooling paradigm. In two-phase immersion, server hardware is fully submerged inside a tank filled with the saturated, dielectric coolant. In this cooling embodiment, heat from the server components causes the dielectric liquid to boil and transition into vapor bubbles. Gravity drives the vapor bubbles up and away from the heated components where the vapor rises into a headspace, contacts an embedded liquid to air heat exchanger, transfers the server heat into the facility water supply and then causes the vapor to condense and precipitate back down to into the tank. Figure 11 below illustrates the basic components of the two-phase immersion liquid cooling loop. In two-phase immersion cooling, 100% of the heat rejection from the servers goes into liquid eliminating the need for any server or data center air movers used in server cooling. Furthermore, the liquid to air heat exchangers (condensers) found in the headspace of immersion tanks utilize materials and hydraulic diameters that are similar in composition to those employed on the primary side of the heat exchangers within a CDU – meaning there is no need to have a CDU in the loop at all.



$$PUE = 1 + \left[\frac{1}{COP_{Pumps}} + \frac{1}{COP_{Cooling\ Tower}} \right] = 1.03$$

Figure 11. Cooling loop and efficiency approximations for a two-phase immersion liquid-cooling of servers in a saturated, dielectric, coolant.

The complete omission of any air movers within the server *cooling PUE* calculation sets two-phase immersion apart from the other two liquid cooling embodiments from an efficiency perspective with cooling *PUE* in the 1.03 range.

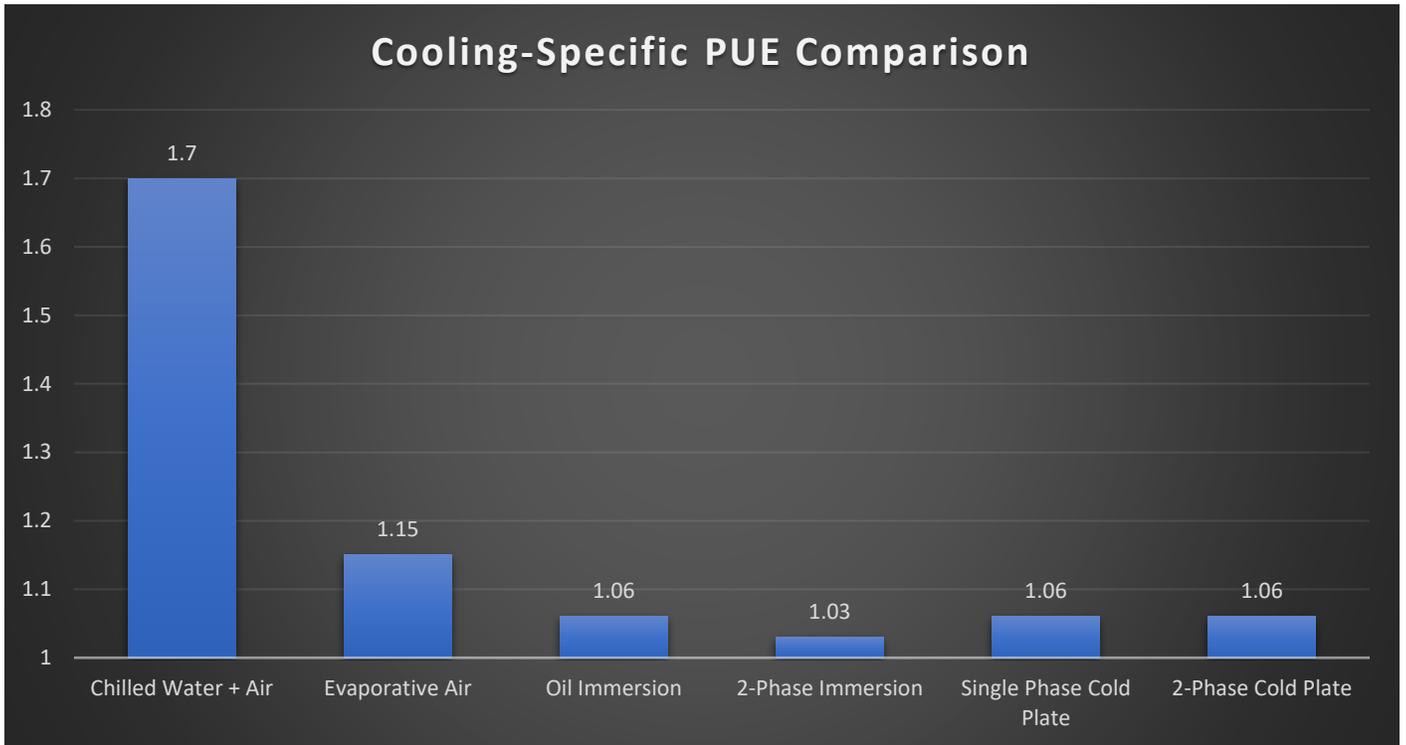


Figure 12. Cooling-specific PUE summary for the 4 liquid-cooling technologies discussed alongside popular air-cooling technologies.

Conclusion on Efficiency: Interestingly, single-phase cold plate, two-phase cold plate, and single-phase immersion liquid cooling have very similar efficiency profiles while two-phase immersion cooling offers significant improvements over all three. The biggest gains in PUE improvement come where chillers are entirely removed from the loops; competing for crowning efficiency in the domain of liquid technologies gets pretty close to asymptotic.

Heat Transfer Performance

The next objective of this discussion is to compare the relative heat transfer performance characteristics of the three forms of liquid cooling (single-phase cold plate, two-phase cold plate, and two-phase immersion). After all, the expectation with any form of liquid cooling is that all air-cooled, glass-ceilings limiting component thermal design power are decisively shattered, allowing for unprecedented new heights in silicon heat density and workload performance.

In a similar fashion to the technique used to analyze efficiency, the best way to compare the relative heat transfer performance of the three liquid cooling flavors is to discretize each element in the heat transfer path and consider their respective impacts to the temperature rise of the end component (in this case a CPU). As before, we will first consider cold plates and then pivot to immersion. Figure 13 below is an illustration of the physical interfaces responsible for transferring heat generated within the CPU die to the fluid that passes through a cold plate.

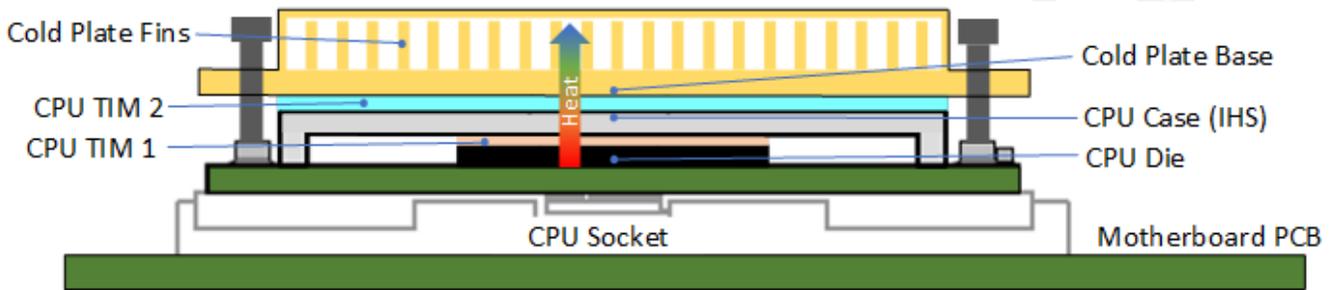


Figure 13. Diagram of the physical entities that map the heat transfer path from CPU die to cooling fluid through cold plate fins.

Thermal engineers often rely upon a term to quantify a system's thermal performance known as the *thermal resistance*. The *thermal resistance* (R_{TH}) of a heat transfer device simply reflects the temperature rise of the heat source over the cooling fluid per unit of heat input.

$$R_{TH} = \frac{T_{component} - T_{fluid}}{Power}$$

The combined thermal resistance (R_{TH}) of a cooling solution can be broken down into individual pieces that sum together in series. Using this technique provides perfect visibility into the relative contributions that individual thermal resistances make to the total temperature rise of a heated component.

The conduction heat transfer path depicted in Figure 11 from the CPU die to the cold plate fins is identical for both single-phase and two-phase cold plate liquid cooling. What changes is the nature of heat transfer from the surface of the cold plate fins into the liquid coolant – typically referred to as “convection” heat transfer.

The amount of energy that can transfer from a surface to a fluid per unit of surface area per degree of temperature difference is called the *heat transfer coefficient* (h). A larger value of h would allow for more heat to transfer from a surface to a fluid for a given finned surface area at a similar temperature difference.

$$h = \frac{q''}{(T_{surface} - T_{fluid})}$$

In virtually all shipping electronic cooling solutions today, the convection heat transfer is the dominant thermal resistance dictating the temperature difference between the heated component and the coolant (whether that coolant is liquid or air). Engineers salivate at the opportunity to increase the value of (h) but are usually constrained by physics to simply increasing the amount of finned surface area that the (h) applies to.

