Case Study

Immersion Cooling | **High-Density Compute Workloads**

Single-Phase Immersion Cooling: The Path to 1000W TDP & Beyond

Type of Solution

Forced Convection Heat Sink (FCHS) package for single-phase immersion cooling

Availability

Worldwide

Highlight

The FCHS package provides a simple, low-power, and cost-effective solution for cooling chips up to the ~1000W+ TDP range in single-phase immersion cooling platforms.

Industry

All, with a focus on high-density compute workloads

Introduction

Liquid cooling has been gaining prominence, driven by the increased need for more performance resulting in the step-up of Thermal Design Power (TDP) of silicon components. This shift is spurred by the imperative to enhance not only Data Center (DC) efficiency and sustainability metrics but also the ability to contend with fluctuating ambient temperatures in edge environments. Immersion cooling effectively addresses all three of these pivotal concerns. It not only delivers superior cooling performance but also reduces infrastructure requirements, all the while providing complete insulation for IT equipment from the surrounding environment. In particular, single-phase immersion cooling, where the cooling fluid remains in its liquid state (unlike two-phase immersion cooling, where it undergoes a phase transition from liquid to gas), offers the cooling advantages needed to demonstrate sustained effectiveness across a spectrum of thermal resistance and Thermal Design Power (TDP) values. This makes it an ideal choice to foster widespread adoption.

To push the boundaries of cooling capability, there is a pressing need for a breakthrough that will propel single-phase immersion cooling beyond the 1000W threshold and expand its cooling capabilities to meet the evolving demands of modern technology.

Through innovative engineering and science from a collaboration between Submer and Intel, single-phase immersion cooling will now be a competitive cooling solution for higher TDP systems.

By harnessing the power of **forced convection**, Submer's R&D teams have demonstrated that single-phase immersion cooling can **achieve thermal resistances that rival those of Direct Liquid Cooling (DLC)**. Moreover, it offers the significant advantage of **reducing the quantity and cost** of components required for comprehensive heat capture and dissipation when compared to a DLC solution at elevated TDP levels.

The introduction of the **Forced Convection Heat Sink** (FCHS) marks the initial milestone on the path toward **effectively cooling Intel silicon exceeding 1000W.**

Moving From Natural to Forced Convection

Natural convection stands as the primary mechanism governing heat transfer within an immersion tank. However, significantly increasing flow rates at the point of injection into the tank yields a marginal enhancement with little impact on CPU and GPU temperatures. The effect of flow rates on the case temperature of the CPU (modelled via Intel TTV) is shown in *Figure 1*.

