


White Paper

The Future of Immersion Cooling: The Path to Cooling 1000W Chips, and Beyond!



The Future of Immersion Cooling: The Path to Cooling 1000W Chips, and Beyond!

- Introduction 1
- The Density Paradox..... 2
- Immersion Cooling:
Near-Term Innovations and Future Directions. 4
 - Meeting Near-Term Needs..... 5
 - Replacement Heat Sinks 6
 - Alternate Thermal Interface Materials (TIM)..... 6
 - Colder Fluids..... 6
 - Cooling 1000W and Beyond 7
 - Hotter but Larger Chips 8
 - New, Purpose-Built Fluids 8
 - Precision Flow Enhancements 9
 - CPU and GPU Performance Criteria 9
 - Servers Designed for Immersion 10
- More than Cooling Power:
The Extended Benefits of Immersion..... 13
 - Sustainability 13
 - Power Savings 14
 - Cost Savings 14
 - Additional Advantages 14

Introduction


Over the last decade, immersion cooling has risen as a pivotal solution in the data center landscape. As traditional air and augmented air-cooling methods grapple with the thermal challenges posed by advanced, high-performance CPUs and GPUs, the shift towards liquid immersion has gained momentum. By placing electronic components directly in non-conductive fluids, this technique offers excellent heat absorption and rejection for the full server. The industry's pivot towards liquid immersion cooling is evident from its growing adoption, a testament to its efficacy in managing the robust thermal outputs of next-generation servers.

“Most data center operators expect air cooling to cede its position as the dominant method in cooling IT hardware within six years”

– Uptime Institute

Primary Drivers of the Switch to Liquid Cooling

Which of these do you think are primary drivers for liquid cooling adoption?
Choose no more than two. (n=740)



(Responses for “Other” are not included.)

Source: Uptime Institute Cooling System Survey 2023

While thermal limits of air cooling are being challenged, the current and upcoming densities are only scratching the surface of what is possible with liquid immersion cooling. To put that in perspective, GRC’s customers have successfully deployed systems supporting up to 400W TDP chips, currently the highest TDP mainstream CPU, and rack densities exceeding 120 kW/rack!

High-Performance Processors Ideal for Immersion Cooling¹



Given this backdrop, the industry now poses an essential question: Over time, can single-phase immersion cooling continue to keep pace with escalating heat loads?

Based on our expertise in the field, the answer is a resounding “Yes.” Immersion cooling is not only adept at handling the heightened heat outputs of today’s processors but is also strategically primed to address the challenges presented by next-generation applications and the increasingly dense ITE that enables them.

In this paper we examine immersion cooling’s trajectory, as well as spotlight its role in addressing a significant concern in the data center industry: the Density Paradox. The subsequent section will dissect this challenge, paving the way for an in-depth exploration of the enhancements on the horizon and the approaches being considered for cooling chips expected to surpass 1000W in the future.

The Density Paradox

When we think about data center cooling, we often consider the total heat load that needs to be cooled. But one metric that may be even more important than the total heat load is density. In the past, we’ve thought of density from the perspective of the whole data center, often evaluated in terms of its power density (W/m^2), hardware density (U/m^2), or from the perspective of rack density ($kW/rack$).

However, neither the average data center density nor the rack density have changed as dramatically as one would expect. One of the primary reasons densities have stayed relatively flat, while hardware continues to get hotter, is due to the limitations of existing power and cooling infrastructure. Most data centers in operation today were designed and built with air cooling in mind, which typically maxes out at $\sim 15 kW/rack$. Newer air-cooling technologies such as rear door heat exchangers (RDHx) can increase that density to 20-25 $kW/rack$, but anything beyond that starts to push the limits of practicality.