Consider the data in Table 2 below showing approximate ranges for the heat transfer coefficient of various forms of convection and notice the final two entries being boiling and condensation processes for water. While boiling organic

fluids like the dielectric coolants used in electronic immersion have substantially lower heat transfer performance than boiling or condensing water, they can, potentially, compete with water in single phase. **For this reason, the most alluring feature of two-phase heat transfer, for many engineers, is the potential for extremely high heat transfer coefficient (h) values compared to single phase heat transfer.**

Table 2. Typical ranges of overall heat transfer coefficient (h) for various forms of convection heat transfer

Type of Convection	Typical Range of h (W/m ² K)
Free Convection, air	2.5 - 20
Forced Convection, air	25 - 150
Free Convection, single-phase oil	100 - 250
Forced Convection, single-phase oil	500 - 1,500
Forced Convection, single-phase water	1,000 - 5,000
Forced Convection Boiling, HFE-7000	800 - 5,000
Pool Boiling, HFE-7000	1,000 - 9,000
Condensation, HFE-7000	1,000 - 3,500

** Typical range for overall heat transfer coefficient (U) established from empirical data of heat transfer geometries common to server industry - not indicative of theoretical or local maximums.*

Having spent more than a decade working with various forms of multi-phase cooling specifically applied to data center servers and after working with many suppliers who ran their own experiments with two-phase liquid cold plates for servers, there is an unfortunate reality about the state of two-phase cold plate cooling today that needs to be noted: the empirical performance of two-phase cold plates using dielectric fluids often undershoots expectations.

Engineers who are even deeper into this field than I am might retort that the above statement simply means that we have had the wrong expectations about organic fluids boiling in cold plates, though I have never heard or seen a proper articulation as to “why” those expectations might be so poorly aligned to reality.

After a recent deep dive into some of the more contemporary publications from academia on the subject of two-phase heat transfer in microchannel cold plates and a prolonged revisit into data I had collected over the past 14 years there are three conclusions on this topic that I would like to share:

- The local heat transfer coefficient for two-phase boiling is highly dependent upon surface roughness characteristics of the channel where edges and surface imperfections are necessary for nucleation (bubble) formation to occur. Polished, smooth, copper surfaces common to the cold plate industry are not, necessarily, ideal for two-phase heat transfer coefficients.
- The local heat transfer coefficient for two-phase boiling in microchannels is highly dependent upon heat flux (q ” the amount of heat dissipated per unit of surface area).
- The local heat transfer coefficient varies significantly across a heated surface due to the variation in temperature gradient and local fluid quality. The average (often referred to as the “effective” or “overall”) heat transfer coefficient will always be significantly less than peak, local, heat transfer coefficients expressed in many marketing headlines. It is the “effective” heat transfer coefficient that truly defines the efficiency of a cooling device and is the focus of the data presented in this paper.

Within the world of boiling, there are multiple grades or levels of boiling (think about the first few bubbles that appear on the bottom of a pot of water on the stove vs a “rolling boil” that occurs after the water is fully heated). There is an entire spectrum of boiling grades ranging from “periodic nucleation” where an occasional bubble appears here and there all the way to “film boiling” where heated surface creates so much vapor that the liquid literally dries out and no

longer contacts the surface. Each grade of boiling has very different heat transfer performance. Furthermore, as fluid moves through a heated channel (i.e. the fins of a cold plate) the grade of boiling is constantly changing from entrance to exit as the fluid will advance towards higher grades the longer it is heated.

When operating in a heated state that is well below the “film boiling” regime (sometimes called “Critical Heat Flux” or CHF), the two-phase heat transfer coefficient increases linearly with heat flux.

Why is this important?

Heat sinks and cold plates are both devices that attempt to improve heat transfer by increasing the surface area that heat is dissipated through into a coolant fluid. By nature, cold plates are designed to decrease heat flux by increasing surface area. Single-phase forced convection, whether air or water, has no first order dependency on heat flux, however, two-phase boiling does.

Figure 14 below uses empirically collected data for the overall heat transfer coefficient of boiling HFE-7000 and single-phase water as a function of heat flux using a high-volume, production, server cold plate.

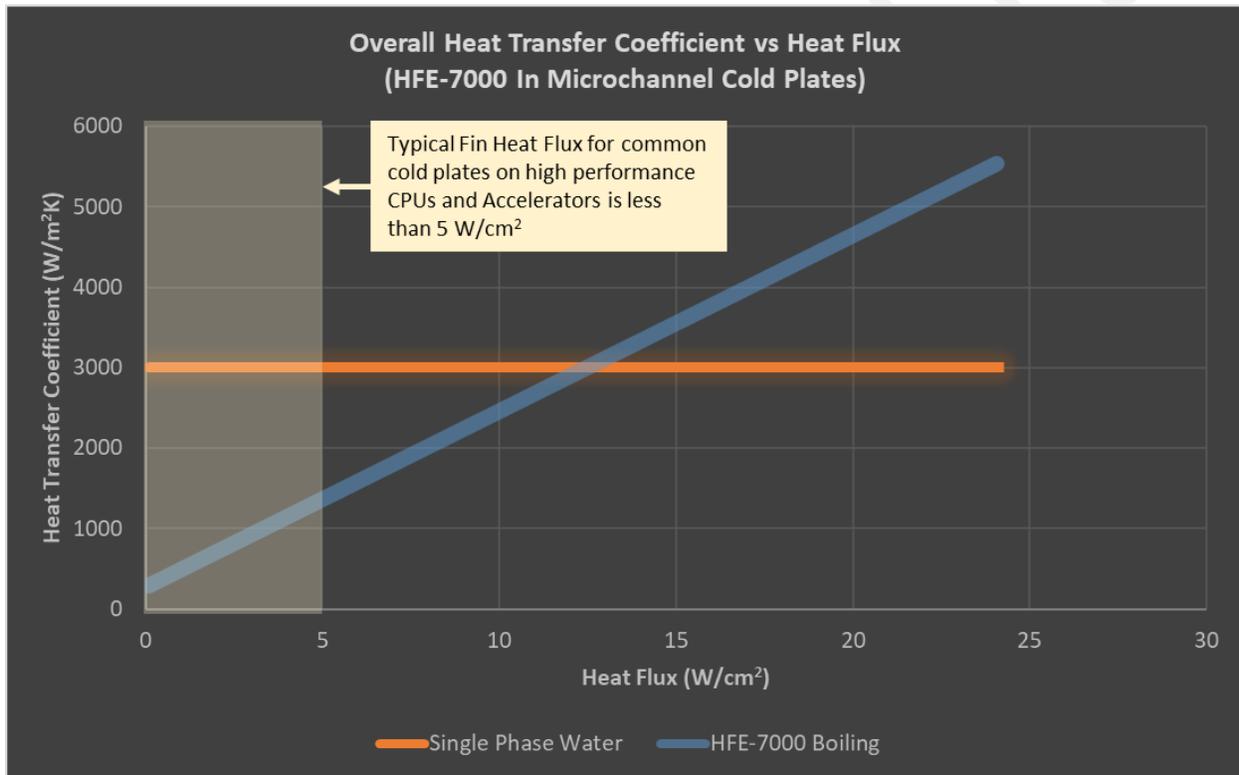


Figure 14. Overall or “effective” heat transfer coefficient of HFE-7000 and single-phase water versus cold plate fin heat flux – smooth copper microchannels.

The shaded region in Figure 14 highlights the typical heat flux for common server cold plate geometries using current generation CPUs. In fact, using one of the most popular server cold plates in production with a CPU TDP of 350 Watts results in just 1.4 W/cm² heat flux. Figure 15 flips the graph and shows how increasing the surface area of a cold plate results in a decaying heat flux and decaying heat transfer coefficient for a 350-Watt CPU load.

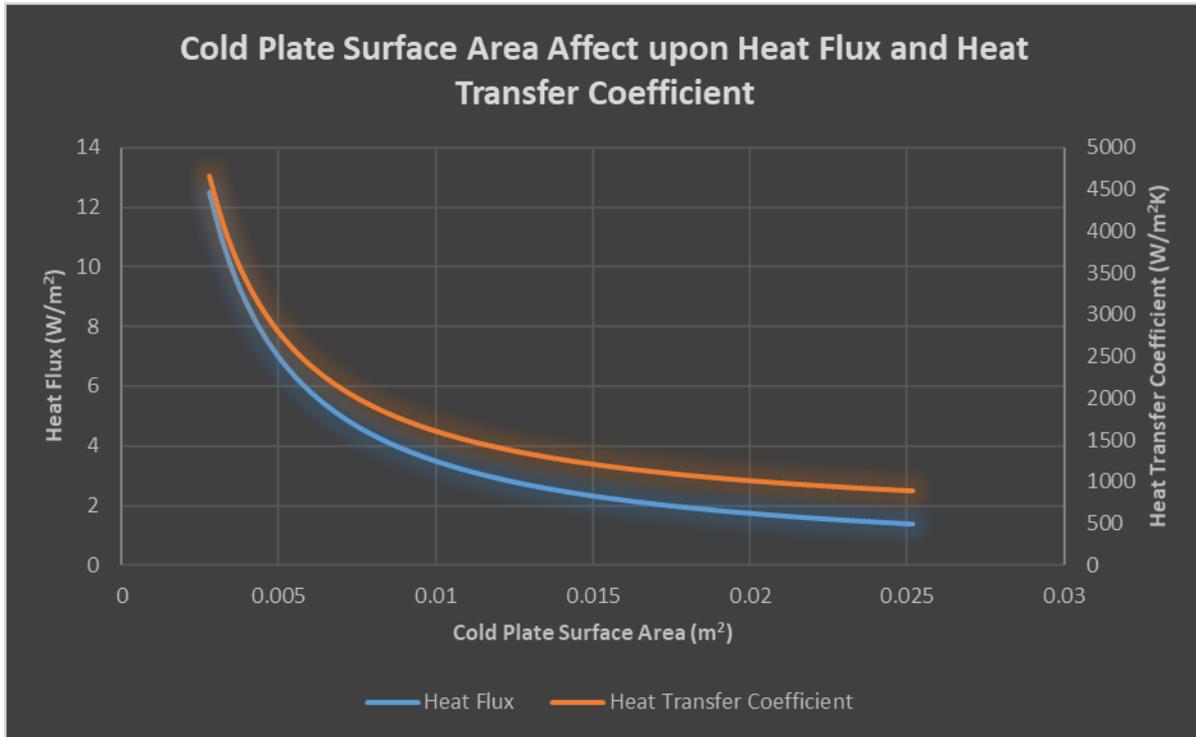


Figure 15. Affect of increasing the finned surface area of a cold plate on the heat flux and heat transfer coefficient for HFE-7000.

