DATA CENTER Frontier

SPECIAL REPORT

New Data Center Efficiency Imperatives: How Immersion Cooling is Evolving Density and Design

Author: Bill Kleyman, Contributing Editor, Data Center Frontier
Draft Date: 06/02/2023
Draft Version: 5.0

brought to you by

Hypertec

Contents

Introduction 3
A cooler way to use immersion cooling in the data center 3

Section 1 – Immersion Cooling: Trends that Transform the Data Center 5
Evolution of Liquid Cooling 5
New business demands translate to new infrastructure challenges 5

Section 2 – Immersion Cooling: Top Five Design and Deployment Considerations 7
Design considerations and working with immersion cooling 7
Top Five Immersion Cooling Myths 10

Section 3 – Immersion Cooling Readiness Checklist 11
What can be submerged? 11
How to become immersion ready 12

Section 4 – The Immersion Cooling Journey: Innovation in Infrastructure 14

Final Thoughts 15

SUMMARY:

You won't need to look too far back to realize just how key critical infrastructure is to our everyday lives. Data center ecosystems work hard to keep everyone connected, operational, and productive. Over the past few years, we've seen density levels rise as organizations are tasked to deliver more capacity with limited space.

Considering this, the evolution of efficiency, density, and functional design must change. Leaders in the data center and colocation space have turned their attention to immersion cooling technologies that have proven highly effective and efficient and deliver new capabilities for critical infrastructure. Looking at costs specifically, data center leaders must zero in on cost efficiency and workload density. It's a big reason many leverage immersion cooling to reduce data center OPEX by over 90% and reduce CAPEX costs by more than 50%.

This paper will explore the fundamental shift happening in our data center industry. Specifically, we'll explore the evolution of immersion cooling and how data center leaders use this technology to achieve new efficiency imperatives.

Introduction

How we design and build critical infrastructure has changed so much over the past few years. One thing is for sure: Data centers aren't going anywhere, and demand is increasing sharply due to various new use cases.

We see continual changes in reach density, in both heat and power, to keep up with high-performance computing (HPC) trends, which include Artificial Intelligence, Automation, and Machine Learning.

Here's the other key challenge: Many current solutions can only sustain power requirements, heating production increases, and future flexibility for so long. There are new demands around cloud computing, big data, and infrastructure power/cooling efficiency. With advancements in new use cases and demand and the rapid growth in data leading the way within many technological categories – working with the right data center technologies has become more critical than ever.

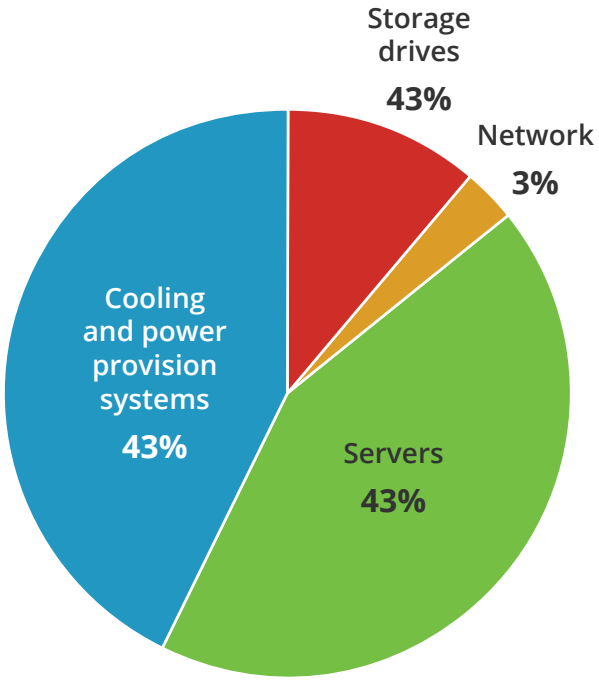

Figure 1. Fraction of U.S. data center electricity use in 2014, by end use. Source: Shehabi 2016.

This is where immersion cooling comes into play. According to a recent study, on average, servers and cooling systems account for the most significant shares of direct electricity use in data centers, followed by storage drives and network devices (Figure 1). Some of the world's largest data centers can contain tens of thousands of IT devices and require over 100 megawatts (MW) of power capacity to power around 80,000 U.S. households (U.S. DOE 2020).

This means new imperatives exist in how we cool, maintain, and optimize our data center facilities.

A COOLER WAY TO USE IMMERSION COOLING IN THE DATA CENTER

Immersion Cooling isn't anything new; however, recent technological advancements have allowed more data centers to adopt the technology.