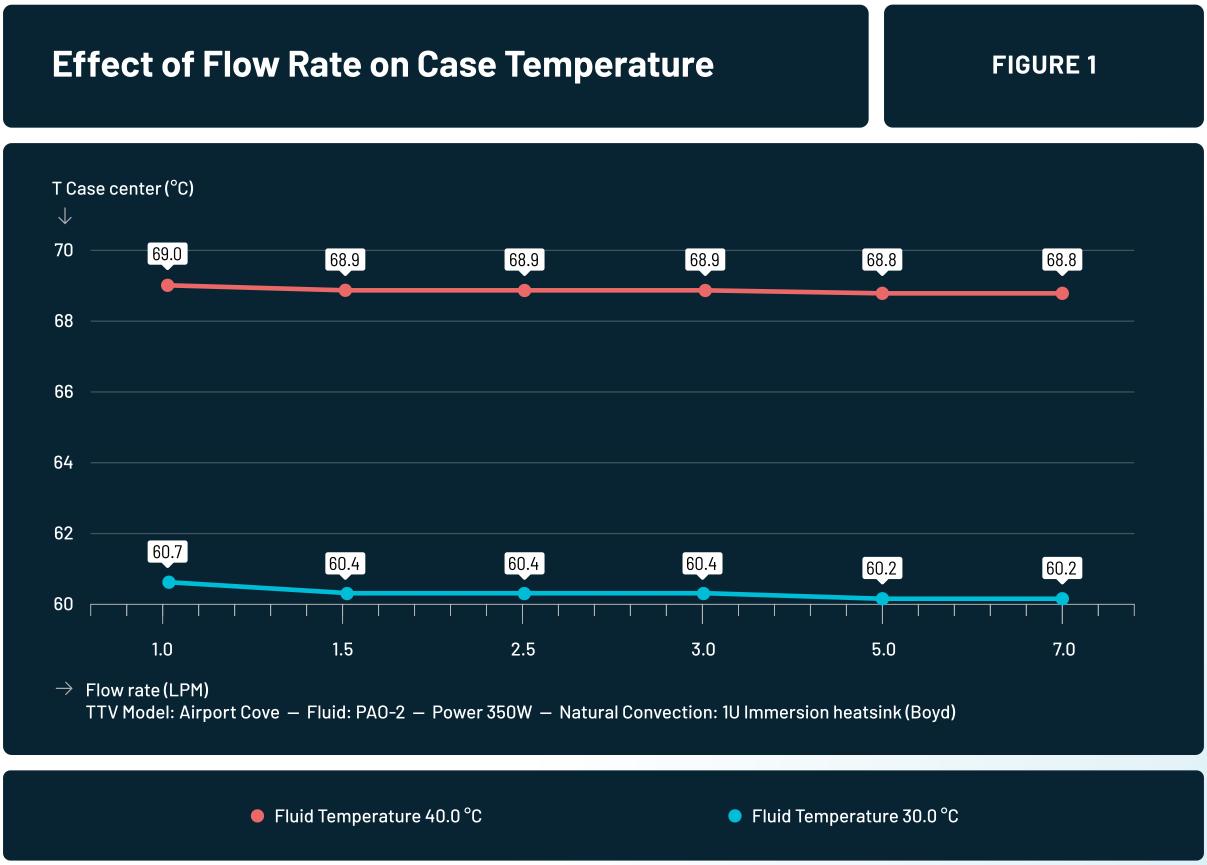

Figure 1: Effect of flow rate on Case Temperature

The following can be interpreted from *Figure 1*:

- Server level flow rate had **little effect on TTV temperature/theta**
- Fluid takes the path of **least resistance** i.e., avoids going between the fins of the heatsink
- True for both 1U air heatsink and immersion-optimized heatsink

When coupled with immersion-optimized heat sinks, natural convection has proved the ability to cool CPUs in the 400W power range. Nevertheless, as TDP levels surpass this threshold, novel methods for enhancing the flow of dielectric fluid through the heat sinks become imperative.

Over the past year, Submer has been diligently striving to comprehend the implications and to devise cost-effective means to achieve this acceleration, with the ultimate goal of supporting TDPs exceeding 1000W. The initial milestone in this journey is the development of the Forced Convection Heat Sink (FCHS).

The Forced Convection Heat Sink (FCHS)

The Forced Convection Heat Sink (FCHS) is an **innovative heat sink attachment** designed for use in immersion cooling systems. It combines the heat transfer efficiency of forced convection with the resiliency of passive cooling. The FCHS attachment for immersion-optimized heat sinks integrates propellers for localized forced convection as illustrated in *Figure 2*.

Figure 2: Forced Convection Heat Sink

The liquid propellers are designed for dielectric cooling fluids, featuring elongated blades with fewer blades than typical propellers. When inactive, the propeller allows fluid to flow smoothly through the heat sink for passive cooling, either due to mechanical failure or intentional shutdown.

Unlike traditional forced convection systems, the FCHS maintains consistent fluid flow direction in both active and passive modes, ensuring uninterrupted cooling. It can be equipped with redundant external convection sources, such as two or more liquid propellers.

Furthermore, the liquid propeller is easily replaceable without the need to detach the heat sink from the hardware. Additional components such as baffles or static fluid directing devices can further improve cooling, particularly in passive mode. This design **enhances cooling efficiency and reliability** in immersion cooling systems.

The main benefits of the FCHS are as follows:

- FCHS is easy to retrofit into any existing server or immersion tank design.
- FCHS enables high TDP CPUs/GPUs in single-phase immersion.
- FCHS is cheap to manufacture and even 3D print.
- The easy-to-swap, fail-proof design still enables natural convection capability in case of propeller failure.
- The FCHS delivers thermal resistances that are competitive with water-based Direct Liquid Cooling (DLC).
- If desired, FCHS can return thermal management to the server by enabling BIOS PWM control.
- FCHS is freely available for integration by fan and server ODMs in their server product offerings.