OK, so having too much effective heat transfer surface area works against the two-phase heat transfer coefficient, why not just scale back the surface area of the cold plate so heat flux goes up?

This is similar to our previous pressure drop versus flowrate conundrum in that both heat transfer coefficient and surface area are equally weighted in the calculation of the convection thermal resistance:

$$R_{convection} = \frac{1}{hA_{surface}}$$

Where $A_{surface}$ is the total surface area of the cold plate fins in contact with the coolant. In single-phase heat transfer, the heat transfer coefficient (h) is more or less independent of heat flux, so the objective is almost always to maximize the value of $A_{surface}$ to create the lowest possible value of $R_{convection}$. Two-phase liquid cooling changes this design paradigm to an extent.

So who wins: heat transfer coefficient or surface area?

Figure 16 below shows the results of increasing cold plate surface area on the cold plate's convection thermal resistance ($R_{convection}$) with the same 350-Watt CPU load. Increasing cold plate surface area, in this case, does still promote a reduced convection thermal resistance, but the effect of the diminishing heat transfer coefficient decays the rate of reduction in thermal resistance as compared to single-phase water.

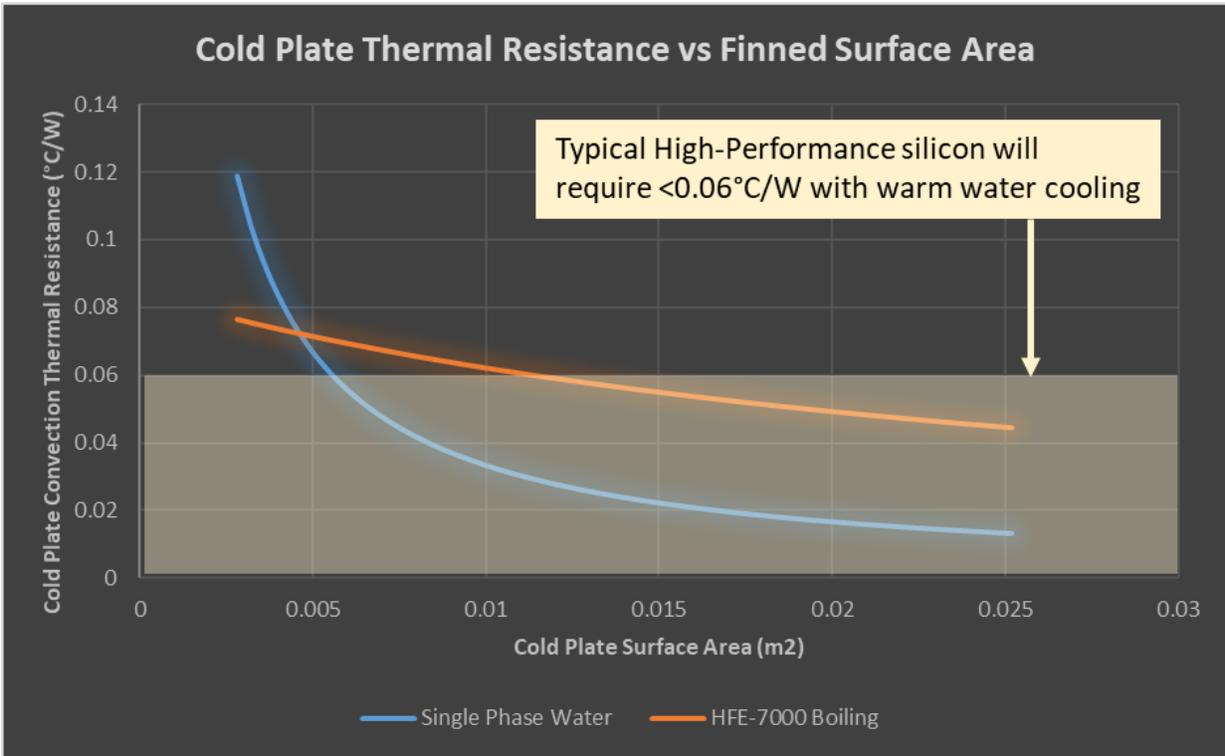


Figure 16. Cold plate convection thermal resistance as a function of finned surface area for a 350-Watt CPU load.

It was previously established that the convection thermal resistance of the cold plate into the cooling fluid is just one component in a series of thermal resistances that map the heat transfer path from an individual CPU to the end heat rejection site being external ambient air (unless your data center lives under the ocean – more on that another day). Assuming that the building facility cooling tower and cooling pumps comprising the facility water supply (FWS) would look the same for both single and two-phase server cooling, the next step is to examine the convection resistance impact in the greater context of the full thermal resistance network. Figures 17 and 18 below show the basic thermal resistance networks for 350-Watt CPUs cooled by single and two-phase liquid loops, respectively.

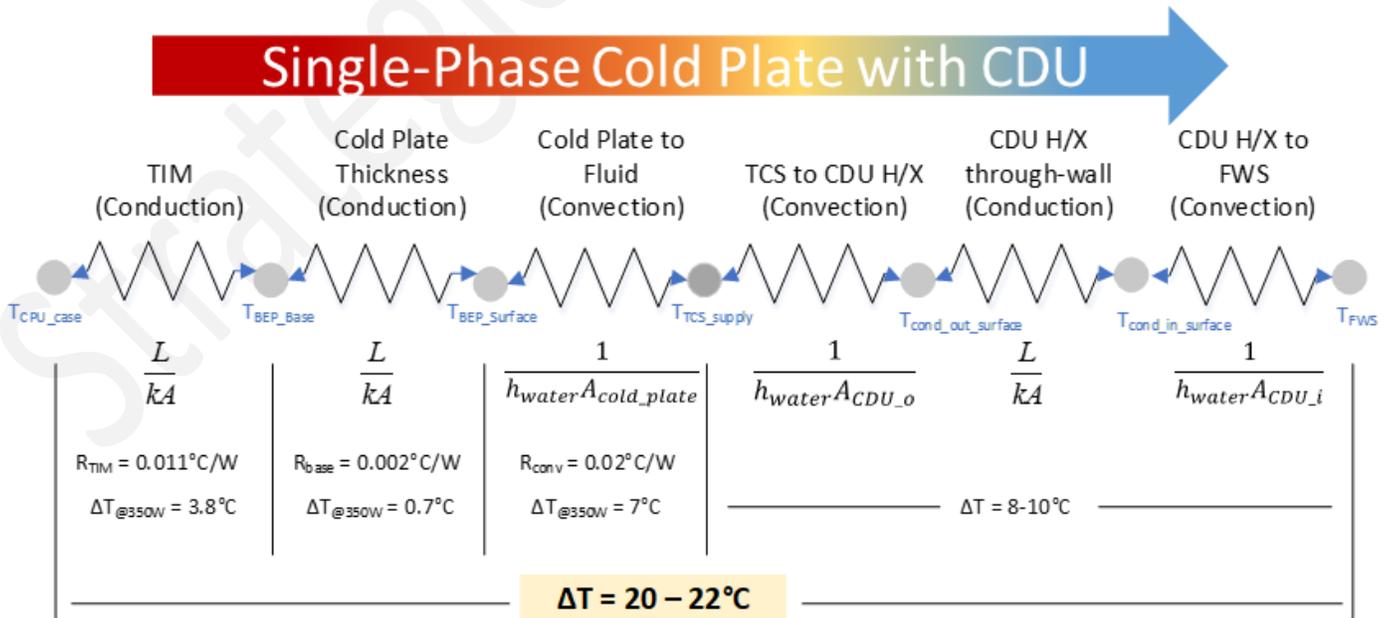


Figure 17. Thermal resistance network from CPU case to Facility Water Supply (FWS) for single-phase water cold plate liquid cooling.