Why is this happening? Rising investments in high-density technology, high-performance computing, and data-driven workloads have pushed data center leaders to design the most reliable and efficient methods to cool their data centers. Today, data centers consume about 2-3 percent of global energy. However, according to the latest findings from the IEA, data centers and data transmission networks are responsible for nearly 1% of energy-related GHG emissions. Since 2010, emissions have grown modestly despite rapidly growing demand for digital services, thanks to energy efficiency improvements, renewable energy purchases by information and communications technology (ICT) companies, and broader decarbonization of electricity grids in many regions. Solutions like immersion cooling help reduce carbon footprint by removing water waste and revolutionizing heat dissipation.

It's critical to note that business and data center developments have pushed immersion cooling into the forefront of infrastructure design. Let's explore a few of these new developments:

- **Data centers are using more power.** According to a recent post from the IEA, digital technologies have direct and indirect effects on energy use and emissions and hold enormous potential to help (or hinder) global clean energy transitions, including through the digitalization of the energy sector. Today, data centers and data transmission networks are responsible for nearly 1% of energy-related GHG emissions. Furthermore, a recent US Department of Energy report indicates that US data centers are projected to increase the amount of energy they consume based on current trend estimates. A trend that's been steadily rising since 2000.

- **HPC, AI, and data-driven solutions require a new approach.** Here's a specific example. High-performance computing (HPC) is being leveraged for research and data analysis, and these systems are being deployed within traditional data center walls. However, cooling for HPC systems requires a different kind of approach. Immersion cooling helps dissipate heat from heavily used HPC systems and supports more significant density levels. The growth in data-driven solutions has directly impacted how data center leaders view cooling critical systems. It's a big reason immersion cooling and immersion-ready technologies have supported emerging use cases.

  Innovation around immersion-cooled data centers and specific use cases doesn't stop at HPC. Emerging requirements around machine learning, financial services, healthcare, CAD modeling, and rendering are exploring new immersion-cooled solutions to maintain sustainability, reliability, and the greatest density levels.

A recent report from Technavio indicates that the adoption of immersion cooling is high as they are considered more efficient than air-based cooling. Liquid-based cooling, a sub-segment of immersion liquid-based cooling, **is the most widely accepted cooling system**. Finally, Dell EMC's study on the efficiency and liquid cooling dove into the cost of cooling data centers and how it impacts operators' overall operating expenses. 50% responded that it had a significant impact, and 78% responded that there was at least some impact on their costs.

The survey also showed us that 25% already use liquid cooling in production environments, and 32% use liquid cooling in some fashion within their data center.

As the survey points out, there are many great ways to use liquid cooling within the data center.

- 39% have liquid cooling technologies within high-density cabinets or racks
- 20% deploy liquid cooling for *high-performance computing (HPC) use cases*
- 17% use it for big data processing
- 15% already have full-package systems where the **majority of the heat is removed through liquid**

More of your customers will be asking for highly dense, compute-intensive solutions. Can today's air-cooled ecosystems keep up? Immersion cooling has pushed boundaries on efficiency and density. Today, these solutions are a bridge to delivering the most advanced workloads to your customers.

Let's dive into how immersion cooling changes data center design. Most of all, to get started, you'll see how immersion cooling solutions today are much different than ever before, making adoption far more feasible.

Section 1 – Immersion Cooling: Trends that Transform the Data Center

It's not surprising that immersion cooling has been around for quite some time since the 60s. But what is it, exactly? We'll get into a few types of immersion cooling in a minute. Still, at a very high level, immersion cooling is a method adopted by data centers to quickly cool IT hardware by directly immersing it in a non-conducive liquid. The heat generated from the components is directly and efficiently transferred to the liquid. It offers the highest level of efficiency plus virtually unlimited capacity. With that in mind, it's important to note that liquid cooling has existed since the 1970s.

EVOLUTION OF LIQUID COOLING

Between 1970 and 1995, much of the liquid cooling was done within mainframe systems. Then, in the 90s, we saw more gaming and custom-built PCs adopt liquid cooling for high-end performance requirements. Between 2005 and 2010, liquid cooling entered the data center with chilled doors. From there, looking into 2010 and beyond, liquid cooling was used in high-performance computing (HPC) environments and designs featuring direct contact and total immersion cooling solutions.

Source: Hypertec

NEW BUSINESS DEMANDS TRANSLATE TO NEW INFRASTRUCTURE CHALLENGES

Across the world, connectivity and infrastructure have changed dramatically. Businesses and the corresponding infrastructure supporting digital connectivity have seen an explosion of newly remote workers working with new applications and leveraging new solutions to stay productive. For example, as the cloud evolution continues, the line between HPC, application virtualization, and even virtual desktop infrastructure (VDI) continues to blur. Technology leaders are looking for platforms that allow organizations to rethink how they deploy these critical enterprise resources to provide the maximum return on their investment and the highest end-user experience levels.