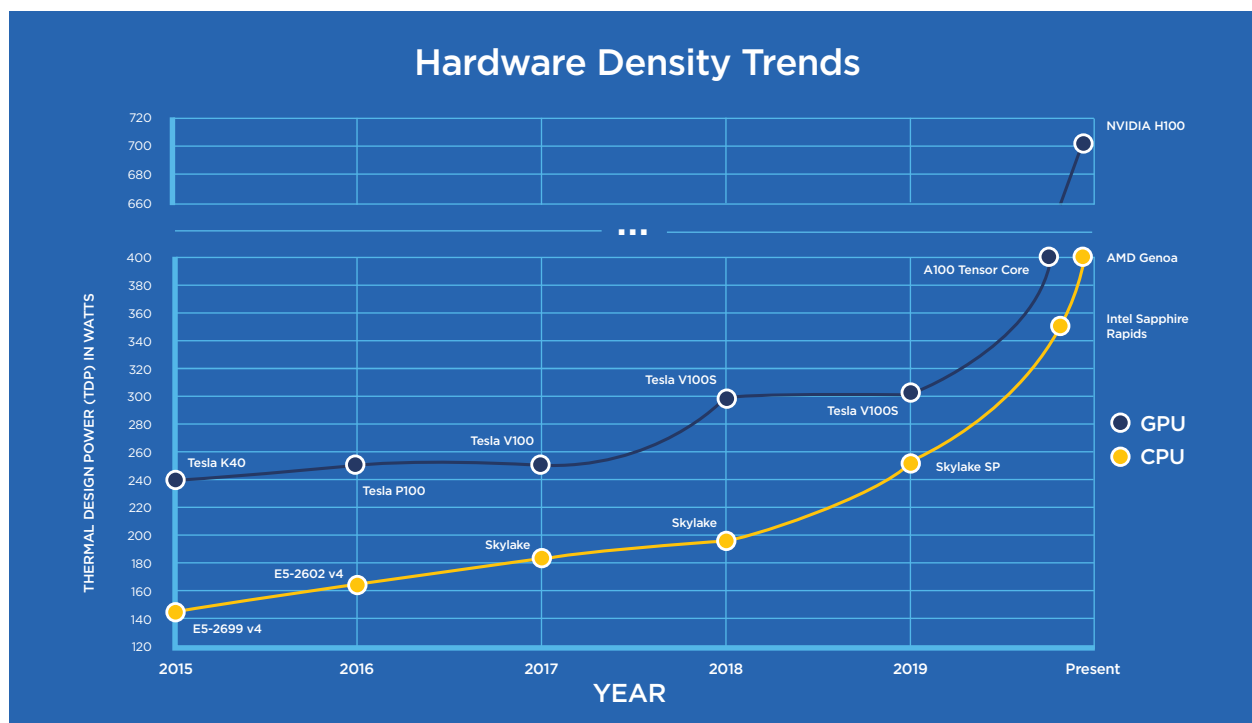
When looking at how data center cooling may evolve in the coming years, we need to drill down into the drivers of thermal density. Specifically, we need to consider three additional factors that are changing dramatically and will bring about a paradigm shift in the way we cool servers in the data center of the future:

¹Not a comprehensive list.

- Hardware chip density, in terms of the number of chips (CPUs/GPUs/accelerators) per server
- Chip power density, in terms of the thermal design power (TDP) of chips
- ESG regulations and company targets forcing dramatic reductions in power usage

These factors are driven by macro trends such as:

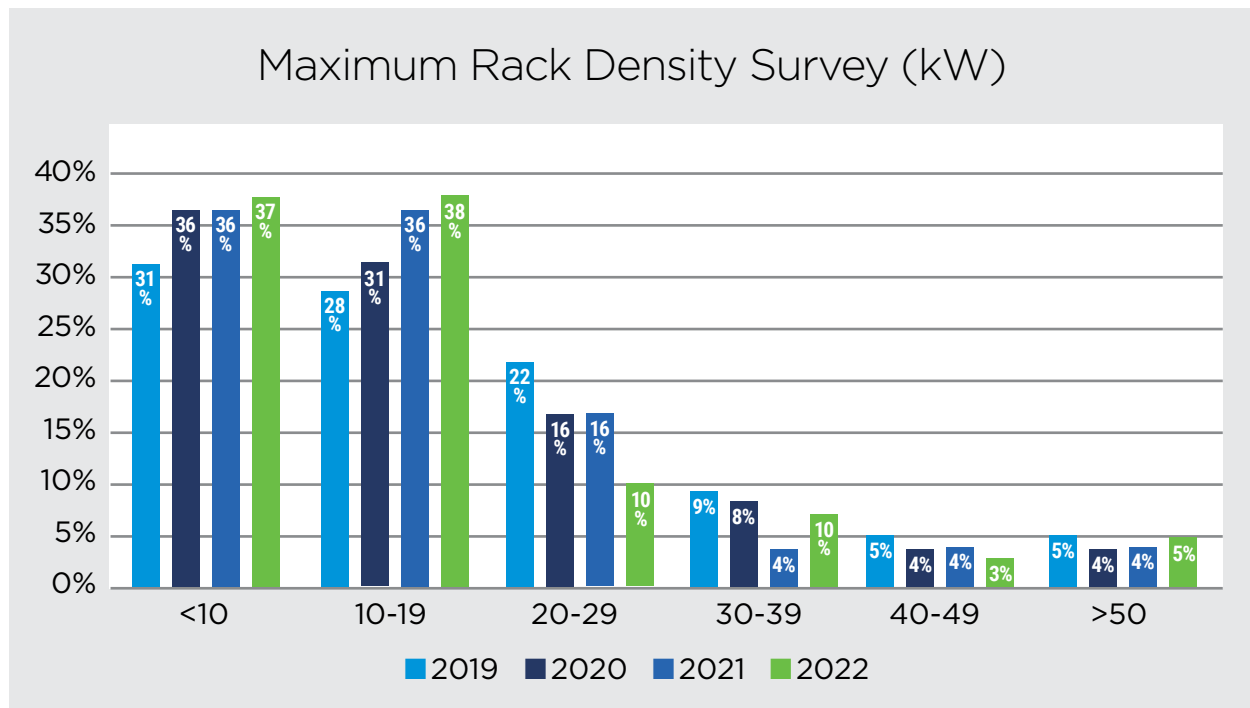
- Power density, or thermal design power (TDP), of chips, pushed by increasing workloads and Moore's law
- Accelerated hardware in mainstream data centers due to advanced applications such as artificial intelligence (AI) and machine learning (ML)



Source: GRC Internal Research

Organizations of all sizes are leaning into processing the massive workloads driven by high-performance computing, AI, data analysis, ML, scientific simulation, and graphics rendering. To accomplish these levels of compute, many data center owners and operators are choosing parallel processing to handle multiple tasks simultaneously and increase efficiency by distributing them across multiple graphics processing units (GPU) and central processing units (CPU) – each of which generates significant heat.

But while there is an increasing amount of hardware with higher TDP CPUs and even higher TDP GPUs making it into mainstream, enterprise-level data centers, the fact is that the majority of servers already deployed or being deployed today are still using fairly low wattage chips (<200W). This is another reason why we see the density paradox: density within the servers has risen, but that of the data center hasn't quite matched this growth.



The relative and absolute number of high-powered servers deployed today represents only a small fraction of the larger server population. Of course, this is bound to change as processor speeds continue to climb and the high-performance servers of today become the general-purpose servers of tomorrow.

Addressing the intricacies of the Density Paradox requires a comprehensive, forward-looking strategy. Single-phase immersion cooling not only meets today's cooling needs but is also adapting to the rapid advancements in chip technology.

Immersion Cooling: Near-Term Innovations and Future Directions.

In our exploration of the evolving processor landscape thus far, we've delved into the increase in power and thermal output, emphasizing the growing challenges for traditional cooling technologies. Immersion cooling has not only proven itself as a contemporary solution but is also extremely well positioned to be the front-runner in overcoming future cooling demands.

Ben Smith, Chief Product Officer at GRC, explains immersion's trajectory: "Our comprehensive testing, combined with valuable insights from our partners and customers, confirms that we can adeptly cool 400W TDP chips using single-phase immersion today. Yet, we believe