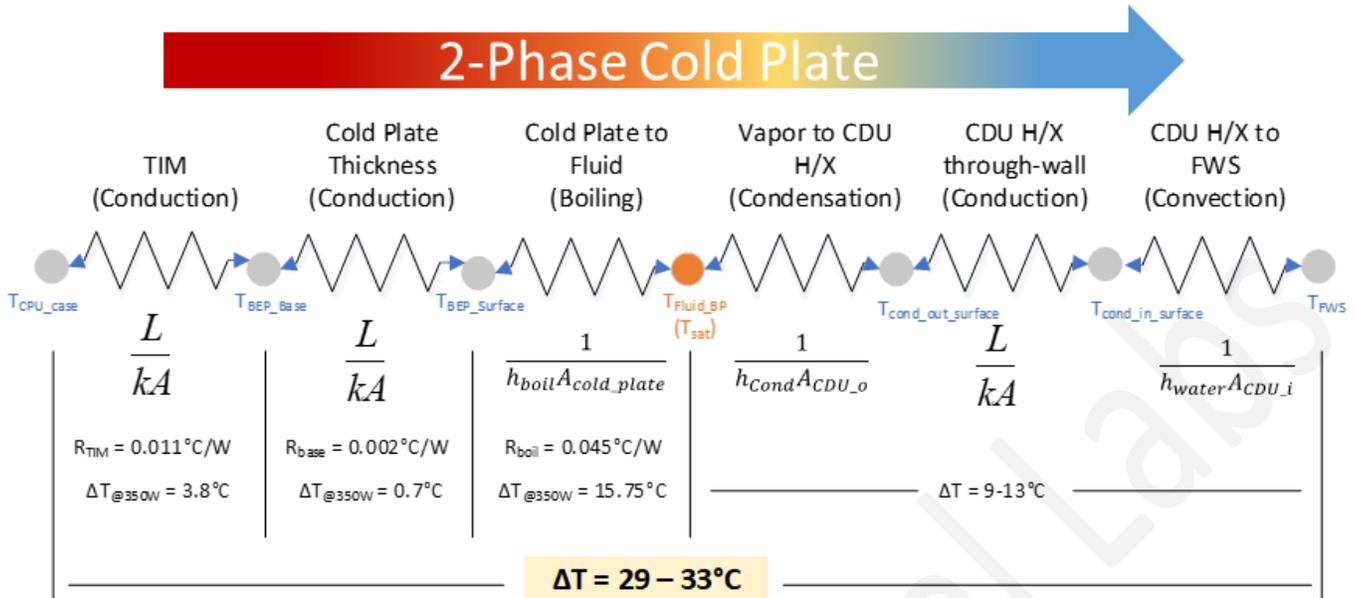


Figure 18. Thermal resistance network from CPU case to Facility Water Supply (FWS) for two-phase cold plate liquid cooling.

While much time has been devoted to discussing the boiling heat transfer coefficient (that’s the part everyone likes to talk about anyway), no time has been spent on the condensation heat transfer coefficient. In the essence of time and space it will suffice to say that condensation of organic fluids typically has a heat transfer coefficient in the 2000-2500 W/m²K range while single-phase water passing through a heat exchanger is more commonly in the 3000-4000 W/m²K range. Condensation heat transfer of organic fluids lags single phase water driving up larger Technology Coolant Supply (TCS) to Facility Water Supply (FWS) temperature gradients.

The only parts of the thermal resistance network that change going from single to two-phase cooling are the cold plate to fluid thermal resistance (boiling heat transfer coefficient) and the coolant to CDU thermal resistance (condensation heat transfer coefficient).

Comparing Figure 17 and Figure 18, the increase in total thermal resistance can cause a 350-Watt TDP CPU to be somewhere between 7 and 13°C hotter when cooled by two-phase cold plate liquid cooling than when cooled by a single-phase cold plate.

This difference in heat transfer performance can exclude two-phase cold plate solutions from using warm water cooling and/or limit the CPU TDP that can be supported. Table 3 below shows the 5 classes of water temperature defined by ASHRAE in their 2011 Liquid-Cooled Thermal Guidelines. Based upon this data, it would be challenging to cool a 350W TDP CPU with two-phase cold plate liquid cooling with ASHRAE water temperatures at W4 and beyond.

Table 3. 2011 ASHRAE Liquid-Cooled Thermal Guidelines

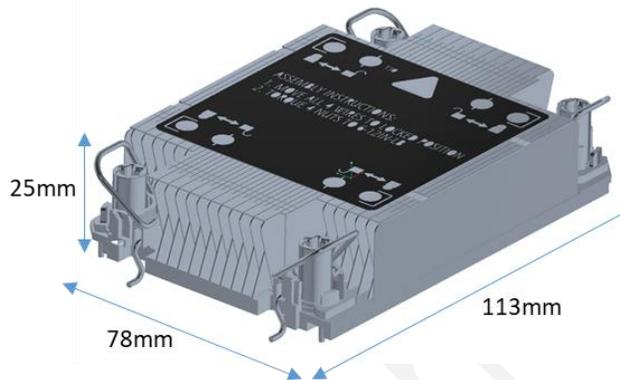
Classes	Typical Infrastructure Design		Facility Supply Water Temp (C)	IT Equipment Availability
	Main Cooling Equipment	Supplemental Cooling Equipment		
W1	Chiller/Cooling Tower	Water-side Economizer Chiller	2 – 17	Now available
W2			2 – 27	
W3	Cooling Tower	Chiller	2 – 32	Not generally available, dependent on future demand
W4	Water-side Economizer (with drycooler or cooling tower)	Nothing	2 – 45	
W5	Building Heating System	Cooling Tower	> 45	Specialized systems

A few things should be noted on the two-phase cold plate data:

1. The smooth-wall copper microchannel cold plates used to create this data should not be considered to be highly optimized for two-phase cooling. Cross-cutting skived fins and surface enhancement coatings like sintered copper powder and other nanostructures (to be discussed in more detail later) are all capable of dramatically enhancing the two-phase heat transfer coefficient and improving performance. At the time of this writing, these techniques are not yet mainstream and/or require expensive, secondary, post-manufacturing operations. They are also far less studied in literature for accurately estimating their performance benefits and tradeoffs.
2. Data collected assumed cold plate space limitations common to the current generation of server CPUs (Intel Ice Lake-SP and AMD Rome).
3. The CPU TDP assumption used to drive the heat flux for thermal resistance calculations was 350 Watts. While this value is substantially higher than mainstream server CPU TDP today, it is, conceivably, on the horizon for CPUs as well as being common in the discrete accelerator world. Accelerators with silicon TDP stretching towards 1000 Watts are in sight and these rising heat fluxes would further improve the heat transfer coefficient for two-phase cold plates.

Single-Phase Immersion Heat Transfer Performance

Single-Phase immersion makes a good showing when comparing overall cooling efficiency against the three other liquid technologies in this discussion. Where it falls behind, however, is in the comparison of heat transfer performance. Rather than cold plates with sub millimeter skived-fin microchannels, single-phase immersion cooling providers often utilize CPU heat sink geometries that look like air-cooled heat sinks. In fact, many use those air-cooled heat sinks as-is. This is because the natural convection heat transfer coefficient for single-phase oil is actually very similar to that of forced convection air – meaning similar amounts of surface area are required to achieve CPU thermal resistance targets. Figure 19 below illustrates the approximate air-cooled, heat sink geometry features for an Intel Ice Lake-SP 1U heat sink alongside a numerical optimization used to tune the fin geometry within that envelope for maximum heat dissipation in single-phase oil immersion.



1U Oil Immersion Heat Sink Performance Vs. Fin Spacing and Thickness

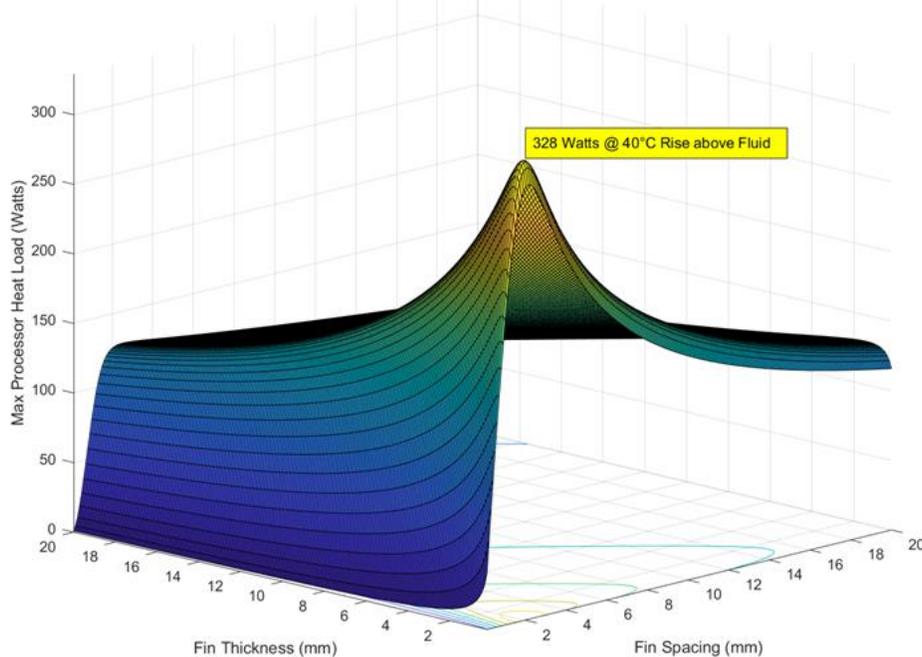


Figure 19. Geometric keep-out for a 1U Intel Ice Lake-SP heat sink shown alongside the numerical optimization for max CPU heat load.