From a business and technology perspective, here are some of the critical challenges facing digital infrastructure:

- **Space and time-to-market are a concern**. In our introduction to the paper, we talked about a crucial statistic around data center availability and vacancy. According to DataCenteHawk, North American primary market vacancy is now averaging 4.4%, its lowest point ever, and in Northern Virginia, that vacancy rate is closer to 1%. As a result, data center leaders have been forced to rethink how they design critical infrastructure and go to market with new solutions. There is a need for space and the need to do more with that space.

- **Doing more with less and finding ways to reduce OpEx**. As noted in the previous point, business leaders are tasked with doing much more with less space. On top of that, there is a demand to lower operational costs by deploying more efficient technologies. Working with immersion cooling solutions can increase density

while reducing cooling OpEx by over 95%. Further, immersion cooling requires less infrastructure, resulting in a 50% reduction in CapEx of build costs.

- **Finding ways to maximize performance and improve data center efficiency.** Energy efficiency and data center management are critical considerations for a few reasons. First, not only are data center administrators working hard to cut costs, but a significant business objective is also to minimize management challenges and improve infrastructure efficiency. To support these new initiatives, immersion cooling has become a technology to explore and implement alongside infrastructure efficiency.

- **Legacy infrastructure is holding technology innovation and business planning back.** Legacy data center solutions have a bullseye on them as they're seen as inefficient solutions slowing down infrastructure modernization. Leaders in the data center space see that although these solutions "work," they aren't adding additional value to the business. It's a key reason why modernization efforts have become front and center for business and technology leaders alike.

  A key point in working with legacy IT is that you don't necessarily have to rip it all out. New IT components can be deployed parallel to legacy IT to ensure that the transition or integration is least disruptive and beneficial to IT, users, and businesses.

This is where we pause and begin to understand the broader adoption of new solutions, like immersion cooling, aimed explicitly at modernizing today's most critical infrastructures. New solutions are being leveraged to deliver highly efficient, high-density immersion cooling solutions to support emerging use cases around HPC, data-driven technologies, IoT, edge computing, and cloud.

In some cases, these are modular single-phase immersion-cooled data center platforms with fully integrated power, cooling, racking, networking, and management experience backed by partnerships with numerous industry leaders. Modern immersion cooling solves a combination of challenges, including the following:

- Sustainability-minded design goals
- Governmental regulation
- Rising power costs
- Technological demands for cooling cutting-edge CPUs and GPUs
- The need for more efficient and less power-hungry cooling systems is clear.

With all of this in mind, let's pivot and explore some of the new advancements around immersion cooling and how they help overcome the above challenges. In a special mini-section, we'll bust some older immersion cooling myths.

Section 2 – Immersion Cooling: Top Five Design and Deployment Considerations

It's no wonder immersion cooling has gained traction within the data center industry. Modern designs do not require leaders to replace their existing solutions. Rather, immersion cooling can sit alongside air-cooled and other types of infrastructure. However, there is a crucial reason why many are looking more closely at immersion cooling for a growing number of use cases. *Energy efficiency and optimizing overall data center performance* is a critical design concern.

New technologies like solid-state have different operating temperatures than their spinning disk counterparts, and this is also true for new use cases around converged systems, HPC, supercomputing, and others. The latest metrics show that regularly operating in the recommended 64.4F and 80.6F range, based on the ANSI/ASHRAE Standard 90.4-2016, will keep your data center nice and healthy. But what about efficiency?

Furthermore, keeping your data center operating at a lower-than-necessary temperature may not help performance. Instead, you're just keeping the temperature and paying for it with no added value.

DESIGN CONSIDERATIONS AND WORKING WITH IMMERSION COOLING

With all this in mind, let's explore the five key considerations when considering immersion cooling.

1. **Legacy infrastructure.** To retrofit or not. If your infrastructure is aging, you must look beyond immersion cooling to bring your data center up to par. However, legacy infrastructure doesn't have to be old to be a challenge. With new use cases around HPC and high-density computing, leaders need more than blowing air to support density. New components can be deployed parallel to legacy IT to ensure that the transition or integration is least disruptive and beneficial to IT, users, and businesses. Here is an important note for adoption. The evolution of chips has gone from large to small (for handheld devices, for example) and now back to large within the data centers. The heat generation is increasing quickly by returning to chips with more dense core capacity, PCIe lanes, and capacity. This density will only grow with new solutions like AI, ML, and now ChatGPT. As new and updated technology is added, adding and integrating immersion cooling would be the least disruptive, cost-efficient, and beneficial.