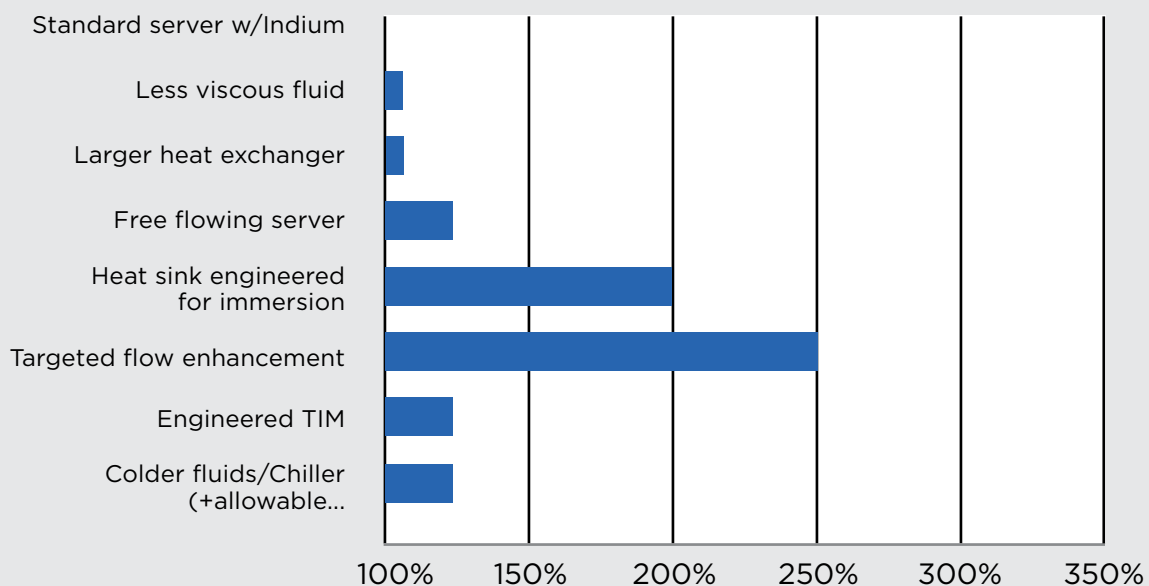
we're just getting started. The forthcoming advancements in cooling fluids, refined cooling techniques for CPUs, immersion-specific server designs, and system-level enhancements will position immersion cooling to confidently manage the thermal outputs of 1000W chips and even more formidable ones in the future."

In this section, we'll delineate the near-term refinements to immersion cooling, many of which have already undergone rigorous validation. Next, we'll outline a road map of immersion cooling advancements as we foresee it adapting to meet the intensified thermal demands of processors in the distant horizon.

"The forthcoming advancements in cooling fluids, refined cooling techniques for CPUs, immersion-specific server designs, and system-level enhancements will position immersion cooling to confidently manage the thermal outputs of 1000W chips and even more formidable ones in the future."

— Ben Smith, GRC Chief Product Officer

Identified Areas for Advancement of Immersion Cooling



Source: GRC Internal Research

Meeting Near-Term Needs

As mentioned earlier, GRC is already cooling a TDP up to 400W with existing systems, fluids, and server technology. And GRC testing, simulation, and analysis show that cooling CPUs beyond 750W TDP can be readily achieved with currently available solutions, even with servers that were designed for air cooling and subsequently adapted for immersion.

Some of the methods that are being applied today to boost the maximum cooling capacity of single-phase liquid immersion cooling solutions include:

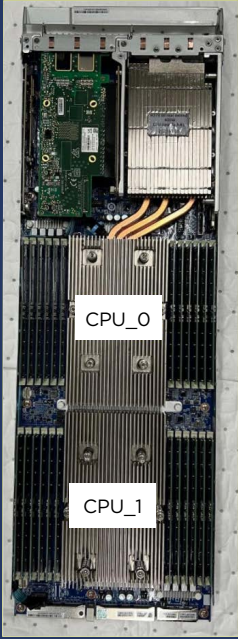
- **Replacement Heat Sinks.**

In a collaboration between GRC, Unicom, and Intel, replacing standard air-cooled heat sinks with immersion-designed alternatives showed up to a 100% performance boost. This improvement stems from a halved thermal resistance. Simply put, thermal resistance, the inverse of peak TDP, determines the temperature difference needed for efficient operation. For instance, a 0.10 C/W resistance means a 500W CPU requires a 50C temperature gap between its case and the coolant. Thus, reducing this resistance allows for greater heat load capacity at the same coolant conditions.

Production Server with Eight 400W CPUs Is Easily Cooled

Immersion-optimized heat sink can pass the thermal test under the condition 400 W Genoa CPU stress test with coolant at 35 °C

- CPU_0 junction temperature is 7-11 degrees lower than the limitation
- CPU_1 junction temperature is 9-12 degrees lower than the limitation



Source: GRC Internal Research

- **Alternate Thermal Interface Materials (TIM).**

Most immersion cooling solutions currently use Indium or other foil based TIMs today. However, there are numerous alternative TIMs already on the market that can help lower thermal resistance and improve performance. Recent testing of alternate materials on Intel Sapphire Rapids processors improved maximum wattage by 25%.

