


BLOG LIBRARY:
TECH COMPARISON

Data Center Cold Wars – Part 3

Cold Plate Versus Single-Phase Immersion Cooling



Submitted by Dhruv Varma, Director of Product Marketing



Something of Note

When we talk immersion cooling, we talk **single-phase immersion cooling, not 2-phase**. There's a world of difference in the mechanics and economics of each.

[Find Out More](#)

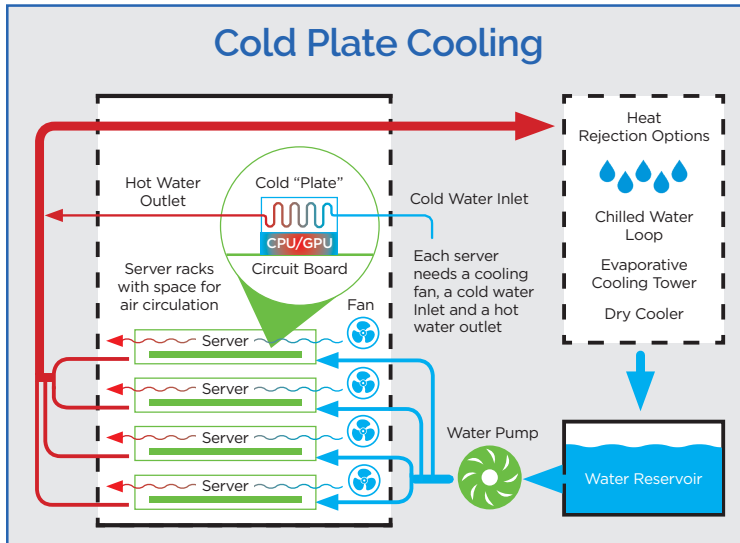
Cooling a data center never used to be this hard. But IT and data center professionals have watched the thermal design power (TDP) of chips rise almost 50% in the last decade, generating more heat and using more power than ever before. Rack density has grown. And hot GPUs are becoming the weapon of choice for tackling high-performance computing (HPC) requirements.

This third in our series of blogs compares the two most emergent technologies designed to cool data centers: cold plate and single-phase liquid immersion cooling. Hopefully, these insights will help you understand the technological differences and choose the best solution for your operation.

How These Two Methods Compare with Each Other		
	Cold Plate	Single-Phase
Complexity	★ ★ ★ ★ ★	★ ★ ★ ★ ★
Upfront Costs	★ ★ ★ ★ ★	★ ★ ★ ★ ★
Efficiency	★ ★ ★ ★ ★	★ ★ ★ ★ ★
Operating Expenses	★ ★ ★ ★ ★	★ ★ ★ ★ ★
Cooling Capacity	★ ★ ★ ★ ★	★ ★ ★ ★ ★
High-Density Performance	★ ★ ★ ★ ★	★ ★ ★ ★ ★
Hardware Reliability	★ ★ ★ ★ ★	★ ★ ★ ★ ★
Location Flexibility	★ ★ ★ ★ ★	★ ★ ★ ★ ★

