

Discovering the Immersion Advantage for AI

**An optimal approach to building
next-gen infrastructures**

Special Report

Table of contents

EXECUTIVE SUMMARY	3
INTRODUCTION	4
THE TRUTH ABOUT AIR COOLING	6
Power usage effectiveness (PUE)	7
Water usage effectiveness (WUE)	7
Environmental, social, and governance (esg)	7
Cost control	7
EMERGING COOLING FOR AI	8
Rear-door heat exchangers	9
Direct-to-chip	9
Immersion cooling	10
DEEP DIVE INTO IMMERSION COOLING	12
It offers better use of energy capacity	12
It eliminates empty space between servers	13
It removes fan and blower noise entirely	13
It cuts component failure rates	13
It increases hardware lifespan	14
It supports multiple generations of processors without retrofit	14
It allows ai infrastructure to live outside the data center	14
THE ROI OF IMMERSION COOLING	16
Immersion cooling compares favorably to other liquid cooling technologies	16
Immersion cooling cuts operating costs now and in the future	16
Immersion cooling has long-term value	17
Immersion cooling is environmentally friendly	17
CONCLUSION	18
REFERENCES	21

Summary

Today's explosive AI growth puts unprecedented pressure on data center operators.

In fewer than two years, operators who were able to power and cool no more than 10 kW per rack are being confronted with servers that consume **up to 10.2 kW** — **each** — and know that the industry is on track to consume **2 kW per socket in a short time**. This massive change in server design, built to support AI, forces data center industry leaders, engineers, and administrators to radically rethink the technologies they'll need to succeed, whether or not they like it.

To cope with this shift, organizations are being forced to rethink their cooling technologies. The days of CRACs and CRAHs are beginning to fade away. Part of the reason is that air cooling can only safely cool up to 40 kW of servers per rack, and AI infrastructure demands densities of 80kW or above. But there's another, less obvious reason why organizations are rethinking their cooling technologies. According to the [Uptime Institute](#), improving energy efficiency has become a primary concern, and air cooling is notorious for wasting energy. Data center operators of all sizes need to improve their energy efficiency -- and alternative cooling technologies are part of their approach.

As you read this paper, you'll explore the impact of AI on data centers today. You'll dig deep into the technologies that are putting pressure on data center operators, explore alternatives to air cooling, and wrap up with an evaluation of why one particular cooling technology is an optimal technology for leading-edge AI infrastructure.

After you finish reading this paper, you'll be on the forefront of understanding cooling technologies for next-generation data centers, and you'll know more about an organization that could become a critical partner for your AI ambitions.



Introduction

As anyone in the data center industry knows, the arrival of AI sent shockwaves through our industry, overturning key assumptions about data center design and data center economics.

Supporting hundreds of use cases, it's infiltrating [numerous industries](#) despite requiring infrastructure more akin to a supercomputer than traditional cloud infrastructure. Built with hundreds or thousands of high-performance GPUs, CPUs, flash drives, and intensive networking infrastructure, AI training clusters are simply an architecture that most data centers can't support with legacy data center designs and technologies.

Because AI demands the most powerful compute infrastructure possible, there's a sort of arms race between many of the key players -- from colocation companies to hyperscalers to silicon suppliers -- to create the next new, most innovative AI infrastructure to power their AI ambitions -- whether that's solving for a single use case, or cornering the market.

However, there's a problem blocking the path toward AI innovation. AI brings a perfect storm of data center energy consumption...and heat generation.

Why?

A few shifts in technology are the root causes of this storm.

To address the demanding computational needs of AI, data centers are witnessing a rise in the adoption of specialized processors known as xPUs (e.g., GPUs, TPUs). Instead of the standard cloud architecture with servers that contain one or two x86 processors, servers to power AI contain x86 processors and several xPUs. xPUs excel at parallel processing tasks, significantly accelerating AI training and inference compared to traditional CPUs.

To reach desired performance levels, AI servers have four, six, eight, or ten processors -- one or two CPUs, and several xPUs. They also require masses of memory, flash drives, and extensive networking infrastructure that goes far beyond any infrastructure designs we've seen, other than supercomputers.



If you contrast this picture with the relatively limited silicon, memory, storage and networking used in cloud infrastructure servers, you can begin to see how the aggressive pursuit of performance improvements leads to a continuous increase in server heat generation. More components = more heat.