As shown in the figure above, maximum CPU heat dissipation is approximately 328 Watts for a 40°C rise of CPU case temperature above fluid or a thermal resistance of 0.12°C/W. This is considerably worse than what is attainable through the other 3 liquid cooling technologies discussed in this paper.

Figure 19 shows that the thermal resistance from CPU to coolant in single-phase immersion is fairly large, however, this weakness is made worse by the fact heat must still transfer out of the oil coolant into the facility water supply via liquid-to-liquid heat exchanger in the CDU. Mineral oil generally having substantially less thermal conductivity and volumetric specific heat capacity than water means that CDU performance is also heavily impacted by the fluid choice. Figure 20 below shows a comparison of an identical liquid to liquid heat exchanger being used for water to water vs oil to water heat exchange. Simulating the performance with the same fluid flowrates and entering temperature conditions, the water to water heat exchanger is able to transfer over 2x more heat than the oil to water heat exchanger:

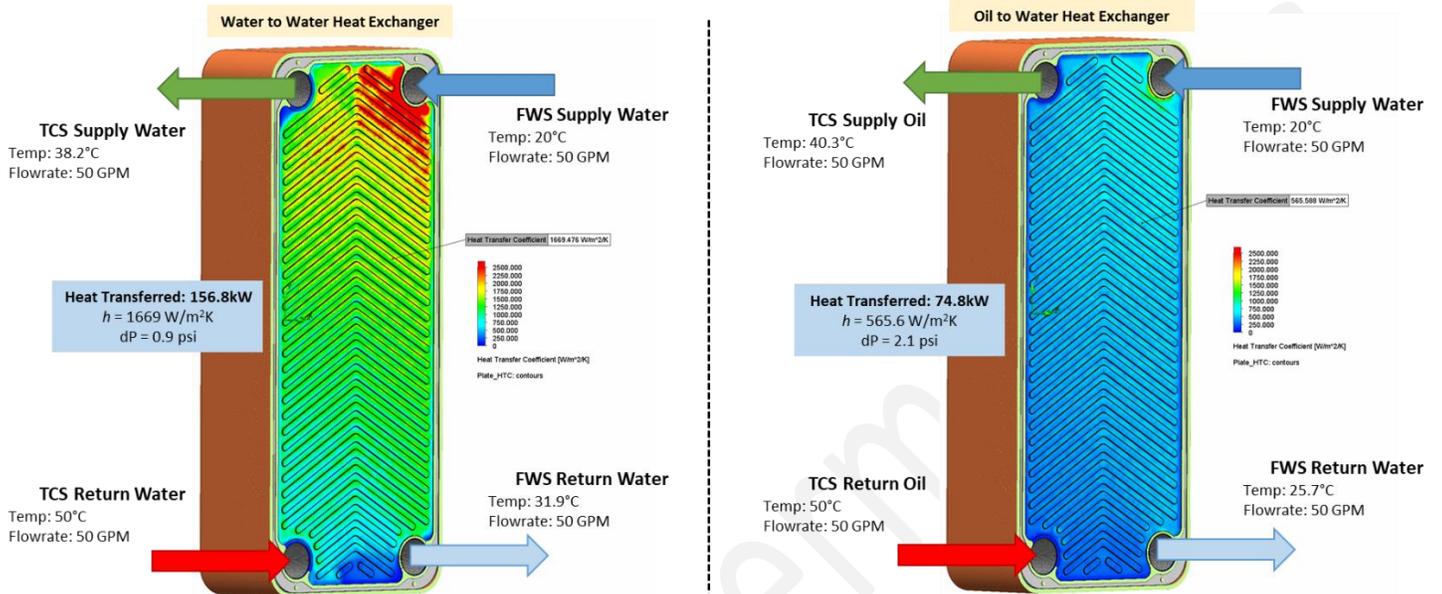


Figure 20. CFD simulation comparing heat transfer performance of a CDU heat exchanger for water vs mineral oil working fluid.

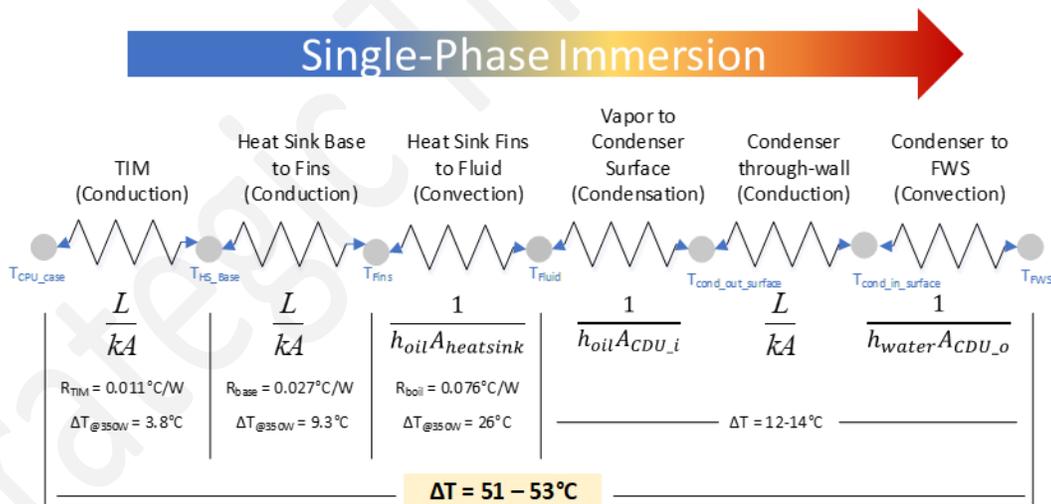


Figure 21. Thermal resistance network from CPU case to Facility Water Supply (FWS) for single-phase immersion liquid cooling

Single-phase immersion oils suffer from 3 primary weaknesses:

- High Viscosity consumes larger amounts of pumping power per unit of fluid
- Low specific heat and density compared with water
- Low fluid thermal conductivity compared with water

While there are a number of different immersion solution providers working directly with top chemical engineering firms to improve some of the thermophysical attributes of their proprietary oil blends, most of them still largely conform to the properties of common mineral oil.

Two-Phase Immersion Heat Transfer Performance

The final topic for heat transfer performance comparison is two-phase liquid immersion wherein the dielectric coolant comes into direct contact with the heated component and captures heat through “pool boiling” as opposed to the “convection boiling” observed in cold plates. It has already been established that because of the complete elimination of the CDU pumps and server fans that two-phase immersion has great energy efficiency potential, but how does it compare with the other two liquid cooling forms from a heat transfer perspective?

Because two-phase liquid heat transfer is so powerful in comparison with air cooling and because the immersion tanks, themselves, provide the fluid containment, it is possible, in many cases, to cool powerful server components without any heat spreader devices whatsoever using bare-die or bare-lid immersion cooling. For higher TDP components, a heat spreader is still required. Figure 22 shows a diagram similar to Figure 17 illustrating the heat transfer path out of CPU die into coolant fluid for two-phase immersion.

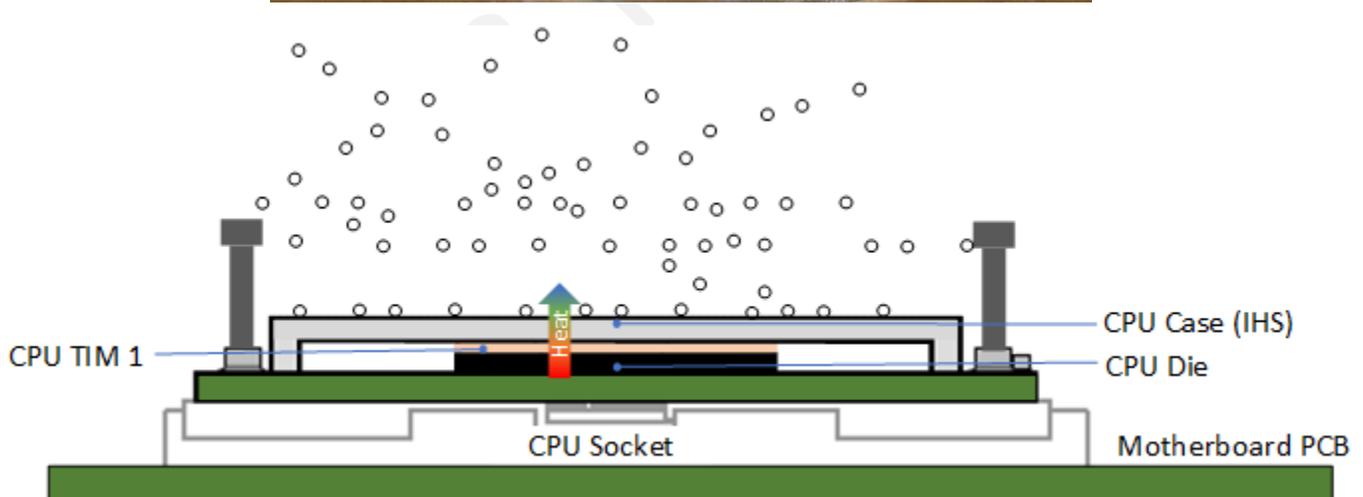
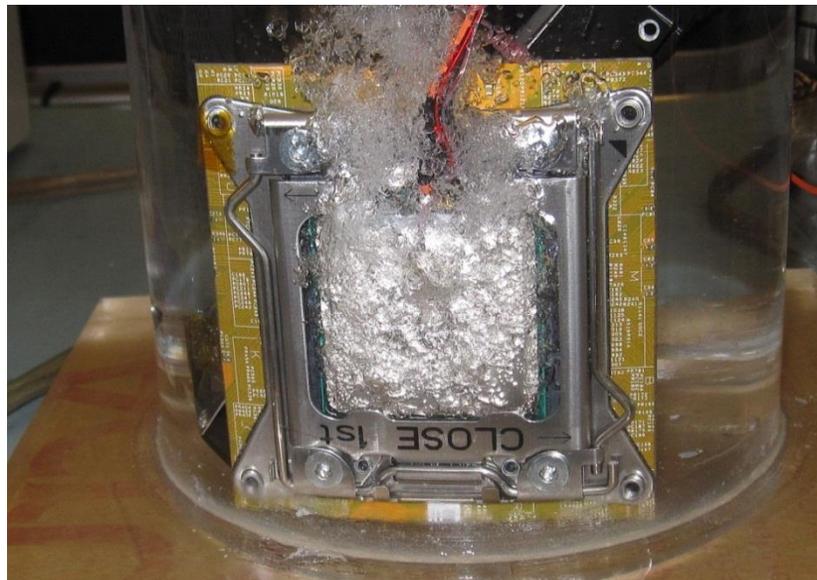