★ - Poor ★ ★ ★ ★ ★ - Excellent

Matchup: How These Competing Cooling Technologies Work

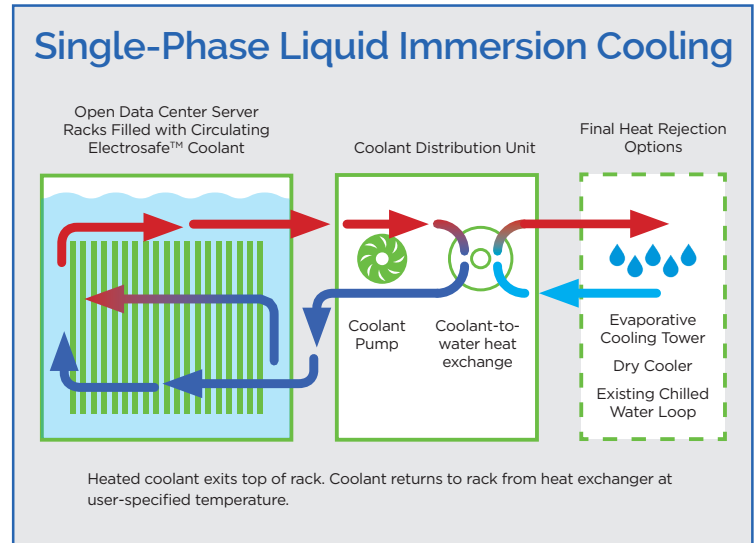


Cold Plate Cooling

Basically, "touch," "direct-to-chip", or "liquid-to-chip" cold plate cooling replaces older, inefficient air-cooled metal-finned heat sinks with liquid-cooled heat sinks. As the name suggests, it attaches a metal plate atop a CPU or GPU, which transfers heat through a heat-spreading material (such as thermal paste) from the chip to the plate. The plate retains the ability to absorb this heat because it is cooled with liquid, which is a much better conductor of heat than air.

The heated liquid circulates from the plate through a coolant distribution unit (CDU) to the facility water loop. This may be connected to a chiller, or even cooling towers, then directed back to the plate. Since cold plate only cools the CPUs (which typically account for 60-70% of the total heat load), air-cooling is used to cool the remaining 30-40% heat load. This makes cold plate technology a hybrid solution that involves both liquid and fan-blown air cooling.

Cold plate is probably the most popular (and one of the oldest) forms of liquid cooling in the world of electronics and IT. Although most recognize the improved performance and efficiency it offers over air-cooling, the cost, complexity and potential risks involved have prevented large-scale adoption of liquid cooling.



Single-Phase Immersion Cooling

Single-phase immersion cooling offers the performance and efficiency benefits of cold plate without the added costs, complexity and risks. With single-phase immersion cooling, servers are installed vertically in a coolant bath of dielectric fluid. Like its two-phase counterpart, the coolant transfers heat through direct contact with server components. Heated coolant then exits the top of the rack and is circulated through a CDU connected to a warm-water loop. This loop incorporates a cooling tower or dry cooler on the other side as the final form of heat removal. In the end, cooled liquid is returned to the rack from a heat exchanger.

Many installations that have switched to single-phase immersion cooling have been impressed with its simplicity, which translates into greater upfront affordability, easier operations and less maintenance.



Compare, Contrast and Be Cool

Now let's break down each of these technologies and see how they compare across eight important categories.

Complexity & Upfront Costs

Cold Plate Cooling


When comparing cold plate cooling to other technologies, you should know that its basic liquid-cooling architecture is only 60% to 70% effective at dissipating heat. Chiller-based air cooling is still required to complete the solution.

Because of this, **data center infrastructure is a lot more complex with cold plate**, often exceeding that of traditional air-cooling. What's more, it often requires a chiller plus specialized engineering depending on the design of the servers. This all adds up to significantly greater upfront costs.

Further complicating things, each heat source (CPU/GPU) typically requires its own cold plate (custom heat sink) and the plumbing to go with it. This is because putting these in series would result in uneven cooling. **The result is often a rat's nest of custom heat sinks and plumbing crammed into a very limited space, requiring custom-made servers.** As an extreme example, ASIC (application-specific integrated circuit) boards can have thousands of chips in each brick-sized chassis. Even today's most basic enterprise server has two CPUs which require their own cold plates and plumbing. Add multiple GPUs to the mix and soon you have 10-20 hoses coming out of each server.

Cold plate technology is a hybrid solution that involves both liquid and fan-blown air cooling.

This fundamental complexity not only impacts upfront infrastructure and hardware costs. It also limits hardware choices and makes server refreshes (typically in three-to five-year cycles) much more involved and costly as well.



The area you're not typically shown by cold plate cooling suppliers. Heat-producing chips in each individual server have to be plumbed.



Example of today's densely packed ASIC miner boards. Each processor would need a water inlet and outlet.

Single-Phase Immersion Cooling

One of the hallmarks of single-phase liquid immersion cooling is that it is arguably the most fundamental form of modern data center cooling around – one that absorbs 100% of the heat load. This simplicity serves as the basis of immersion's main value propositions of reduced upfront and operating costs.

How simple is it? Data center systems using GRC's ICeraQ™ micro-modular immersion cooling solutions have only three moving parts: a coolant pump; a water pump; and a cooling tower or dry-cooler fan. Unlike cold-plate-cooled operations, there's no airflow engineering. And, GRC data centers have no need for expensive infrastructure, such as chillers, air handlers and raised floors.