To make matters worse, for some time, silicon designers, in lockstep with generation-on-generation performance gains, had been forced to consume more power. For years, as each new generation of processor is launched, it comes with increased power consumption and increased Thermal Design Power (TDP) -- the maximum amount of heat a processor is designed to dissipate.


Not that long ago, a typical x86 processor had a TDP of 300W and a GPU, 150W. What's the current situation? Here's a chart of thermal design power (TDP) specifications for some of the most popular AI GPUs.

NAME	TDP
Intel Gaudi 2	600W
AMD Instinct MI300	600W
AMD Instinct MI300A	Peak of 760W
NVIDIA H200	700W

Can you imagine a server with ten 700W processors -- a 7 kW server? Amazingly, that processor's TDP and server's power consumption is almost antiquated. The next generation of Intel CPU/GPU processors [will consume 1500W](#). The latest NVIDIA DGX B200 server demands up to [14.3 kW](#) of power.

This explosive increase in processor count, coupled with increasing TDP, is a perfect storm for data center thermal management. Tens of thousands of data centers have been designed to cope with 10-15kW of power and heat per rack. Today, they're being asked to support 10-14 kW per server -- and anticipating increasing density and TDP with no end in sight.

This shift exposes a fundamental challenge: conventional air cooling simply cannot keep up.



The truth about air cooling

It's true that for decades, air cooling has been the mainstay for data center thermal management, packing six, eight, even ten high-TDP processors together in a 2U server creates thermal bottlenecks that traditional **air cooling** struggles to address. Air-based cooling systems also struggle to keep pace with the ever-increasing heat generated by the proliferation of high-TDP components, raising concerns about thermal throttling and potential premature equipment failure.

To make matters worse, air cooling is inherently inefficient, and increasingly, as organizations are asked to do more with less, resource utilization becomes a top priority. Data center operators are looking for ways to be laser-focused on operational efficiency, sustainability, and cost control. They, their customers, and their investors are looking for ways to improve:

Power Usage Effectiveness (PUE)

Power Usage Effectiveness is a measure of energy efficiency. Though flawed, it's the best measure the industry has, and the majority of data centers struggle with sub-optimal PUE, leading to higher energy costs and environmental impact. According to research, the average PUE is 1.58, which means that, for every 1000W spent for computing, 580W are wasted. Air cooling is one of the critical causes of sub-optimal PUE. Energy spent on air cooling, whether through fans, blowers, or compressors, is energy not spent on compute capacity.

Water Usage Effectiveness (WUE)

Refrigerant-based air cooling is expensive, leading data center providers to adopt evaporative cooling systems, which can consume vast amounts of water while polluting discharge water with contaminants. Water waste is a practical issue, due to limited water supplies, as well as a political issue. Data center providers need designs that reduce or eliminate water usage entirely.

Environmental, Social, and Governance (ESG)

Air-cooling refrigerants are under regulatory scrutiny while at the same time, organizations are increasingly focused on improving their ESG profiles by cutting wasteful operations and their carbon footprints. More efficient data centers are a must for companies seeking to minimize their environmental impact.

Cost Control

As demand for AI-ready data centers rises, so do costs associated with power and cooling retrofits in older facilities as well as construction costs for new facilities. Organizations need technologies that reduce disruption and help them deliver AI infrastructure as quickly as possible while maximizing cost efficiencies.

Emerging cooling for AI


For these reasons, leaders throughout the industry believe that, for AI infrastructure at least, [air-cooling technologies are becoming obsolete](#), and data centers that rely on them are also becoming obsolete. Leading organizations, from Meta to Amazon, have been forced to [radically rethink their data center designs](#) to keep up with the leading edge of AI infrastructure.

How are they rethinking their designs?

Liquid cooling. It's a necessity for building and utilizing AI infrastructure at scale because it offers the correct mix of cooling capacity, environmental advantages, and cost benefits. Proven for decades with supercomputers, the growth of AI means that it's rapidly becoming a mainstream technology that's displacing air cooling. Unlike air, liquid cooling can support 50kW, 100kW, and even 150kW racks, at scale, with substantial efficiency, cost, and sustainability advantages. Market analysis reports that the market for liquid cooling in data centers is growing at an [over 25% CAGR](#) because it's in demand for all sorts of intensive data center infrastructure, including AI infrastructure.

Liquid cooling technologies come in a range of forms. Some are used in conjunction with air cooling, others replace air cooling. Liquid cooling technologies are an emerging development in the data center industry, and there's fierce competition among vendors, all aiming to prove themselves as the one solution for AI data center infrastructure.