Figure 22. Diagram of the physical entities that map the heat transfer path from CPU die to cooling fluid through direct immersion – bare IHS.

One of the benefits that immersion’s “pool boiling” has over convection boiling in microchannels is that while vapor bubbles in convection microchannels grow and displace liquid contact at the heat transfer surface, in pool boiling, vapor bubbles immediately jettison from the heat transfer surface making room for fresh fluid to come in and create more bubbles. Furthermore, there is no sensible (single-phase) heating component in pool boiling whereas liquid pumped through a microchannel will, invariably, have a portion of the channel length where the fluid is heating sensibly before it begins to boil which robs some of the effective surface area.

For these reasons, pool boiling tends to transition into optimal boiling regimes at lower heat fluxes than forced convection boiling. The downside of pool boiling is that it also hits the *final* boiling regime known as “dry-out” or critical heat flux (CHF) at lower heat fluxes than forced convection boiling.

Boiling enhancement is also much more accessible through traditional manufacturing techniques in pool boiling than it is in microchannel cold plates. Sintered copper powder is commonly used in heat pipe manufacturing to create a high-porosity wicking surface for the internal working fluid as shown in Figure 23 below.



Figure 23. Sintered copper powder wicking structure from conventional heat pipe.

This type of microstructure has high thermal conductivity, has a high porosity (ideal for bubble formation), and bonds easily to copper plates used for CPU heat spreaders. Heat spreaders using this technology can leverage standard CPU retention mechanisms and be manufactured at low cost with conventional techniques as shown in figure 23 below.

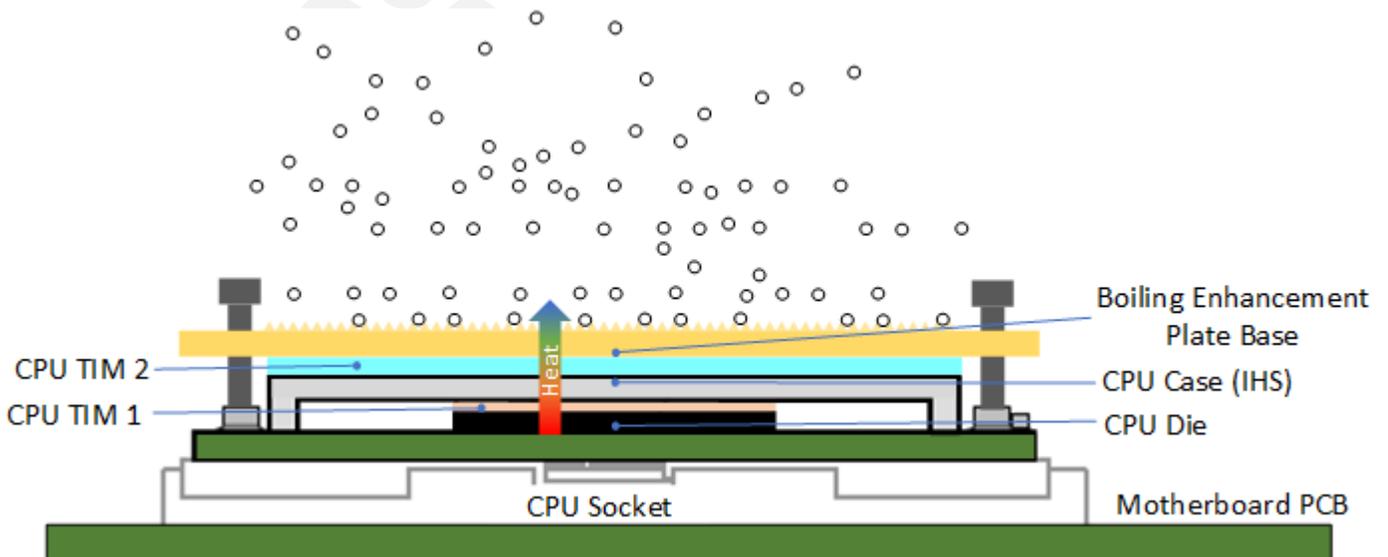
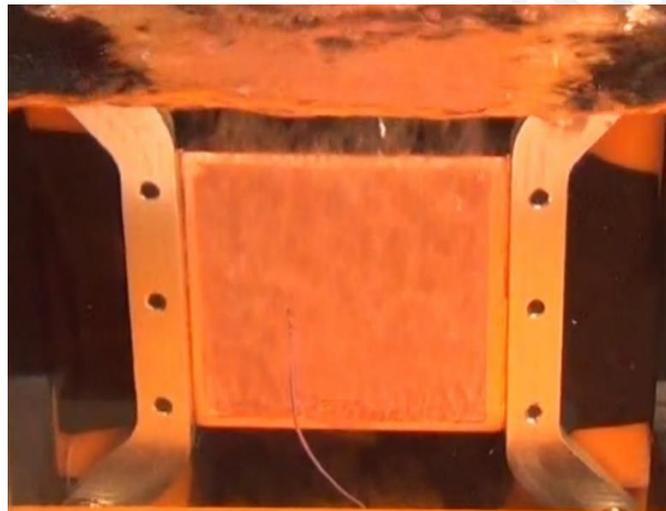


Figure 24. Boiling Enhancement Plate (BEP) utilizing sintered copper powder on copper plate affixed with standard LGA retention mechanism.

Figure 25 below shows empirical data comparing boiling heat transfer coefficient comparisons between a bare CPU Integrated Heat Spreader (IHS) as displayed in Figure 22 and the same CPU heater attached to the boiling enhancement plate depicted in Figure 24 with the 3M Novec 649 fluid.

The results in Figure 25 show that an enhanced copper surface can create heat transfer coefficients that are nearly 3x that of a smooth copper surface. Furthermore, the dipping trend in heat transfer coefficient observed at higher heat fluxes indicates an early transition towards film boiling reflecting immersion's reduced capacity for critical heat flux compared to forced convection.

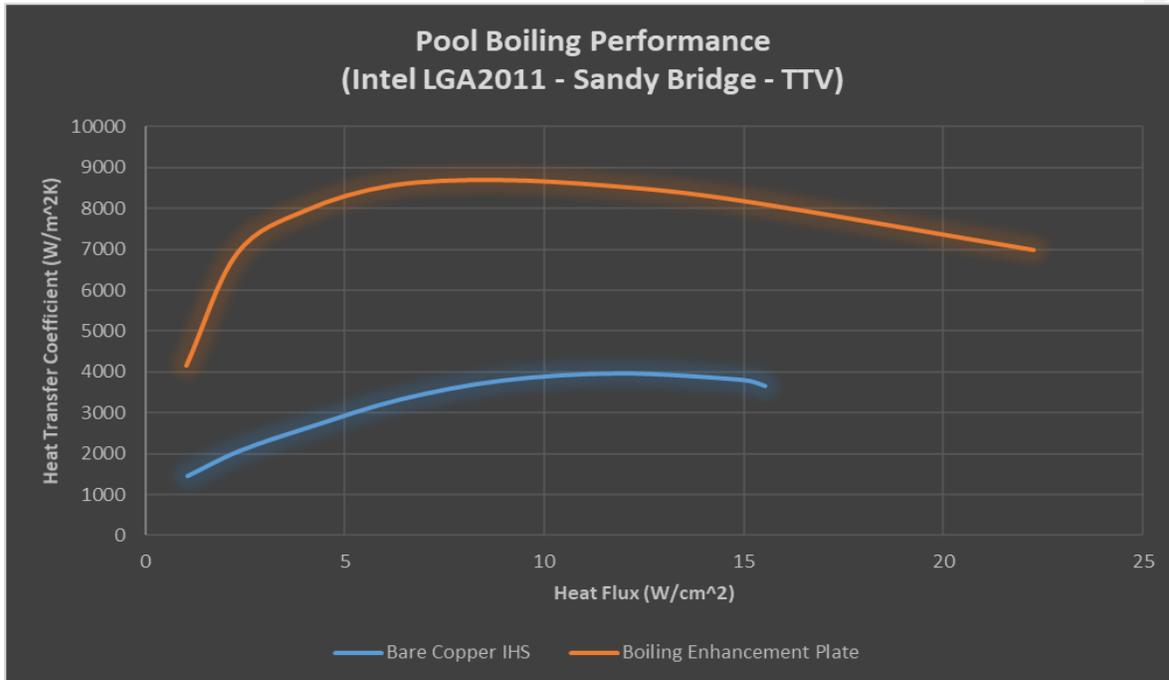


Figure 25. Comparison of effective heat transfer coefficient vs heat flux for bare copper and the boiling enhancement plate in 3M's Novec 649 fluid.