- **Colder Fluids.**

In most environments, immersion cooling systems efficiently provide chiller-free cooling. However, as CPU Tcase temperatures approach ambient wet bulb and dry bulb levels, the option of using non-refrigerant-based outdoor heat rejection might become

unfeasible. Many contemporary data centers are already equipped with chilled water systems. Immersion technologies can capitalize on this existing infrastructure in retrofit situations, drawing only the necessary cooling to guarantee optimal server performance. Though cooler fluids can enhance wattage support by approximately 30%, resorting to this method should be seen as a backup to other strategies, ensuring the sustainability advantages of transitioning to immersion cooling remain paramount.

Cooling 1000W and Beyond

While immersion cooling technology in its current form seems more than sufficient to meet the needs of the near future, it will need to continue evolving as chip and hardware densities grow beyond their current capacities. Much like the evolution we saw with air cooling, over the years, liquid immersion cooling will evolve to deliver performance far beyond what is believed possible today.

In fact, we are confident that we will see single-phase liquid immersion cooling performing with servers using chips 1000W and higher. This technology is designed specifically to address challenges posed by high-power, high-density compute environments, in an efficient, cost-effective, and scalable manner. GRC has been successfully deploying immersion cooling solutions in high-performance data centers across the globe since 2009. However, yesterday's high-performance computing is today's enterprise application.

As data centers move further along the higher power/density curve, immersion cooling is uniquely positioned to serve a broad spectrum of use cases.

Single-Phase Liquid Immersion Cooling Use Cases

- High-density computing environments where servers are packed densely and generate a significant amount of heat in a confined space.
- GPUs performing compute-intensive workloads such as machine learning, artificial intelligence, and scientific simulations.
- Edge computing and remote data centers with limited cooling infrastructure or harsh environments with increased temperature variations.
- Environmentally prudent organizations looking to improve energy efficiency to cut down carbon emissions and OPEX.


The path to cooling over 1000W chips with immersion will likely require multiple pieces of the puzzle coming together. As the whole ecosystem grows and evolves, it will require tight collaboration between key stakeholders to deliver unparalleled value to end-users.

Here is how we anticipate the various components evolving:

Hotter but Larger Chips

Data centers are incorporating larger, more powerful chips, including multicore, which are packing multiple processing units on a single chip to enhance computational capabilities and handle multiple tasks simultaneously. As we approach the physical limitations that prevent chips from getting denser, the size of chips is increasing to pack in more computational power. And with that increase in chip size comes an increase in wattage, as the chips use more power. TDP and chip size will in many cases grow together, though likely at different rates and the cooling difficulty factor will not necessarily scale proportionally to TDP. If the size and hence the surface area of the chips increases faster than their TDP, this would effectively reduce the power density (W/cm^2). So while the total heat load would still increase, the lower concentration of the heat load may help reduce the cooling difficulty.

Multicore CPU Example




New, Purpose-Built Fluids

Currently, most single-phase immersion systems repurpose off-the-shelf fluids within the systems. And while there are a wide range of compatible fluids on the market, the development of purpose-built fluids is still in the early stages. GRC is working hand-in-hand with multiple fluid providers to further the development of higher-performance engineered fluids. The goal is to deliver the highest performance possible, including further advancements when the heat sink is aligned with the fluid. The fluids that exist today are more than sufficient to meet the current needs of the market. By the time high-end applications need more cooling power, the innovations these teams are working on will result in fluids with improved performance.

PAO4 provides an improvement at the same cost, even when the heat sink was designed for PAO6.

GRC ElectroSafe® Fluid Partners Investing in Advanced Immersion Fluids



“When we’re making these immersion-specific fluids, we’re really dialed into the unique cooling needs of today’s advanced processors. We work closely with partners like GRC and their clients, and also with server manufacturers, so we truly understand what the market needs. Keeping in mind both what’s happening now and what’s coming up in terms of heat from these processors, we’re all in on making sure our fluids keep immersion tech ready for any thermal challenge.”

— ENEOS Corporation, Mr. Toru Konishi, Company President,
Lubricants Company

Precision Flow Enhancements







The cooling capacities of immersion cooling systems will exceed the demands of the hardware they support, today. This additional headroom has allowed for rather simplistic system designs where whole servers, if not racks, could effectively be treated as uniform heat loads. However, as individual components such as CPUs get more powerful (TDP increases) and require lower operating temperatures (Tcase decreases), directing the coldest fluid to the hottest components first will result in improvement. Initial studies show over 2x improvement with directed flow. Hence it seems inevitable that immersion cooling systems themselves will evolve to enable more precisely targeted flow than what is seen in today’s systems.