   This integration or transition can be done via management tools, systems integration, and even parallel build-out of new use cases and deployments. Most of all – this does not have to be a complicated process, and the entire process can be very empowering for all associates, IT leaders, and business stakeholders. By integrating and evolving from legacy IT, you learn more about your organization and the capabilities you require to support an emerging market.

   That being said, there are a few key considerations. First, there is a chance that a retrofit will make your data center far more complex. A retrofit could work if you already have a few vendors in your data center and manage the infrastructure well. However, in some cases, you'll need to look at advanced retrofits to reduce the number of vendors in your facility. In other cases, you will need to explore immersion-born designs.

FEATURES	RETROFIT	ADVANCED RETROFIT	IMMERSION-BORN DESIGN
Availibility	Multiple Vendors	Some Vendors + Hypertec	Hypertec Exclusive
Base	OTS Air-Cooled Server	COTS Air-Cooled Server	Immersion-Born
Vanity-Free Design	✖	✖	✔
Ease of Installation and Serviceability	Low	Medium	High
Performance Availibility (OC or Turbo Lock)	✖	Hypertec Exclusive	✔
Supports Both AC & OCP Power	✖	✔	✔
Sustainable Composite Material	✖	✖	✔
Custom Immersion Cooling Heatsink	✖	✔	✔
Intel Xeon and AMD EPYC 3rd Gen Compatibility	✔	✔	✔
Intel Xeon W/SP and AMD EPYC 4th Gen Ready	✔	✔	✔
CPU Density Per SmartPod (48U)	192	192	288

2. **When the airflow is no longer enough.** As data center airflow management reached mainstream status in the past few years, the evolution of this field has focused on fine-tuning all the developments of the preceding decade.

   So – what's changed? During the '90s and mid-2000s, designers and operators worried about the ability of air-cooling technologies to cool increasingly power-hungry servers. With design densities approaching or exceeding 5-10 kilowatts (kW) per cabinet, some believed that operators would have to resort to rear-door heat exchangers and other in-row cooling mechanisms to keep up with the increasing densities. At this point, many realized that traditional airflow might not be enough for AI, HPC, and other data-driven workloads.

   At the data center level, improvements to individual systems entail the implementation of accelerator processors. More specifically, graphics processing units (GPUs), application-specific integrated circuits (ASICs), and field-programmable gate arrays (FPGAS). Today's HPC and AI systems are highly dense. Depending on the design, traditional air-cooled solutions might not handle the density you need, requiring you to spread the workload and utilize more data center floor space. Or, you could turn to immersion cooling. In data center design, immersion cooling can increase server density by 10x or more, handling up to 144 Nodes, 288 CPUs, and 98kW per 48U cooling tank.

3. **Designing around efficiency and performance.** Everyone wants to be as efficient as possible. However, the challenge is ensuring that you're not losing cooling power as you try to save on energy costs. During the evolution of liquid cooling, we saw direct-to-chip liquid cooling solutions. Today, high-density, modular single-phase immersion cooling solutions have all the necessary compute, network, and storage components. As mentioned earlier, liquids are far more efficient at removing heat than air cooling. For use cases where efficiency and performance are critical, purpose-built immersion cooling designs must be considered.

4. **Vendor considerations.** In the past, immersion cooling was more of a puzzle-piece configuration. You would install only what you needed, retrofit parts of your data center, and operate independently with a given vendor. Today, immersion cooling solutions are integrated, encompassing your entire use case in a single immersion-cooled tank. As exciting as that might seem, this is also an excellent opportunity to get

to know your vendors and what they can do for your data center. Although these new types of integrated cooling tanks are far easier to deploy into a data center, validating and challenging your vendors is still essential. During the evaluation process, ensure that the design fits your overall architecture and supports business operations moving forward.

5. **Understanding cost, long-term and short-term.** A look at the cost from a different perspective will be critical. That's why we're spending some extra time on this section. Because of the data center industry's unprecedented rise in power consumption, we've seen a significant increase in operational and power costs. It has become a challenge for end-users to manage and conserve power in data centers.

   This is where we need to look at immersion cooling in the near and long term. Generally speaking, air cooling systems cost less due to their more straightforward operation; however, the conversation around cost and cooling efficiency changes as rack density increases. Furthermore, our industry sees air-cooled data center systems' clear economic and density limitations.

   In a recent study by Schneider Electric, we see that for like densities (10kW/rack), the data center cost of an air-cooled and liquid-cooled data center is roughly equal. However, immersion cooling also enables IT compaction; with compaction, there is an opportunity for CapEx and OpEx savings. Compared to the traditional data center at 10 kW/rack, a 2x compaction (20 kW/rack) results in a first cost savings of 10%. When 4x compaction is assumed (40 kW/rack), savings go up to 14%.