Overall, the industry agrees that liquid cooling is the future. Vendors from [Intel](#) to [NVIDIA](#) are building chips requiring liquid cooling. Leading colocation companies like [Digital Realty](#) and [CyrusOne](#) are rapidly implementing different forms of liquid cooling. Server manufacturers, including Hypertec, support liquid cooling. There's a convergence of agreement in the industry -- and it's important to understand the pros and cons of different technologies. Let's explore three primary technologies used to cool AI infrastructure.






Rear-door heat exchangers

Datacenter rear-door heat exchangers (RDHx) are liquid cooling devices, installed at the back of server racks, directly removing heat from the exhaust airflow. They offer several advantages, including improved energy efficiency by supporting higher chilled water temperatures, potentially eliminating chillers. They can support the higher cooling capacities needed for hot and dense AI computing loads. They come with reduced maintenance compared to CRAC units, and they're easily scalable to incrementally add cooling where needed, reducing the need for larger retrofits. However, RDHx systems require additional plumbing infrastructure, may not be cost-effective for lower rack densities, and can have air leakage issues, reducing efficiency, if not properly sealed. While beneficial for high-density environments as a bridge to using liquid to cool a rack or two, their long-term, at-scale suitability depends on factors like data center layout, coolant distribution infrastructure, and projected future capacity needs.

Direct-to-chip

These products circulate a liquid coolant directly to processors and heat-generating components within servers, enabling superior heat transfer in air-cooled environments. By removing heat at the source through cold plates or micro-channels, this technology allows higher processor densities and frequencies without thermal throttling, cuts cooling energy consumption, enables tighter server packing densities, and extends component lifespans. However, direct-to-chip cooling comes with complexities and risks. Servers could be damaged by leaks. Cooling multiple CPUs, memory chips, and GPUs in a single chassis introduces complex coolant distribution pipes to all the components, increasing failure points. Many solutions require very high water quality. Finally, any direct-to-chip products require future R&D engineering expenditures which must be passed on to customers, as each new chiplet requires a unique cooling plate design. As the use of direct-to-chip cooling scales from dozens of servers to hundreds or thousands of servers in a data center, risks increase substantially.



Immersion cooling

Finally, we come to immersion cooling. Though it's been a niche solution, used principally in high-performance computing environments, including supercomputers, it's rapidly becoming a widely adopted solution for AI infrastructure because of its inherent simplicity and value. Immersion cooling is a different paradigm for data center cooling. Unlike other technologies, it doesn't rely on air cooling at all. Immersion cooling relies on a dielectric fluid, either single-phase or two-phase, which surrounds all the computing infrastructure and is

1

The tanks are filled with dielectric fluid, and the IT equipment is immersed in the fluid.

2

The tanks have a closed-loop cooling system that pumps hot liquid coolant to a heat exchanger, which can dissipate heat into the air or water.



in constant contact with every component. Immersion-cooled data centers are built with tanks, typically constructed from corrosion-resistant materials like stainless steel or specialized polymers.

3

The tanks often incorporate monitoring systems to track parameters like liquid temperature, flow rates, and component performance.

At first glance, immersion cooling is certainly distinctive compared with other technologies.

How is it relevant for AI infrastructure, and does it offer advantages over other liquid cooling systems?



Deep dive into immersion cooling

Unlike other liquid cooling technologies, Immersion cooling for servers offers the ability to directly immerse all components, including CPUs, GPUs, memory modules, and other heat-generating components, in a dielectric liquid coolant. This direct contact between the components and the liquid coolant enables highly efficient heat removal, a significant improvement over other cooling methods because it eliminates the need for intermediate heat transfer mechanisms, including heat sinks, heat exchangers, and fans, which, for hot and dense infrastructure, can contribute to thermal bottlenecks and inefficiencies. The liquid coolant used in immersion-cooled environments effectively removes heat from even the most densely packed and high-power components, enabling higher computing densities and potentially higher clock speeds without the risk of overheating or thermal throttling.

Immersion cooling unlocks several other critical advantages.

It offers better use of energy capacity

Immersion cooling offers superior energy efficiency. By directly submerging servers in a dielectric liquid, heat is dissipated more effectively. This translates into better utilization of available energy capacity, as less power is wasted on cooling. Immersion allows data centers to cut their power utilization for the same amount of compute capacity, or increase compute capacity within a limited power envelope. These efficiencies maximize return on energy expenditure.