CPU and GPU Performance Criteria

The manufacturers of CPUs and GPUs give strict guidance regarding Tcase temperature to assure performance of their devices. As this is the boundary between their thermal management and the Server OEMs and Infrastructure Cooling solutions boundary. The vast majority of devices have TCase maximum values around 75°C with a very small percentage needing

significantly lower TCase values may be as much as 20°C lower or in the vicinity of 55°C. The lower the Tcase values the more difficult the challenge for cooling regardless of the technology. Ideally Tcase values should be kept as high as possible to promote cooling with warmer fluid temperatures whether air or immersion liquids. High Tcase values can also help greatly increase the heat reuse potential of the cooling infrastructure as higher heat quality is an important factor in the diversity and benefit of heat reuse. The below chart gives a general idea of implication of TCase versus TDP and how they relate to cooling difficulty. As Tcase gets lower and TDP increases, many of the enhancement discussed in this paper will need to be incorporated to manage the thermal imperative of particular CPUs or GPUs

Relationship of TDP to Cooling Difficulty


Scenarios	Tcase ↓	Tcase ~	Tcase ↑
TDP Power Steady	 Difficult	 Normal	 Easy
TDP Power Increase	 Hard	 Difficult	 Normal

Source: GRC Internal Research

Servers Designed for Immersion

As the immersion ecosystem grows, there will be a growing number of servers that are designed and built specifically for immersion cooling, not merely adaptations of servers designed for air cooling. There are already several OEMs that are designing servers specifically for immersion and GRC is actively working to build the future of data center cooling with these and other forward-looking manufacturers to design and build servers intended to be cooled solely by immersion.


An immersion-designed server will provide multiple opportunities for enhancing performance. In fact, it is probably the area where the biggest gains will be made in the near future. Heat sinks, component placement for maximum flow, better TIM, and even directed flow methods will rely on server OEMs creating these servers designed for immersion rather than modified from air designs. Any server designed for air cooling will work and indeed benefit from being modified and run in immersion systems. When air is no longer an option and we are left purposefully designing for different types of cooling, a design optimized for immersion will enable much higher performance CPUs and GPUs than today’s server designs.



Purpose-built, servers designed for immersion maximize fluid flow creating large improvements over servers designed for air-cooling.

• **Improved Flow and Thermal Design**

Apart from the immersion cooling systems evolving to enhance targeted flow to the hottest components, optimizing the design of the server itself can significantly amplify the results. Simply removing components such as server and PSU fans, while improving layout to have the hottest components towards the bottom of the server can yield significant improvements. In the chart below, a simple change of removing unused backplanes and drives from the server improved thermal resistance by 30%. Servers built for higher-wattage chips or immersion have better flow.




Source: GRC Internal Research

- **Engineered Heat Sinks**

While GRC along with partners such as Unicom and Intel are already developing and deploying replacement heat sinks for air-cooled servers, these heat sinks are designed to be retrofittable and hence constrained by the geometry of the existing server/chassis and are optimized for cost and ease of replacement. However, designing a server from the ground up, for immersion, will allow further optimization of heat sinks to reduce thermal resistance with immersion. Even small changes in heat sink design can make big improvements in performance. A quick iteration of a basic immersion heat sink design showed a performance increase of 15%. One test of a server with 1U high blades available on the market today with two CPUs arranged vertically (one directly above the other), used thermally shadowed processors—and the shadowed processor was colder due to an extended heat sink, which lowered the thermal resistance between the base of the heat sink and the cooling fluid.

- **New TIM Development**

As previously stated, solid TIMs such as Indium foil have historically served immersion-cooled servers well and GRC is testing other TIMs that are readily available and could help enhance thermal performance (see chart below). However, current testing is focused on existing solutions and not on the development of new materials. Much like the fluids, this is also an area that is ripe for innovation. As the overall immersion ecosystem grows, there is likely to be growing focus and innovation in this area. However, given the wide variety of TIMs already in the market, it may be a while before such innovation is “needed”.