The Testing Process

Intel® Airport Cove Thermal Test Vehicle (TTV) was used to test the FCHS, as shown in *Figure 3*.

Both the TTV surface and the base of the heat sink were cleaned with isopropyl alcohol. Once dry, a 2" x 2.5" sheet of 0.006" thick indium foil was centered on the TTV.

Each heat sink was then securely mounted by gradually tightening the retaining nuts in a criss-cross sequence to ensure the even application of heat sink pressure on the TIM and TTV.

To power the TTVs, an external 0-60VDC power supply was utilized. The heatsink fans (propellers) were powered by a separate external DC power supply. A gear flow meter was used to measure the flow rate of the bulk dielectric.

To measure TTV temperatures, 3 T-type thermocouples were embedded in the Airport Cove Integrated Heat Spreader (IHS). K-type thermocouples were used to measure the inlet and outlet dielectric fluid temperatures.

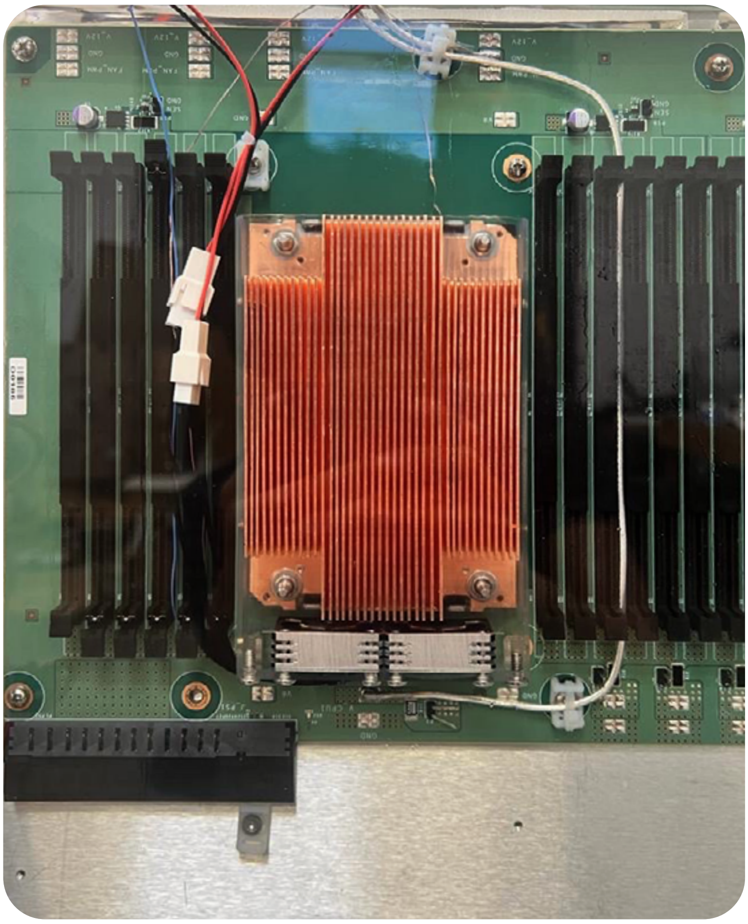

Figure 3: FCHS cooling Intel(r) Airport Cove Thermal Test Vehicle (TTV)

All FCHS packages were tested at 30°C and 40°C inlet fluid temperatures. The fans were supplied with a fixed 8VDC, and their total power consumption per heat sink ranged between 17.8W-18.6W, which varied with the viscosity of the dielectric fluid; at higher dielectric temperatures, the fans consumed less power. A power sweep from 100W to 1000W on the TTVs was executed as part of the test plan for each fluid temperature. Testing was executed both at Submers' R&D facilities located in Barcelona, Spain and Sausalito, California by independent teams to validate performance data.

Each test was allowed to run until temperature and power readings indicated a steady-state condition, at which point data was recorded.