   If we examine cost more broadly, we'll see areas where immersion cooling impacts technology alongside the business. In the most recent AFCOM State of the Data Center report for 2023, we noticed that OPEX and CAPEX costs were rising. According to the report, most respondents reported increased operational expenditures in 2022 (57%). The top drivers of OPEX increases are increasing equipment service and personnel costs, followed by rising energy costs. Similarly, most respondents (53%) reported increased capital expenditures in 2022. The top drivers of CAPEX increases include supply chain challenges, investment in existing facilities construction, investment in an IT refresh, and investment in new facilities construction.

   Immersion cooling aims to reduce costs for both OPEX and CAPEX. For example, with a single-phase immersion cooling solution, you could see a 95% reduction in cooling OPEX. With a PUE of 1.03 (certified by a 3rd party), many will realize an ROI of less than one year, even considering power savings. Similarly, single-phase immersion cooling is rapidly deployable in any greenfield or raw data center

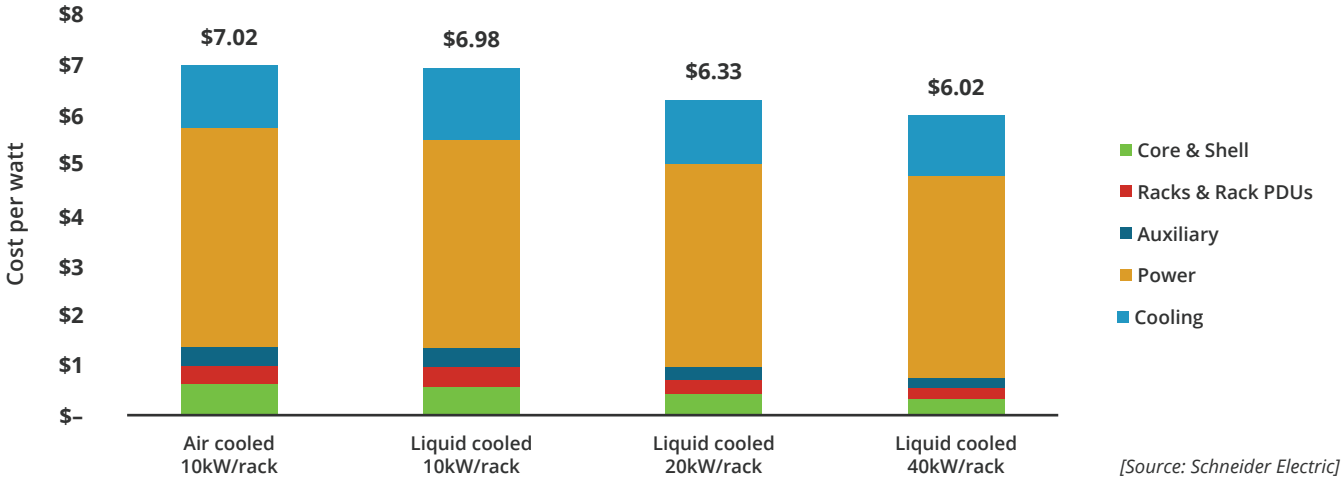

[Source: Schneider Electric]

environment. Operators will not need raised floors or cold aisles. Further, there is minimum retrofitting required for existing data centers. As a result, data center business leaders will see up to a 50% reduction in CAPEX building costs.

As an essential point, immersion cooling allows you to leverage a simpler compute ecosystem. There will be no moving parts, dust particles, and vibrations for immersion systems. Due to the liquid's uniformity and viscosity, the single-phase immersion design has less thermal and mechanical stress. As a result of creating a more streamlined and simplified ecosystem, operators will see a 30% increase in hardware lifespan.

Because immersion cooling continues to gain traction in the data center industry, there are still a few unknowns about the technology. Let's break down the five biggest myths related to immersion cooling.

TOP FIVE IMMERSION COOLING MYTHS

1. **Water treatment and biodegradability.** New integrated immersion cooling solutions leverage coils that do not use any special water treatment package that isn't already used in the data center cooling loop. Furthermore, modern single-phase liquid cooling solutions work with fully biodegradable liquids. Again, there is no toxicity in any of the liquids being used.

2. **Combustible oils.** Today, most immersion liquids and even closed-loop liquids contain oils. If the oils are heated to evaporation or condensation, they technically are flammable; hence the labels affixed to those liquids. Enclosed loops and single-phase immersion solutions are the safest regarding combustibility. Solutions like Hypertec leverage a single-phase immersion cooling where servers are submerged in a thermally conductive dielectric liquid or coolant, a much better heat conductor than air, water, or oil. The coolant never changes state, and the liquid used in leading immersion cooling solutions will not be combustible.