It eliminates empty space between servers

Traditional air-cooled data centers often suffer from an inefficient use of space. Many facilities cannot support AI compute infrastructure at scale because of cooling limitations, forcing them to spread servers across many racks, leaving empty spaces. Conventional data centers also waste space due to the need for hot and cold aisles as well as clearances for airflow. Immersion cooling eliminates these limitations and constraints, enabling servers to be packed more densely within the same footprint. This higher density translates into increased computing power per square foot, making it an ideal solution for AI workloads that require massive computational resources or in situations where space is limited. By maximizing the utilization of available space, immersion cooling also allows data centers to scale their capabilities more efficiently, reducing the need for costly expansions or new facilities.

It removes fan and blower noise entirely

Conventional air-cooled data centers can be notoriously noisy due to the fans and blowers required for cooling. It's normal for data centers to reach sound levels of over [90 decibels \(dbA\)](#). This noise not only creates an unpleasant working environment, but also contributes to increased stress levels and potential hearing damage for personnel. Immersion cooling, on the other hand, operates almost silently, as the liquid coolant dissipates heat without the need for mechanical air movement. This quieter operation not only improves the working conditions for data center staff but also reduces the potential for noise-related issues that could impact sensitive equipment or neighboring areas.

It cuts component failure rates

Environmental parameters are important within data centers, high temperatures, humidity, and particulate matter can contribute to increased component failure rates. Immersion cooling mitigates these risks by providing a controlled, isolated environment for servers, shielding them from external contaminants and reducing thermal stress. By maintaining optimal operating temperatures and eliminating exposure to moisture and particulates, immersion cooling can significantly reduce component failure rates, leading to improved reliability and uptime.

It increases hardware lifespan

The combination of reduced thermal stress and a controlled environment in immersion cooling systems can extend the useful lifespan of hardware components. By maintaining optimal operating temperatures and minimizing exposure to contaminants, immersion cooling can slow down the degradation of components, allowing them to operate reliably for longer periods. This extended lifespan translates into reduced refresh rates, lowering the overall total cost of ownership (TCO) and minimizing the environmental impact associated with frequent hardware replacements.


It supports multiple generations of processors without retrofit


As AI workloads continue to evolve and demand ever-increasing computational power, immersion cooling offers a futureproof solution. Unlike traditional air cooling systems or liquid cooling technologies like D2C that may require extensive retrofitting or replacements to accommodate new generations of processors, immersion cooling systems can adapt more readily to changing hardware requirements. The liquid coolant can effectively dissipate heat from a wide range of processor architectures and power envelopes, allowing data centers to seamlessly integrate new hardware without the need for costly infrastructure overhauls.

It allows AI infrastructure to live outside the data center

Unlike air-cooled and liquid-assisted air-cooled hardware environments that require massive infrastructure including chillers and coolant distribution, immersion cooling enables organizations to rearchitect where their infrastructure resides. It's becoming commonplace for organizations to deploy new AI infrastructure using **mobile, modular containers** that are cooled with immersion cooling. Others are able to deploy immersion-cooled data center capacity in everything from telco huts to warehouses without worrying about the functional, environmental, and operational constraints posed by conventional cooling. These immersion-cooled deployments, whether deployed in difficult climates, at the edge of a major city, or in a colocation facility's parking lot, give technology companies a new way to rapidly implement AI infrastructure without the cost and complexity of retrofitting existing space.

In summary, immersion cooling's ability to support higher densities, improve cooling efficiency, extend hardware lifespan, reduce noise, provide futureproofing, and provide site options makes it a compelling solution for building the high-performance, scalable infrastructure required for AI workloads. And market leaders agree. [Intel](#), [Microsoft](#), and others like the [Open Compute Project](#) have validated immersion cooling for AI infrastructure.





The ROI of immersion cooling

But does it deliver measurable ROI?

It does, in four critical ways.

01

Immersion cooling compares favorably to other liquid cooling technologies.


While D2C cooling offers excellent thermal management for high-performance computing, it comes with substantial upfront and ongoing costs. The intricate nature of D2C systems requires specialized infrastructure and ongoing maintenance. Subsequent generations of server hardware come with the added expenses of direct-to-chip cooling components. Immersion cooling, on the other hand, utilizes simpler and more readily available components and doesn't add costs to the servers -- in fact, it may reduce the costs of the servers by eliminating unnecessary components like fans and chassis enclosures.

