eBOOK

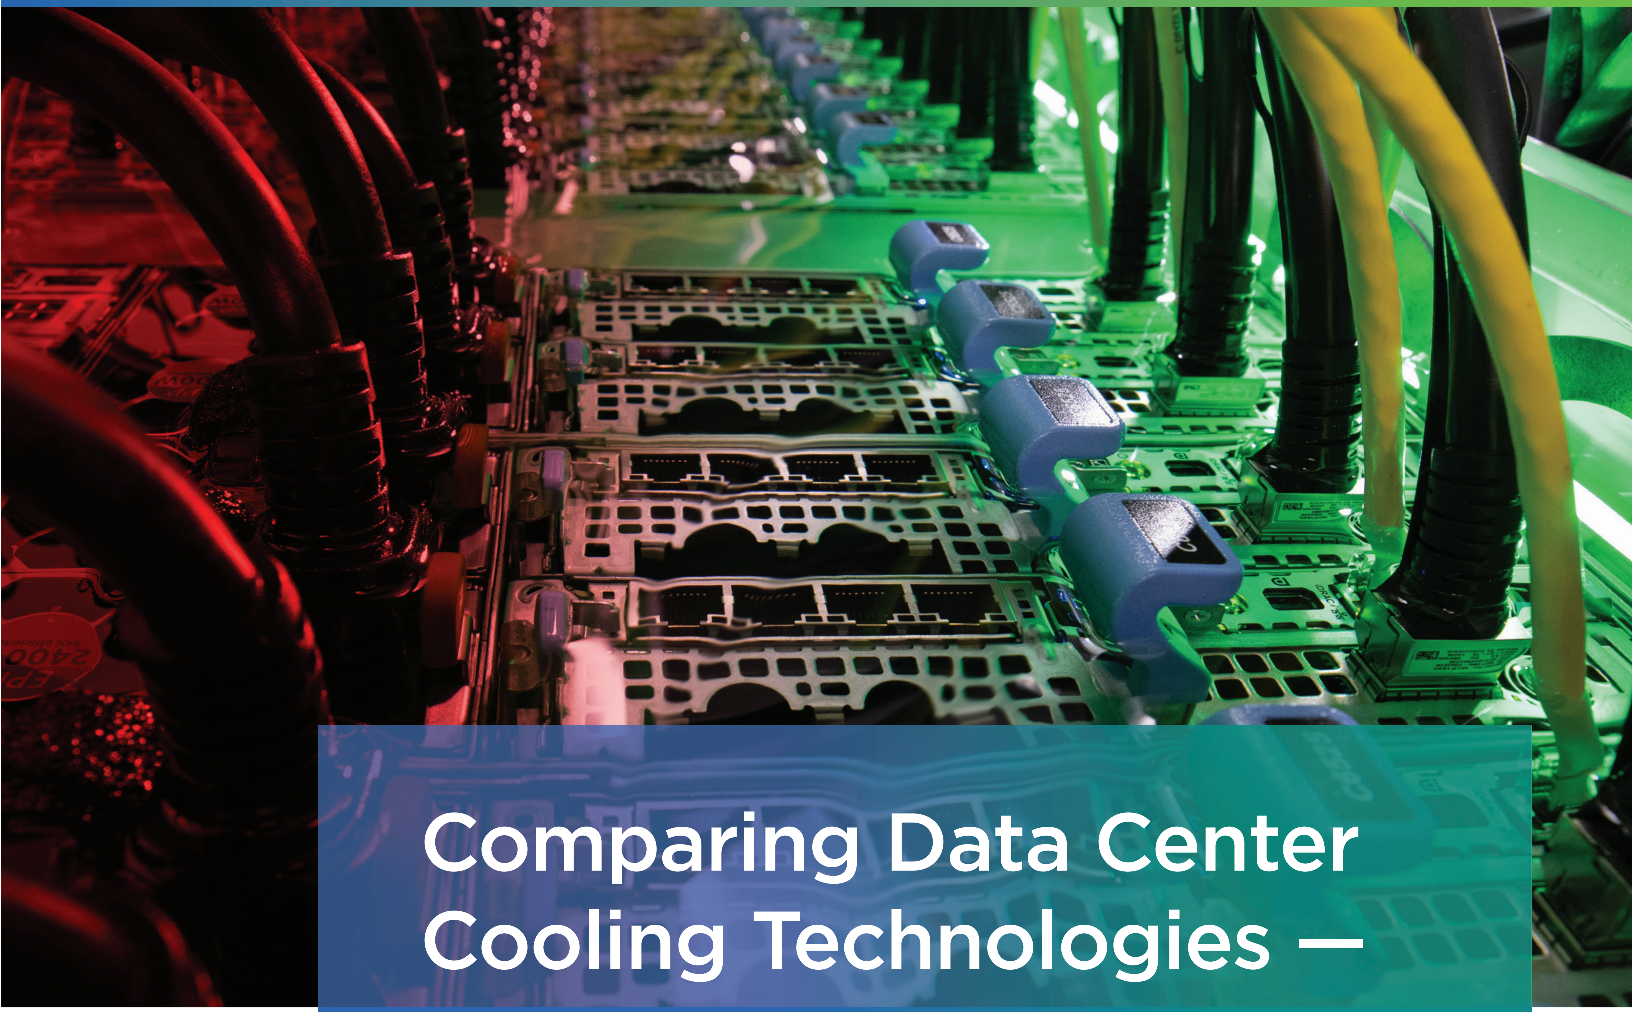

Comparing Data Center Cooling Technologies — Which Is Best for Your Operation?