3. **Servicing, maintenance, and complexity of running air and immersion cooling systems.** This is where we look at next-generation immersion cooling solutions. Even in advanced retrofit situations, installation and serviceability are much easier than in previous immersion cooling designs. For immersion-born solutions, serviceability is easier via easy access to all critical components.

4. **Air cooling is understood better and more widely used.** For the time being, that is correct. However, modern immersion cooling systems rely on standard components in a data center, including power, connectivity, and, depending on the unit, a water loop.

   In traditional data center deployments, engineers can use the return water instead of having a chilled water loop, reducing PUE. New immersion cooling systems are engineered to need hookups similar to traditional infrastructure found in data centers while providing better benefits.

   Another critical point is that the data center industry is one of the last places where immersion cooling is being understood further. Interestingly, liquid cooling is utilized for other everyday use cases most people don't know or forget about. Large-scale computing is one of the last energy-intensive operations not cooled by liquids. Everything from nuclear reactors, car engines, paper mills, and more are already cooled with liquids.

5. **Risk with liquids being so close to the cabinets.** There is a perception of risk with liquids in the data center where the liquid in the cabinet may cause an outage. However, the fluids used are dielectric, meaning they do not conduct electricity. Therefore, this liquid is safe for technology components. Also, the fluids used in single-phase immersion cooling do not leave residues.

Section 3 – Immersion Cooling Readiness Checklist

So far, we've covered the challenges in the data center industry that are pushing digital infrastructure leaders to examine immersion cooling solutions. We have also discussed a few myths that have been overcome with modern immersion cooling designs.

That said, let's explore an immersion cooling readiness checklist.

WHAT CAN BE SUBMERGED?

The first step in any immersion cooling project is understanding what from your environment can be submerged.

	Air	Immersion
Compute/Processing	Y	Y
Fabric (Copper, Single Mode Fiber)	Y	Y
Network Switching	Y	Y
Solid State Storage and NVMe Drives	Y	Y
Spinning Media	Y	Spinning media, if sealed (HDD)
Power Supplies	Y	Y

When working with immersion cooling, it's important to note that there are different types to consider. For the sake of brevity, we'll focus on direct immersion cooling, and let's define that briefly.

> **Direct Liquid Cooling (immersion).** An immersion cooling system directly connects the heat-producing components with a cooling liquid. Heat travels directly to the liquid through convection. Heat can be removed from the coolant in several ways. The coolant can be pumped to a heat exchanger, then cooled and recirculated. Alternatively, heat can transfer passively from the liquid to the enclosure, with the enclosure cooling by natural convection to the surrounding air; this is how power transformers on utility power lines are cooled.

Within the immersion cooling category are two types to consider:

- ✔ 1. Single-phase immersion cooling
- ✔ 2. Two-phase immersion cooling.

SINGLE PHASE COOLING: Like its two-phase counterpart, heat is transferred to the coolant through direct contact with server components. But, **unlike two-phase immersion cooling, the coolant does not boil off**. Instead, it remains in the liquid phase (hence the name) and is cooled via a heat exchanger in a cooling distribution unit (CDU). With the help of a heat exchanger, heat is transferred to a cooler water circuit. *The technique guarantees little to no risk of the coolant evaporating.*

Single-Phase Immersion Cooling

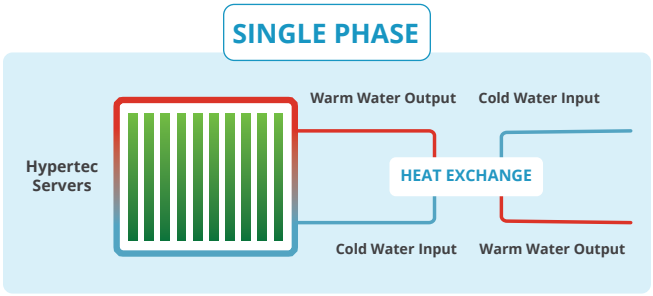

TWO-PHASE COOLING: In a two-phase immersion cooled system, electronic components are submerged into a dielectric heat-transfer liquid bath, a much better heat conductor than air, water, or oil.

Two-Phase Immersion Cooling

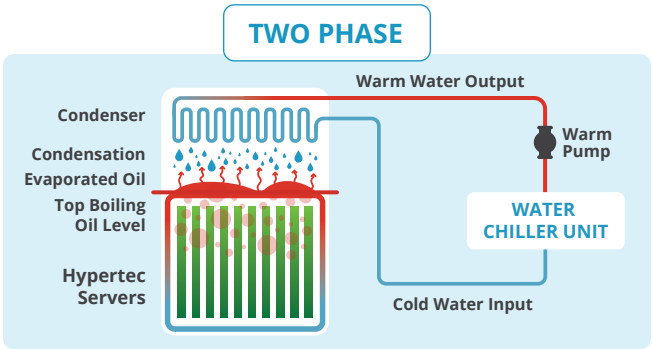

HYPERTEC'S CHOICE: *Picking the right solution* will revolve around your use case. However, following an assessment of single-phase and two-phase immersion systems, immersion cooling like Hypertec has decided to focus its roadmap on single-phase immersion technology entirely. The following table summarizes Hypertec's comparison of the two technologies during their assessment.