02

Immersion cooling cuts operating costs now and in the future

Operating costs are a substantial expense for any data center operator, and for many, costs are already spiraling out of control due to demands for new infrastructure and macroeconomic impacts like inflation. Immersion cooling helps to mitigate increasing operational costs:

- a.** Immersion cooling eliminates the need for costly and power-hungry air handling units (AHUs), expensive and heavily regulated refrigerants, and complex ductwork that requires maintenance.
- b.** Instead, by directly transferring heat from servers to coolant, it reduces cooling infrastructure complexity and saves energy. Lower electricity bills and fewer points of failure cut operational costs.
- c.** Also, by keeping servers cooler and more consistently cooled, immersion cooling reduces thermal stress and wear. This translates to a longer lifespan for expensive AI hardware, lowering replacement costs.
- d.** To make matters even better, servers submerged in dielectric fluid are shielded from dust and moisture, common culprits in hardware failure. This translates to less downtime for maintenance and repairs, leading to improved operational cost efficiency.
- e.** Finally, immersion cooling allows for closer packing of servers due to its superior heat removal capacity. This translates to maximizing data center floor space for AI workloads, opening up opportunities for alternate site selection, reducing real estate costs, and improving value/square meter.



03

Immersion cooling has long-term value

Immersion cooling offers a superior solution for long-term value. Rather than having to be redesigned for every new processor design, immersion cooling works with all manner of processors and other components. It provides data center operators with a futureproofed solution that can support many generations of servers, even the hottest ones.

04

Immersion cooling is environmentally friendly

Facilities want to improve their sustainability profile, for ethical as well as economic reasons. The space reductions, energy utilization reductions, and reduced materials costs in immersion-cooled environments translate into operations with a reduced carbon footprint. Immersion cooling can deliver a PUE as low as 1.02, and single-phase immersion cooling can utilize [biofriendly, renewable cooling fluids](#) that are green and safe, reducing future costs of environmental remediation. These advantages help data center operators avoid the hard and soft costs of regulatory scrutiny.



CONCLUSION

AI is breaking down old data center paradigms because, to put it simply, it can't be done at scale with legacy cooling technologies.

Immersion cooling is the future of the data center, and at Hypertec, our immersion-born approach to building compute infrastructure means that we're at the forefront of this change.





Why did Hypertec become an immersion specialist?

Years ago, even before the explosive growth of demand for AI, we recognized that processors were getting **hotter than ever before**, pressures from regulators and ESG initiatives were making it harder and more expensive to use **refrigerant-based air cooling at scale**, and organizations wanted cooling alternatives that brought **superior performance, resilience, sustainability, and value**.

Immersion was the logical solution to all those challenges. And then OpenAI and others pressed the accelerator on the growth of hot and dense infrastructure that simply wasn't feasibly cooled by air.

At Hypertec, we are immersion-cooled first and only. All of our R&D goes into improving the performance, reliability, and value of immersion-born servers. Having recently launched our **6th generation**, we have a great track record of building hardware exclusively for immersion environments. That track record has translated to customer deployments, across a wide range of environments and use cases, that deliver value others can't match.

The shift to immersion cooling results in over 50% power savings, significantly cutting down on energy consumption and reducing the carbon footprint. Additionally, immersion cooling dramatically reduces the space required; two immersion cooling tanks can replace up to 20 traditional server racks. This translates to a 60% reduction in physical space and a significant decrease in operational costs.

Moreover, the substantial space savings from immersion cooling free up valuable floor space for additional infrastructure or other critical operations. The reduced space and power requirements enhance overall efficiency, enabling data centers to handle over five times more compute power without the need for additional CRAC units, as the air heat load is eliminated.

Furthermore, GPUs typically do not operate at their [fullest capacity due to overheating](#), often caused by inadequate airflow in the computer case or malfunctioning fans. Without proper ventilation, the GPU cannot cool down efficiently. In contrast, Immersion cooling absorbs heat more effectively than air. It eliminates hotspots, reduces the thermal resistance between the GPU and the cooling medium, and ensures uniform temperature distribution. As a result, GPUs operate to their full 100% capacity, delivering a remarkable and nonstop performance.