But here's the number that really matters: **single-phase liquid immersion cooling can reduce data center CAPEX up to 30% over cold plate cooling.**

Which Can Take the Heat? Heat absorption ability of each method	
Cold Plate	70%
Single-Phase Immersion Cooling	100%

KEY Takeaway:

Because it employs both liquid-based and air-cooled systems, cold plate is more complex and costly than single-phase immersion cooling.

Efficiency & Operating Expenses

Cold Plate Cooling

As explained earlier, liquid-to-chip technology will get you (at most) three quarters of the way to a fully cooled data center. Because it needs expensive and inefficient air conditioning to take you the rest of the way, it is intrinsically less efficient than single-phase immersion cooling. Along with server fans, it also requires chillers, air handlers and other air conditioning equipment, which demand more electricity and regular maintenance, thus kicking up operating expenses.

Most modern cold plate systems deliver a PUE somewhere in the neighborhood of 1.15.


Single-Phase Immersion Cooling

Single-phase liquid immersion cooling couldn't be more efficient. Literally 100% of the heat is picked up by the coolant. That means there's zero residual heat requiring less efficient air-cooling.

Simple in design, with few moving parts, systems like GRC's [ICEraQ](#) and [ICEtank single-phase liquid immersion cooling systems](#) are 80% more energy-efficient than cold plate methods, delivering a PUE (PUE) of 1.03.² That delta between cold plate and single-phase immersion pPUE is amplified by the 10-30% server power reduction immersion cooling enables through fan removal. The end result equates to a total data center energy cut of up to 35%.

Since immersion cooling eliminates the need for all air-cooling infrastructure and equipment – such as chillers, air handlers or humidity control systems, the annual maintenance cost associated with those is also eliminated.

Minimal consumables and vastly reduced electricity usage are what help single-phase immersion cooling deliver a 40% cut in OPEX for most data centers.



Single-phase immersion cooling is up to 80% more energy-efficient than cold plate systems.¹

Single-phase liquid immersion cooling can cut data center OPEX by 50% compared to cold plate.

Less Is More

Compared to Cold Plate, Single-Phase Liquid Immersion Cooling offers:

30%	Less CAPEX
80%	Less Cooling Energy
10-30%	Less Server Power
35%	Less Total Data Center Energy
40%	Less Total OPEX (including maintenance costs)

KEY Takeaway:

Cold plate cooling has a higher PUE and is more expensive to operate.

Cooling Capacity & High-Density Performance

Cold Plate Cooling

Unlike the dielectric coolant used in single-phase immersion cooling, direct-to-chip or liquid-to-chip methods typically employ water or a water/glycol mix. These offer higher thermal conductivity than dielectric coolants, but are good conductors of electricity, nevertheless. Higher conductivity theoretically helps support higher heat flux. Yet this has little to no impact on the power density cold plate can support on a per-server or per-rack basis.

Since the electrically conductive coolant needs to be contained within heat sinks and plumbed individually to chips, it is at times physically impossible to fit and plumb multiple cold plates to every CPU and GPU within a high-density server.

Even though cold plate solutions are great for cooling isolated and localized hot spots, they don't do so well with multiple heat-producing chips within confined spaces. Plus, given the fact that there are multiple heat-producing chips within each server – and multiple servers in a rack, cold plate solutions can get very complicated at scale.

Immersion cooling, on the other hand, uses a dielectric coolant that is a good conductor of heat but not electricity. Therefore, the coolant can directly contact all components within a whole rack of servers to capture 100% of the heat and cool each chip effectively.

For these reasons – and the fact that plates do not contact all heat-producing components, cold plate can only chew up some 70% of a typical data center's server heat. Air-cooling has to do the rest. **What's more, cold plate typically maxes out at about 50 to 60 kW in rack density, making it unsuitable for many next-gen apps.** Again, the real constraint here is two-fold: a limitation on how many custom heat sinks you can plug into what becomes very crowded boards; and the air cooling required for the remaining 30-40% of the heat load.