CRITERIA	SINGLE PHASE	TWO PHASE
Power Usage Effectiveness (PUE)	✔	✔
Fluid Loss	✔	✖
Toxicity of FLuid	✔	✖
Biodegradability of Fluid	✔	✖
Material Compatibility	✔	✔
Cost	✔	✖
Maintenance	✔	✖

The last part of this checklist is to understand if you are immersion-cooling ready. Before diving in, working with a good partner who can assess your current data center infrastructure and how it can be adapted to support immersion cooling is essential. To get there, consider the following:

How to become immersion ready

1. **It'll be critical for you to know your workload and your use case.** The business value must meet technology capability to gain the most out of immersion cooling. There are many use cases where immersion cooling can make an impact. However, the type of servers, tanks, and overall architecture might differ. You will need to understand your market and the direction of your customer. If your business works more closely with data-intensive or GPU-intensive workloads, immersion cooling would be a good consideration. Remember, optimal workloads for immersion cooling include crypto-mining, visual and graphics processing, GPU-based workloads requiring constant computation, AI and ML workloads, and data science.

   As an essential point, the rise of ChatGPT has also highlighted immersion cooling. Training long-language models (LLM) and ChatGPT-like bots are very resource intensive. A single Google search can power a 100w lightbulb for 11 seconds, and a single ChatGPT session is between 50-100 times more powerful than that. Immersion cooling is a great tool to dissipate GPU heat required to operate new workloads like ChatGPT 3.5 and ChatGPT 4.

2. **Do the data center math.** A crucial part of becoming immersion ready is the environmental calculations of the workloads, the BTU, and how many physical nodes can enter the tank. For example, you can place 144 servers in one 48U rack at about 100Kw per rack.

   Let's stay on this point for another minute. We discussed this in Section 2, but single-phase immersion cooling offers several savings benefits. When deployed, immersion cooling like those from Hypertec can reduce cooling OpEx by 95% with a Power Usage Effectiveness of 1.03 (certified by a 3rd party). For these systems, an ROI of less than one year is accomplished,

even considering the electricity savings. When used with immersion-born servers such as the Ciara Trident, you will reduce the amount of compute and storage required, lowering the total cost of ownership by up to 33%.

Now let's look at some real math. Using an Immersion Cooling Savings Calculator, we see that for a facility using about 1MW of power, paying about $.10 per kWh, annual savings come out to about $500K annually. These calculations are based on average indirect adiabatic cooling versus immersion cooling data center PUE. Because the PUE of immersion cooling is so low, data centers of all sizes can see direct savings from single-phase immersion cooling.

3. **You'll need to do an architectural review.** You'll need to determine if your facility can support immersion cooling or if you'll need to retrofit. For data center operators, it's easier if you own the building. However, you can also work with your colocation partner to see if immersion cooling will work for you.

4. **Working with the right partner.** The right partner will go a long way in helping you become immersion cooling ready. That said, you don't need to become an immersion cooling expert to have a solution deployed. A leading partner can help you create a white-glove service to manage your immersion cooling ecosystem. This includes site assessments, installation planning, tank layout design, factory assembly, software and OS installation, onsite installation, and, very importantly, cabling and labeling.

With all of this in mind, let's wrap up but looking at a few immersion cooling solutions and how to get started on the journey.

Section 4 – The Immersion Cooling Journey: Innovation in Infrastructure

Before you jump on the immersion cooling bandwagon, it's essential to understand your use cases and where immersion cooling can make a powerful impact. Remember, depending on the component and the solution, a liquid can remove only a portion of the heat. That said, not all equipment will or should be immersion-cooled. As mentioned earlier, the exciting part is that immersion cooling can create an all-encompassing solution where every component in your design is adequately cooled. So, designing around your use case is critical.

For those organizations looking to leverage immersion-born technology, it's vital to look at immersion servers currently leading the market. For example, The Ciara Trident line of servers was the first to be developed specifically for single-phase immersion cooling. These servers offer ultra-high-density compute performance, up to 288 CPUs per 48U (144 nodes/18,432 cores). They are engineered to accommodate next-generation CPU, GPU, and FPGA cards.