With the accessibility of high heat transfer coefficient-enabling boiling enhancement plates (BEPs), how then does two-phase immersion compare with the other two liquid forms discussed in this paper from an overall thermal resistance network perspective?

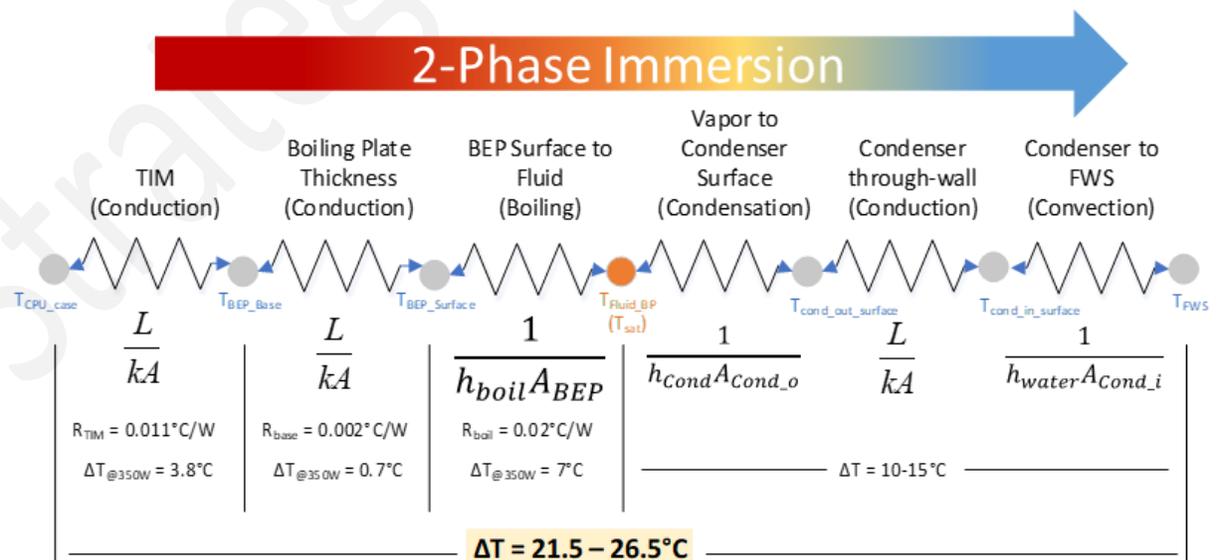


Figure 26. Thermal resistance network from CPU case to Facility Water Supply (FWS) for two-phase immersion liquid cooling

Figure 26 above shows the thermal resistance network from CPU to facility water supply (FWS) for two-phase immersion liquid cooling. Utilizing a boiling enhancement plate, two-phase immersion is good competition for single-phase water in net thermal resistance from CPU to liquid coolant. Practical sizing for a boiling enhancement plate that leverages a standard retention mechanism and moderate surface area improvements (pin fins) shows nearly identical performance with contemporary single-phase water cold plates.

Though the condensation heat transfer coefficient for organic fluids slightly lags that of single-phase water, what ultimately inhibits two-phase immersion from truly separating itself from single-phase water in total thermal resistance from CPU to FWS is the volumetric limits on condenser surface area in immersion tanks.

Condensers in two-phase immersion tanks most often rely upon surface-enhanced copper tubing as a rugged, economical, means to create usable heat transfer surface area as depicted in figure 27 below. The challenge is that the surface area density of copper tubing is relatively low. Furthermore, these tubes are usually welded in place to prevent vapor escape meaning that they need to be fully clear of the space directly overhead of the immersed servers to allow for those servers to be removed during a service event.

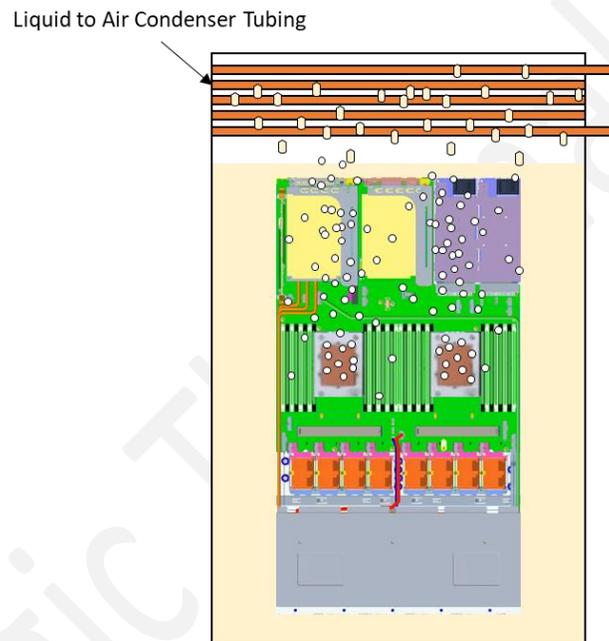


Figure 27. Two-phase immersion cooling diagram depicting the configuration and placement of condenser tubes.

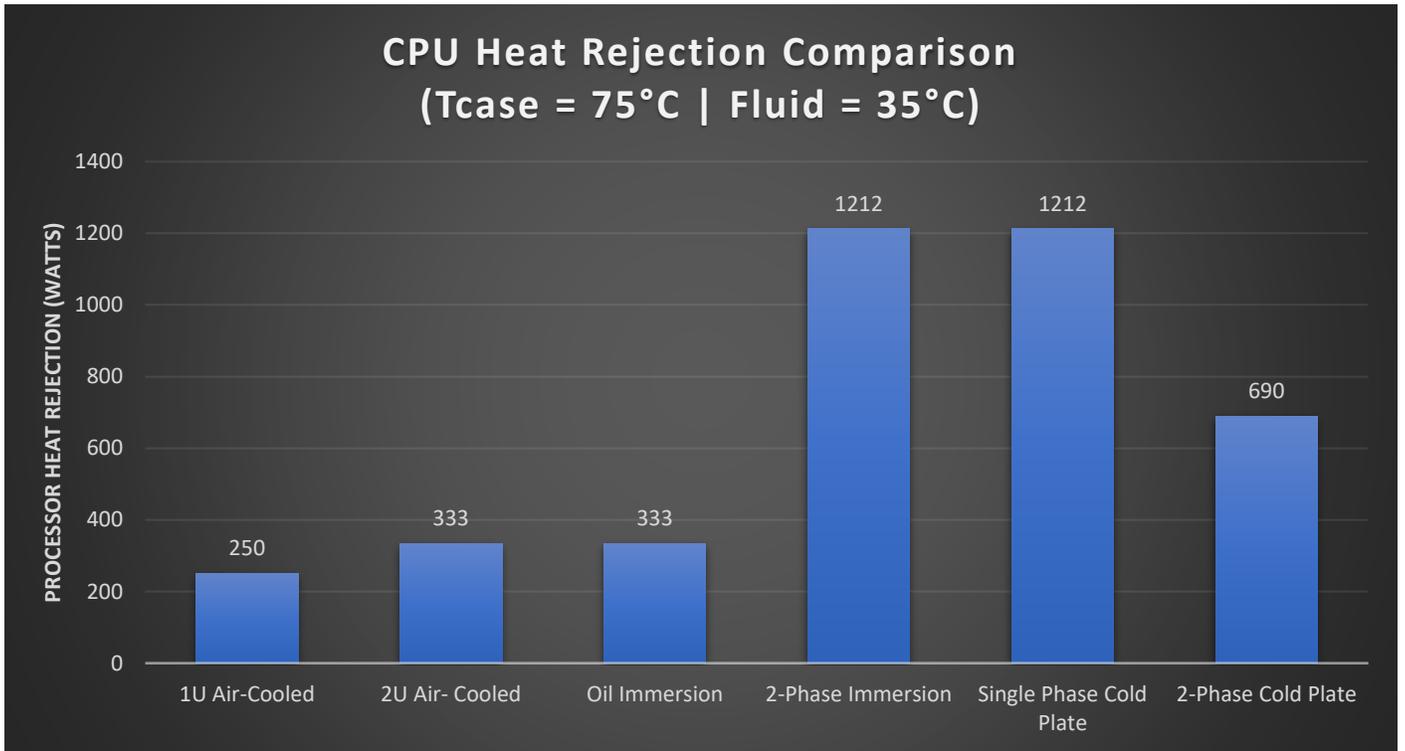


Figure 28. Cooling technology comparison for max CPU heat rejection capability

Conclusion on Heat Transfer Performance: Single-Phase cold plate using water ties with two-phase immersion cooling for maximum component heat rejection capability while two-phase cold plate lags slightly behind. Single-phase oil immersion does not compete for component heat transfer potential with the other three liquid technologies but is able to contend with air-cooling solutions at a similar volume.