To date, a total of 6 different heat sinks have been tested as part of the FCHS package:

- 2x skived copper heat sinks of differing fin thickness and pitch
- 2x vapor chamber heat sinks with brazed Ni-plated copper fins of differing thickness and pitch
- 1x copper pin-fin heat sink

The Results

TTV Temperature for Different Configurations from 100W to 1000W

FIGURE 4

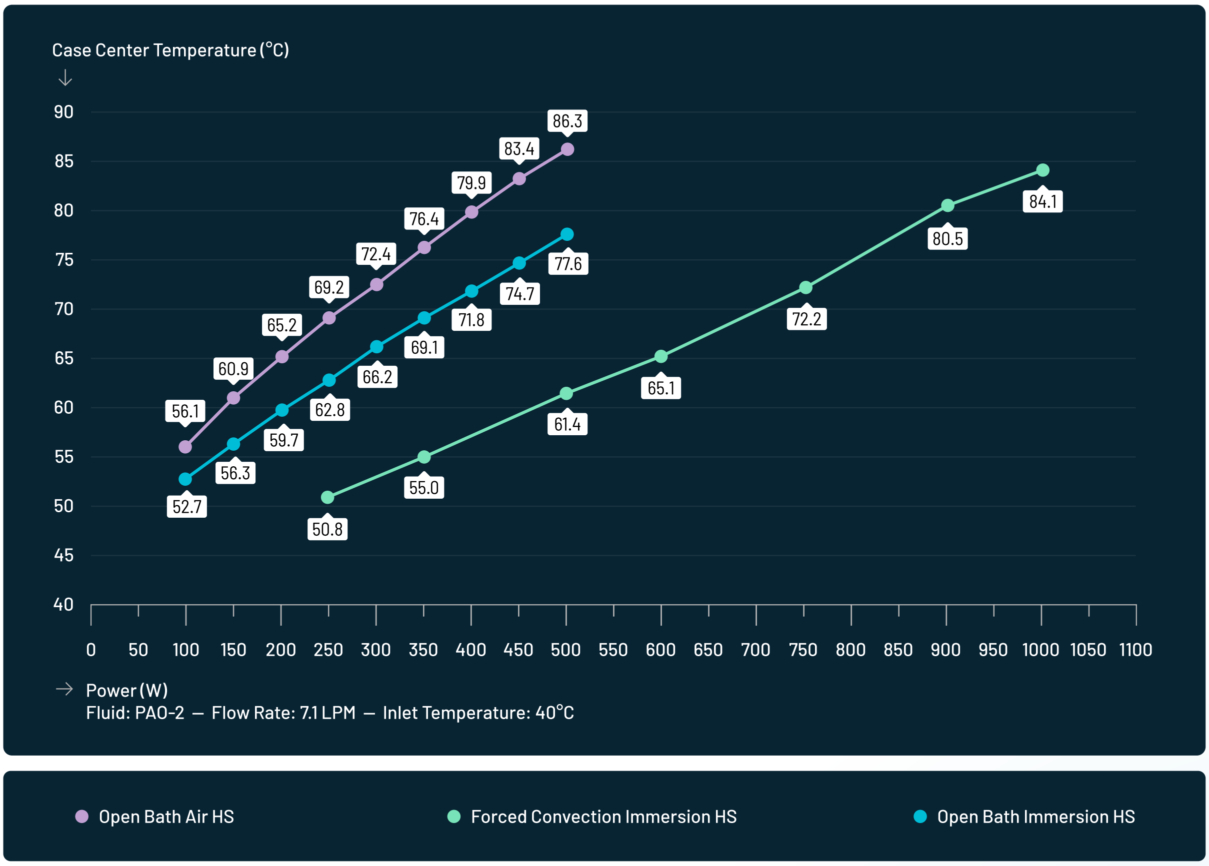

Figure 4: Case Temperature vs TDP

TTV Thermal Resistance for Different Configurations from 100W to 1000W

FIGURE 5

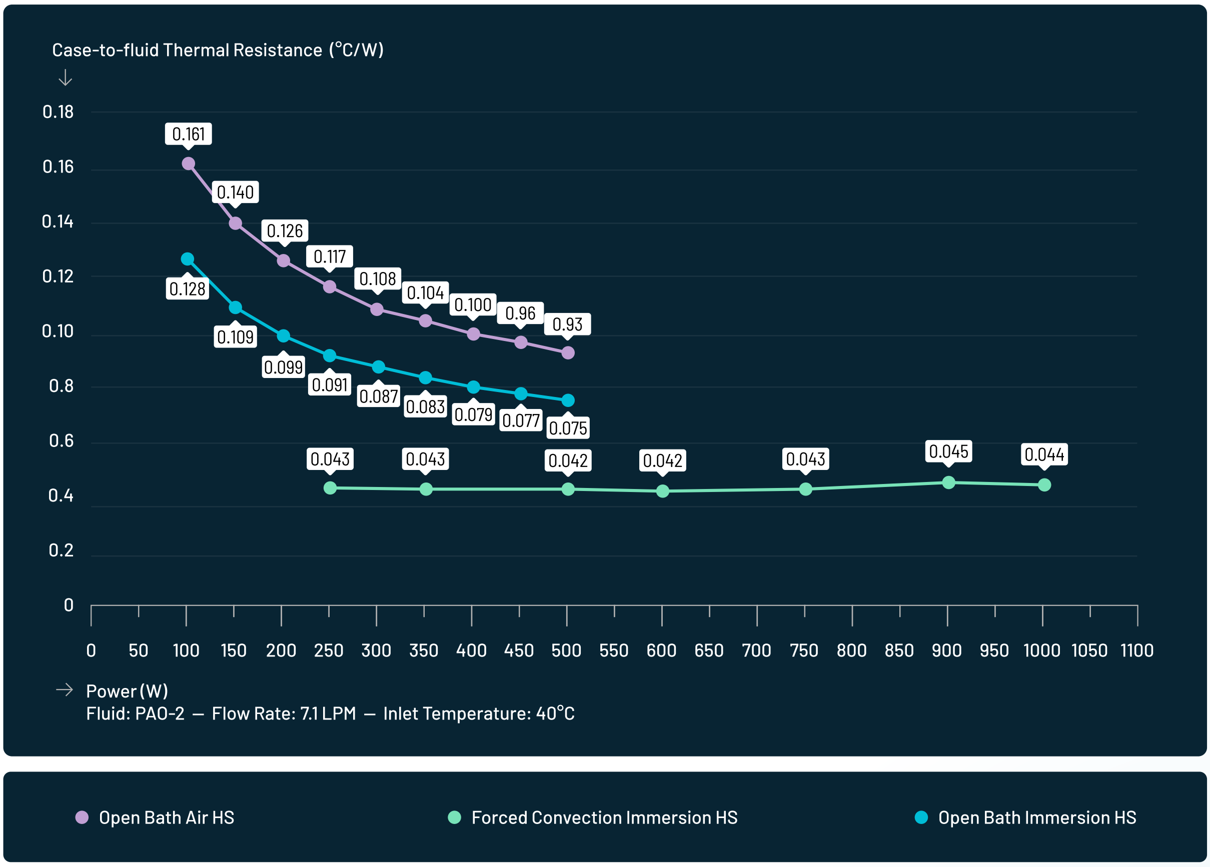

Figure 5: Case-to-fluid Thermal Resistance vs. TDP

The above charts (see Figure 4, Figure 5) compare Case-to-Fluid temperatures and Thermal Resistance for a power sweep going from 100W to 1000W on Intel Airport Cove Thermal Test Vehicle (TTV) using:

- Non-optimized air heat sinks
- Immersion-optimized heat sinks
- Forced Convection Heat Sink (FCHS)
- There is still room for substantial heat sink and FCHS optimization.
  Improved case-to-fluid temperatures and thetas of <0.04 are well within reach.

Real-World Cooling of an Overclocked Intel® Xeon® Processor

To validate the FCHS's ability to cool an actual CPU, it was paired with a heavily overclocked Intel® Xeon® w9-3495X workstation processor and tested at full load using Prime95 CPU stress testing software. With an inlet dielectric fluid temperature of 27°C, the steady state CPU power fluctuated around 800 Watts and the average DTS temperature hovered around 88°C, resulting in a junction-to-fluid theta of 0.076 °C/W.