Looking ahead, the evolution of next-generation silicon will pose increasing challenges for data center managers. As these advanced processors demand higher power densities and generate more heat, traditional cooling methods will become increasingly inadequate. There is an urgent need to explore and adopt innovative cooling solutions like immersion cooling to address these challenges effectively. By considering and implementing such options now, data centers can stay ahead of the curve, ensuring they are well-prepared to handle the escalating demands of future technologies.

In summary, immersion cooling stands as a transformative solution in the data center landscape. By offering energy efficiency, space optimization, and hardware performance, it paves the way for future-ready data centers. Embracing this technology will not only increase computational capabilities but also contribute significantly to sustainability efforts.

To learn more about the Hypertec difference, visit us at hypertec.com



References



- Donnellan, D. (n.d.). *Executive summary: Uptime Institute Global Data Center Survey 2023*.
- Future of Data Center cooling*. (n.d.). Digital Realty. Retrieved June 4, 2024, from <https://www.digitalrealty.com/resources/articles/future-of-data-center-cooling>
- Intel dives into the future of cooling*. (n.d.). Intel. Retrieved June 4, 2024, from <https://www.intel.com/content/www/us/en/newsroom/news/intel-dives-into-future-of-cooling.html>
- Miller, R. (n.d.). *Meta Plans Shift to Liquid Cooling for its Data Center Infrastructure*. Data Center Frontier. Retrieved June 4, 2024, from <https://www.datacenterfrontier.com/cooling/article/11436915/meta-plans-shift-to-liquid-cooling-for-its-data-center-infrastructure>
- Moss, S. (n.d.). *Nvidia's CEO confirms upcoming system will be liquid cooled*. DCD. Retrieved June 4, 2024, from <https://www.datacenterdynamics.com/en/news/nvidias-ceo-confirms-next-dgx-will-be-liquid-cooled/>
- NVIDIA DGX B200 datasheet*. (n.d.). NVIDIA. Retrieved June 4, 2024, from <https://resources.nvidia.com/en-us-dgx-systems/dgx-b200-datasheet>
- Patrizio, A. (2023, November 16). *High CPU temps are here to stay*. *Network World*. <https://www.networkworld.com/article/1247693/high-cpu-temps-are-here-to-stay.html>
- Products*. (n.d.). Hypertec. Retrieved June 4, 2024, from <https://hypertec.com/immersion-cooling/>
- Project, O. C. (2023). *Investigation of immersion cooling for data center hardware: Progress, opportunities and challenges* [Video]. In YouTube. <https://www.youtube.com/watch?v=M7y1c2PQLlw>
- Thompson, S. (2024, May 19). *Intel's new Falcon shores GPU to draw massive 1500W power*. *eTeknix*. <https://www.eteknix.com/intels-new-falcon-shores-gpu-to-draw-massive-1500w-power/>
- Ware, A. (2024, February 19). *Intel advances liquid cooling solutions for datacenters, saving water and energy*. *HPCwire*. <https://www.hpcwire.com/2024/02/19/intel-advances-liquid-cooling-solutions-for-datacenters-saving-water-and-energy/>
- Why data centers could hit obsolescence sooner than you think*. (2024, March 20). <https://www.us.jill.com/en/trends-and-insights/investor/why-data-centers-could-hit-obsolescence-sooner-than-you-think>
- Zissou: New datacenter, server, and software architectures for immersion-cooled systems*. (2021, June 28). Microsoft Research. <https://www.microsoft.com/en-us/research/project/zissou-new-datacenter-server-and-software-architectures-for-immersion-cooled-systems/>
- (N.d.). Retrieved June 4, 2024, from <https://www.cyrusone.com/resources/blogs/latest-liquid-cooling-tech-makes-solid-impact-at-cyrusones-lebanon-facility>
- CE. (n.d.). Retrieved June 4, 2024, from <https://www.carrierenterprise.com/hvac-news/what-to-know-about-the-r410-refrigerant-phase-out>
- Catherine. (2023, June 29). *Hypertec Launches TRIDENT iC Series 6th Gen Immersion-Born Servers*. *Hypertec*. <https://hypertec.com/hypertec-launches-new-ciara-trident-ic-series-immersion-born-servers/>
- Akshat V., November 18, 2023, *Fix Graphics Card Overheating* [High GPU, VRAM, VRM Temperature], <https://graphicscardhub.com/graphics-card-overheating/>



Special Report Immersion Cooling



Get In Touch

-  + 1 514.745.4540
-  hypertec.com/immersion-cooling/
-  9300, Route Trans Canada
St-Laurent, Québec, H4S 1K5
Canada