Source: GRC Internal Research

In addition to the thermal benefits of designing servers specifically for immersion, there are also notable cost and reliability advantages. Eliminating extraneous components, such as server fans, not only reduces costs but also enhances reliability. This eliminates potential failure points and sources of vibrations that could lead to failures. Moreover, servers designed for immersion are optimized for coolant flow, minimizing the presence of unnecessary metal and plastic bays that could impede flow. Using fewer materials also helps in lowering raw material and shipping expenses, as well as reducing the server's carbon footprint.

The ongoing evolution of processors underscores the need for future-ready, immersion cooling solutions. Taking advantage of the above technological advancements, immersion will lead the way in meeting this need, not only effectively dissipating the heat, but also by offering reductions in power consumption, more streamlined data center designs, and a significant reduction in the total cost of ownership. This trajectory is being championed by a collaborative ecosystem of innovative companies, each contributing to the developments necessary for the efficient cooling of tomorrow's high-performance chips.

Intel, ever at the forefront of processor technology, reiterates this vision: “As leaders in CPU innovation, Intel is acutely aware of the looming thermal challenges as chip technology pushes forward. Immersion cooling has already showcased its capabilities for today's demands, and with continued innovations on the horizon, it stands as the preferred solution to meet the formidable cooling needs of the future.” This shared vision among industry leaders underscores the pivotal role of immersion cooling in the future of data centers.

“At Intel, we have been testing immersion cooling systems for several years and have seen tremendous progress with this innovative technology. We are confident that we are on a path towards single-phase immersion cooling being future proofed up to and beyond 1000W.”

— Mohan J. Kumar, Fellow at Intel Corporation



More than Cooling Power: The Extended Benefits of Immersion

Choosing the right data center cooling technology is a critical business decision, one that can have lasting and significant impact on an organization's top and bottom line. And while the thermal challenge right in front of us may make this seem like a pure physics problem, this is in fact a strategic business challenge. Hence, it's important to step back from the laser focus on thermal efficacy of a cooling solution and look at the business impact of the various solutions as well.

With single-phase liquid immersion cooling you can get to the future of data center cooling today delivering the cooling capacity demanded by the widest array of data center types and computing environments; while realizing sustainability benefits, power savings, cost savings, and other advantages compared to other cooling solutions.

Sustainability

Immersion cooling enhances energy efficiency, which leads not only to reduced energy consumption but also a significant reduction in carbon footprint. This cooling method permits

denser server configurations, promoting a compact data center design. It conserves water and extends the hardware lifespan, contributing to resource conservation. Furthermore, almost complete waste heat recovery is achievable, emphasizing its environmental advantages.

Power Savings

Immersion cooling can decrease total data center energy consumption by up to half. By eliminating components associated with air-cooling, it offers up to a 90% cut in cooling energy, delivering an extremely efficient pPUE of <1.03.

Cost Savings

Immersion cooling leads to significant financial benefits, with potential reductions in both CapEx and OpEx by up to 50%. The streamlined design minimizes infrastructure needs and operational costs, while the lower energy consumption translates to substantial power savings. A reduced data center footprint also ensures further savings on real estate.

Additional Advantages

The simplicity of immersion cooling systems results in enhanced reliability with fewer potential failure points. They operate silently, negate technician heat stress, and provide protection against environmental contaminants. The coolant is durable and eco-friendly, with no evaporation or loss of cooling capacity over time. In critical situations, immersion cooling offers an extended ride-through time, providing a safety net during unexpected outages.

The above points only provide a brief glimpse into the benefits of immersion cooling. For a comprehensive exploration of this technology, please download our eBook 'The Definitive Guide to Immersion Cooling'.

With mature products and more than 14 years of real-world testing in data centers, GRC understands liquid immersion cooling. We enable data centers of all sizes and configurations to unlock computing power while maintaining optimal performance and reliability.



Connect with a GRC Expert to Position Your Data Center for Success Today and into the Future

Email: ContactUs@grcooling.com **Visit:** grcooling.com

Call: +1.512.692.8003

GRC, 11525 Stonehollow Dr. Suite A-135, Austin, TX 78758, United States

©2023 GRC, Green Revolution Cooling, and The Immersion Cooling Authority are each registered trademarks of Green Revolution Cooling, Inc.

Green Revolution Cooling, Inc. ("GRC") believes this information to be accurate; however, GRC does not make any representation or warranty, express or implied, as to the accuracy or completeness of any such information and shall have no liability for the consequences of the use of such information.