Staying on the topic of servers, we also need to consider sustainability. The Ciara Trident server line is designed and manufactured for optimized liquid flow within single-phase immersion cooling, decreasing energy and water consumption and maximizing heat capture. The architecture leads to a reduction in carbon emissions by up to 40%. In addition, Hypertec will be using recycled materials to build the next-generation Ciara Trident systems, further reducing its carbon footprint.

Finally, immersion and air cooling will not be mutually exclusive in designing a good solution, and this means working with the right partners to design your ecosystem will be critical. Whether for HPC, rendering, edge computing analytics, AI, machine learning, or deep learning, immersion cooling has helped all these use cases thrive in the business world.

Data center leaders have also felt the impact of immersion cooling. "Partnering with Hypertec was a logical next step in offering our customers even more sustainable, cost-saving energy reducing from immersion cooling," said Patrick Quirk, CTO at Nautilus Data Technologies. "The Hypertec immersion cooling solutions open the door for Nautilus customers to push the limits of increased Kilowatts for unrivaled power without increased electricity consumption."

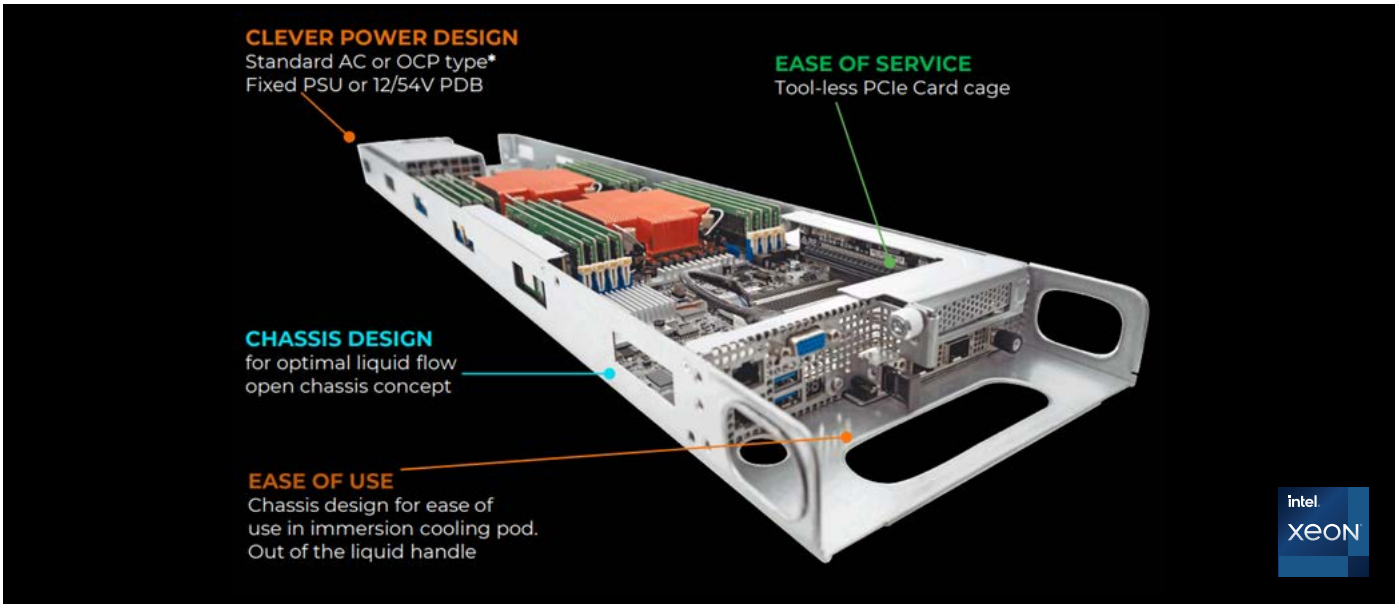

FINAL THOUGHTS

Data center technologies will only continue to evolve, and there is a certainty that data and data-driven systems will become the primary workload for many organizations. To support these use cases, leaders must rethink how they deploy critical infrastructure and what they use to deliver capacity, density, and efficiency.

We've covered many topics in this paper relating to single-phase immersion cooling. This includes the following:

- Growing data center energy consumption metrics
- Evolving workloads including, HPC, AI, machine learning, deep learning, and more
- Working with legacy infrastructure and how to retrofit
- Critical differences between two-phase and single-phase immersion cooling solutions
- How to become immersion cooling ready
- Innovation in infrastructure and immersion cooling

To start your journey, it would be best to ask reflective questions about your business and current infrastructure designs. If you aim to work with high-density, data-driven workloads, immersion cooling is optimal for business and technological advantages. And when you can get to the data faster, you can leverage technology to better position your business. Immersion cooling can be a crucial driver for technology and business.