Figure 6 FCHS cooling overclocked Intel® Xeon® w9-3495X

The Roadmap Toward the 1000W+ TDP

The journey towards achieving a 1000W+ TDP represents a formidable technological quest. As the demands on computing hardware continue to surge, fueled by innovations in artificial intelligence, data analytics, and high-performance computing, the need for cooling solutions capable of handling such colossal power densities becomes increasingly urgent. It involves the convergence of cutting-edge materials, advanced fluid dynamics, and innovative engineering strategies.

Submer and Intel are dedicated to shattering the current limitations that constrain high (TDP) thresholds in single-phase immersion cooling, thereby extending the advantages of immersion to address the challenges faced by a broader spectrum of users.

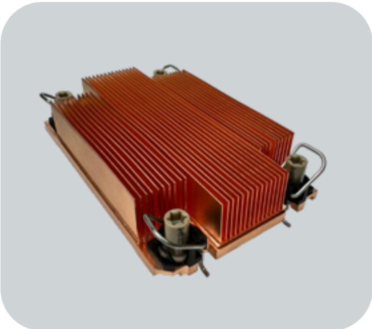

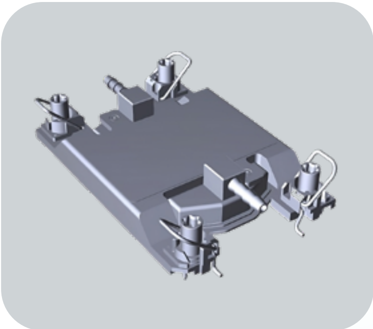

UP TO 500W	UP TO 1000W+	BEYOND 1000W
Immersion Optimized Heat Sink	Forced Convection Optimized Heat Sink (FCHS)	Immersion Cold Plate

Conclusion

This Intel and Submer technological breakthrough stands to substantially reduce the quantity and costs of server components that deal with heat capture and dissipation, setting single-phase immersion cooling apart as a frontrunner among competitive liquid cooling technologies.

As well as its immediate benefits for the state of datacenter cooling today, the forward-looking FCHS package establishes a solid technological runway for future high-density single-phase immersion deployments.

Rigorous testing has proved that this innovative solution can effectively cool Intel Silicon up to 800W+ in single-phase immersion and return thermal management to the server. Furthermore, the FCHS add-on can be produced at substantially lower price points versus water-based DLC components.

Next Steps

The subsequent phase in the FCHS roadmap entails delivering the design to the industry, establishing a robust network of suppliers, and thereby creating an opportunity for any server manufacturer to seamlessly incorporate the FCHS package into their designs.

Following on from the milestones achieved with the FCHS package, Submer will subsequently focus on 800W+ TDPs and collaboration for the potential of immersion-optimized cold plates.

"An immersion heat sink utilizing forced convection is a key innovation in taking single-phase immersion cooling beyond the current barriers allowing single-phase immersion not only to be a solution of today but also a solution of the future!"

Mohan J Kumar,
Intel Fellow

"Many have challenged the technological runway of single-phase immersion cooling. The Forced Convection Heat Sink is the undeniable proof that immersion is here to compete head-on with other liquid cooling technologies, including Direct Liquid Cooled water-based cold plates."

Daniel Pope,
Co-Founder and CEO at Submer

About Submer

Founded in 2015, Submer provides best-in-class technology, enabling datacenters around the world to leverage the power of immersion cooling for HPC, hyperscale, data centers, Edge, AI, DL, and blockchain applications. Submer is headquartered in Barcelona, Spain, and operates a Gigafactory in Houston, Texas. For more information, visit submer.com.

About Intel

Intel (Nasdaq: INTC) is an industry leader, creating world-changing technology that enables global progress and enriches lives. Inspired by Moore's Law, we continuously work to advance the design and manufacturing of semiconductors to help address our customers' greatest challenges. By embedding intelligence in the cloud, network, edge and every kind of computing device, we unleash the potential of data to transform business and society for the better. To learn more about Intel's innovations, go to newsroom.intel.com and intel.com.

Authors:

Daniel Pope - CEO and Co-Founder, Submer

Mark Miyoshi - Director of R&D, Submer