Discussion on WUE

The term *Water Usage Effectiveness (WUE)* has seen elevated importance in data center power and efficiency discussions of late due to the explosive growth in data center water consumption. While WUE is beyond the scope of this whitepaper it can actually be addressed in fairly short order.

Even within the context of data center liquid cooling, “Water Usage” does not reflect the fluids (water, dielectric oil, or other) used for directly cooling server hardware as the fluids in these loops are effectively ‘static’ and not intended to be lost or replaced using water from a utility provider. Rather, “Water Usage” refers to the cooling water that is evaporated in a facility cooling tower to reject heat from the facility water supply (FWS). The amount of water required to be evaporated by the facility cooling tower is a function of the quantity of heat being rejected into the atmospheric air and the allowable temperature rise of the FWS above the ambient air temperature.

Assuming that the total amount of data center heat being rejected and the ambient air temperature are the same, the data center cooling design that can operate with the hottest temperature water will require the least amount of evaporated water in the cooling tower.

The cooling solution that can create the smallest temperature difference between critical components (such as CPUs) and the facility water supply will allow for the hottest facility water supply temperature.

Therefore, the cooling technology with the lowest overall thermal resistance will have the best Water Usage Effectiveness (WUE); in this case, single-phase water followed up closely by two-phase immersion.

Strategic Thermal Labs

Brief discussion on two-phase liquid-cooling challenges

While this author is, personally, well-traveled in the complexities associated with implementing two-phase cooling on server products we will not belabor that subject in this paper. Highly creative, problem solving engineers working at various intersection points of the data center ecosystem are continuously unveiling innovative solutions to address these challenges and make two-phase cooling more accessible. Having said that, two-phase cooling is, by no means, trivial from an implementation perspective. There are significant obstacles. Some do not have pleasant solutions, others don't, yet, have solutions at all.

Here is a quick summary of some of the key challenges with two-phase server liquid cooling:

- Proprietary dielectric fluids used in two-phase cooling have a small supply base and are extremely expensive.
 - This is a bigger challenge with immersion simply because it is difficult to prevent air voids in hardware that needless waste expensive liquid volume.
- While these fluids are expensive, they are even more difficult to contain.
 - Having high vapor pressures and low surface tension, the fluid will wick through incredibly small openings at high rates (think the micro gaps between wire insulation and embedded conductors) and extreme care must be taken to create near-hermetic seals around every fluid-carrying vessel.
 - Most rubber or elastomeric materials commonly used for O-rings and gaskets are incompatible with these fluids. They have strong solvent properties and will dissolve loose hydrocarbons (oils) from seals leaving voids for the vapor to egress through.
 - There is virtually no supply base for hermetically sealed connector types commonly used to carry high speed signals in the data center (SFP, QSFP, LC/SC/MPO Fiber, etc.) and these require custom potting operations to prevent significant fluid loss over time.
- Finally, and specifically for immersion, there is not an 'ideal' fluid chemistry, today. Each of the most infamous dielectric fluids used for two-phase immersion cooling have unique, substantial, draw backs.
 - HFE fluids have a very high dielectric constant causing signal integrity problems for high-speed signals common in servers.
 - FKs (Novec 649) is chemically incompatible with water and moisture within the system (either liberated by immersed plastics or brought in through vapor exchange with moist external air in the head space during service events can create a weak acid that will destroy hardware unless cleansed.
 - FC-72 (Fluorinert) has fantastic chemical properties 'within' the immersion tank but has an extremely high atmospheric lifetime and accompanying Global Warming Potential (GWP). Fluorinert has a GWP nearly 5000 times higher than CO₂ and is similar to that of R-134a, the refrigerant most often used in automobile air conditioners requiring strict handling procedures by licensed technicians. Saving carbon consumed in the generation of power by utilizing a higher efficiency cooling solution is fantastic unless those environmental savings are all but washed away by the accidental release of this powerful greenhouse gas through leaks or mishandling.

Nothing on this list is unsolvable. If it were easy, everyone would be doing it.

Final Thoughts

What was intended to be a concise review of the pros and cons of two-phase liquid cooling in servers compared to traditional single-phase technologies turned into something much longer but, hopefully it was valuable just the same.

The primary benefit of two-phase cold plate liquid cooling is the complete elimination of risk associated with leaking coolant on valuable electronics. As the price of silicon components continues to skyrocket (I'm looking at you, Intel, with some of your Xeon-SP CPUs listing at over \$13k per unit!) the cost-benefit analysis shifts towards an even greater need to ensure that a server's cooling solution does no harm to its precious commodities. As expensive as server components are becoming, they pale in comparison to the value of the workload being executed on top of them which further incentivizes damage-free cooling solutions. Having said that, two-phase cold plate cooling solutions lag single-phase water-based solutions significantly in heat transfer performance today and have a high degree of material complexity required to ensure that expensive fluids are not lost during operation. The reduced thermal performance may put a lower cap on component TDP that can be supported and/or FWS temperature limits.

Two-phase immersion could be considered somewhat of a "holy grail" for server cooling. It provides similar levels of heat transfer performance as single-phase water, innately allows for 100% capture of server heat, and an end-to-end efficiency that solutions relying upon server-side pumps and fans will not be able to match. The downside of two-phase immersion, of course, is that it has the highest degree of complexity with respect to mitigating ongoing vapor loss and addressing the aforementioned fluid chemistry limitations. Resolving these complexities for two-phase immersion is, today, a very expensive endeavor and TCO models continue to struggle against alternative cooling options. With that said, heavily resourced R&D teams at virtually all global hyperscale customers have two-phase immersion somewhere on their radar. The more of these high-performance teams that pick up, iterate, and cost reduce the infrastructure around two-phase immersion, the easier it will be for the rest of the market to reap the benefit.

Single-phase immersion carries the benefits of a non-conductive liquid coolant but at heat transfer performance levels that are similar to air, albeit at higher efficiencies. The greatest weakness of single-phase immersion will be its struggle to support high-power silicon devices on the horizon while also supporting chiller-free operation in many geographical locations. The greatest strength of single-phase immersion is actually its TCO structure which was not discussed in this paper. Look for that to be discussed in a future publication.

While single-phase water solutions carry an inherent risk of equipment damage in the event of a leak, the amount of servers equipped with water cooling in the field has exploded over the past 5 years and, yet, very little empirical data supports the likelihood of such leaks occurring. Single-phase water is also capable of setting the bar for best-in-class heat transfer performance, has a massive supply base of compatible equipment spanning a large plurality of industries, and has a relatively low complexity barrier.

In conclusion, each form of liquid cooling is ripe with unique strengths and weaknesses. As a thermal engineer, it is extremely exciting to see the day finally come where silicon power increases are finally starting to require new, creativity, cooling solutions. I do expect to see significant progress on two-phase cold plates in the coming months/years to create surface geometries that close the gap with the other two liquid formats. It will be interesting to watch how various segments across the data center space navigate the tradeoffs between the various liquid options, even as many end users are still desperately clinging to air-cooling for as long as they can.

Thank you for reading! If you stayed with me for this entire discussion, I sincerely applaud your commitment and implore you to engage in some additional dialogue on the subject. Shoot me an email or start up a technical discussion in the comment section!

Stay tuned for our next publication on objectively exploring the limits of air-cooling in the data center!

Information about this Paper and its Author

Author

Austin M. Shelnett, P.E., President at Strategic Thermal Labs, LLC

Inquiries

Please contact Strategic Thermal Labs, LLC at info@strategicthermal.com if you would like to discuss this report or its contents.

Citations

This paper can be cited by accredited press, analysts, corporate marketing, and other publications but must be cited in-context, displaying author's name, title, and "Strategic Thermal Labs".

Disclosures

Strategic Thermal Labs provides independent research, analysis, design, and test services to many high-tech companies mentioned in this paper. No employees at Strategic Thermal Labs hold any equity positions with any companies or technology solution providers cited in this document as of the date of initial publication.

Disclaimer

This document consists of only the findings of Strategic Thermal Labs based on independent research and study and should not be construed as statements of fact or absolute conclusion. The findings expressed herein are subject to change without notice to any party.

Strategic Thermal Labs makes all reasonable efforts to obtain and present accurate information as presented in this document (the "Data"); however, Strategic Thermal Labs does not endorse or approve the Data and does not guarantee the accurateness or completeness of the Data. Additionally, the Data presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors.

STRATEGIC THERMAL LABS ASSUMES NO RESPONSIBILITY FOR, AND EXPRESSLY DISCLAIMS ANY LIABILITY FOR, ANY CONSEQUENCES RESULTING FROM THE DISTRIBUTION OR USE OF THE DATA. UNDER NO CIRCUMSTANCES SHALL STRATEGIC THERMAL LABS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, CONSEQUENTIAL, OR SPECIAL DAMAGES ARISING OUT OF OR IN CONNECTION WITH THE DATA. STRATEGIC THERMAL LABS MAKES NO WARRANTIES (EXPRESS OR IMPLIED) RELATING TO THE DATA.