Single-Phase Immersion Cooling

Although GRC's ElectroSafe™ coolant does not have the heat-carrying capacity of water, it is still 1,000X more efficient than air. With 300+ gallons of coolant in a rack, plus the managed flow and convection at work, this allows us to **easily and effectively cool over 100 kW per rack without compressor-based cooling (with warm water).** Theoretically, we can range up to **200 kW with a chilled-water system.**

So, while it's possible that cold plate could support higher heat flux, the metric has no practical application, as no existing or anticipated chips have come close to the limitations of immersion cooling. Imagine buying headphones that support a wider frequency range, but one that can't be heard by human ears. Yes, the spec is true. Yet is it useful and/or worth the higher cost and complexity?



Read our white paper on tackling the data center density dilemma >> grcooling.com/learning-center

KEY Takeaway:

Single-phase immersion can cool 100 kW/rack or more versus 50-60 kW/rack for cold plate.

Reliability & Location Flexibility

Cold Plate Cooling

The complexity of cold plate solutions can present significant reliability issues. First, the sheer number of intricate parts and fittings creates numerous failure points. Furthermore, leaks can be catastrophic. Remember, we're talking conductive water here: the enemy of electrical components.

Since server fans are still required, this introduces vibration and poses yet another point of failure. What's more, IT assets are exposed to environmental assaults (e.g., moisture and airborne particulates), which can hasten deterioration and impact MTBF (Mean Time Between Failures).

As for location flexibility, two things work against cold plate when compared to immersion cooling. **Cold plate is simply more complicated. It requires air conditioning infrastructure that is very energy-and cost-inefficient at a smaller scale.** It magnifies capital, power and site constraints. Plus, the fact that components are exposed to air can make these systems a no-go for deployments in harsh environments.

Single-Phase Immersion Cooling

The cooling systems we've developed and perfected at GRC involve the total immersion of IT assets. This affords full protection from the heat, moisture, oxidation, and dust that can create real reliability problems. In addition, since no air flow is required, rooms or modular structures where the racks are enclosed can be completely sealed off if needed.

That's how we give you the flexibility to locate your data center virtually anywhere on the planet, no matter how harsh the environment. **Further, the 50% lower energy consumption and minimal site and space requirements allow you to maximize IT capacity within limited power and space envelopes.**

Case in point: our ICEtank™ solutions sit inside an ISO container and are completely sealed off from the outside world. Not only can they be delivered in as little as 10 weeks, they can also go practically anywhere. Their simple, fully enclosed design makes them ideal for edge deployments.

Case Study: See how our ICEtank solution helped the U.S. Air Force >> grcooling.com/learning-center

We also offer a **"plug-and-play" solution called the ICeraQ™ Micro** – a 24U pre-configured rack that comes with an integrated CDU and pump. It lets you drop computing power in the most unlikely places with minimal site requirements. You just need power, water and a level floor.

KEY Takeaway:

Cold plate's complex architecture presents more points of failure and does not provide location flexibility.

Conclusion: Single-Phase Immersion Cooling Finishes First

Cold plate solutions have been around for some time. Although they can be a great spot-fix for isolated heating issues, their cost, complexity and risk have been obstacles to larger scale adoption. Single-phase immersion cooling not only overcomes all the limitations of cold plate. It delivers superior thermal performance and cooling of servers/components that wouldn't be feasible with cold plate. And it features a modular form factor that is designed for scale and enterprise-grade data centers.

These factors together have made cold plate increasingly ineffective, and immersion an obvious solution for bringing data centers into the future.

Additional Recommendations from Our Blog Library

 grcooling.com/blog



[Data Center Cold Wars – Part 1: Air-Based Cooling Versus Single-Phase Immersion Cooling](#)



[Data Center Cold Wars – Part 2: Two-Phase Versus Single-Phase Immersion Cooling](#)



[Data Center Cold Wars – Part 4: RDHx Versus Single-Phase Immersion Cooling](#)



[Cold Hard Facts of Immersion Cooling – Part 1](#) [Part 2](#)



[When Does Liquid-Immersion Cooling Make Sense? – Part 1](#) [Part 2](#) [Part 3](#)



Single-Minded About Growing You Data Center
with Immersion Cooling? Let's Talk!

Call **+1 512.692.8003** • Email **info@grcooling.com** • Visit **grcooling.com**

1 Based on a pPUE (partial PUE) of 1.03 vs. 1.15.
2 Based on a pPUE (partial PUE) of 1.03 vs. 1.15.