GRC

THE IMMERSION COOLING AUTHORITY®

Table of Contents

Introduction 1

Cooling Technologies 2
- Legacy Air-Cooling (CRAC's/CRAHs) 3
- Rear Door Heat Exchanger (RDHx) Cooling 3
- Cold Plate Cooling 4
- Two-Phase Immersion Cooling 4
- Single-Phase Immersion Coolong 5

Complexity and Up-Front Costs 6
- Air Cooling vs. Single-Phase Cooling 7
- RDHx vs. Single-phase Cooling 8
- Cold Plate Cooling vs Single-Phase Cooling 9
- Two-Phase Cooling vs. Single-Phase Cooling 9

Efficiency & Operating Expenses 10
- Air Cooling vs. Single-Phase Cooling 11
- RDHx vs. Single-phase Cooling 12
- Cold Plate Cooling vs Single-Phase Cooling1 12
- Two-Phase Cooling vs. Single-Phase Cooling 12

Cooling Capacity and High-Density Performance 13
- Air Cooling vs. Single-Phase Cooling 14
- RDHx vs. Single-phase Cooling 15
- Cold Plate Cooling vs Single-Phase Cooling 15
- Two-Phase Cooling vs. Single-Phase Cooling 15

Reliability and Location Flexibility 16
- Air Cooling vs. Single-Phase Cooling 17
- RDHx vs. Single-phase Cooling 18
- Cold Plate Cooling vs Single-Phase Cooling 18
- Two-Phase Cooling vs. Single-Phase Cooling 19

Sustainability 20
- Power Consumption 21
- Water Use 21
- E-Waste 21
- Heat Reuse 21

Conclusion 22

Introduction

Cooling a data center has never been as challenging as it is today. 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. Even conventional computing operations are frequently pushing the 30 kW/rack barrier now. HPC apps like AI are becoming commonplace. And hot GPUs are becoming the weapon of choice for tackling high-performance computing (HPC) requirements. Because compute is such an integral part of business growth, IT executives are under renewed pressure to have a fail-proof game plan for scaling up and maintaining efficient and dependable operations.

In this eBook, we will compare our single-phase immersion cooling solution with several other popular cooling technologies across five different categories:

Let's start with a brief description of how each of these cooling technologies work.

Cooling Technologies

Cooling Technologies

Legacy Air Cooling (CRACs/CRAHs)

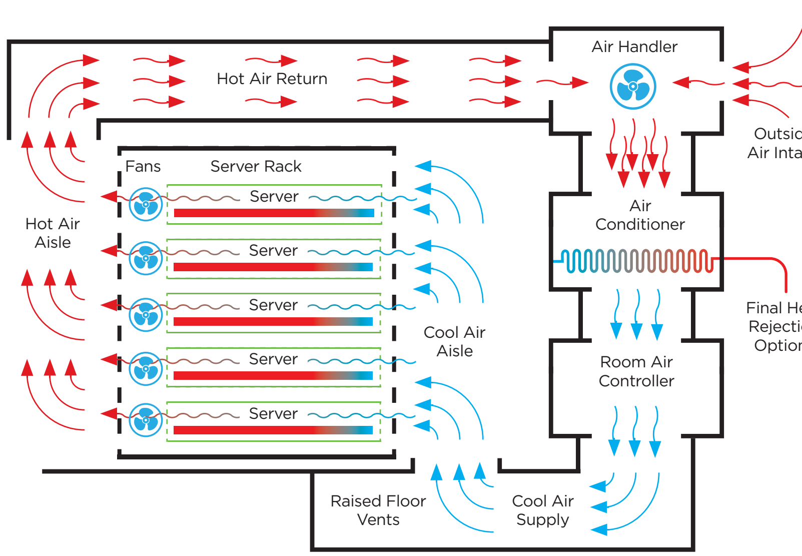

In its simplest form, an air-cooled data center brings outside air in through intakes on air handlers. This air is chilled by a computer room air conditioning (CRAC) unit before it is forced beneath a raised floor up into the "cold aisle" of the server racks. This cold air moves through and cools the servers, then exits via the "hot aisle," where it is contained and vented through a plenum that returns it to the air handlers. Cold-water chillers and cooling towers can also be used.

Rear Door Heat Exchanger (RDHx) Cooling

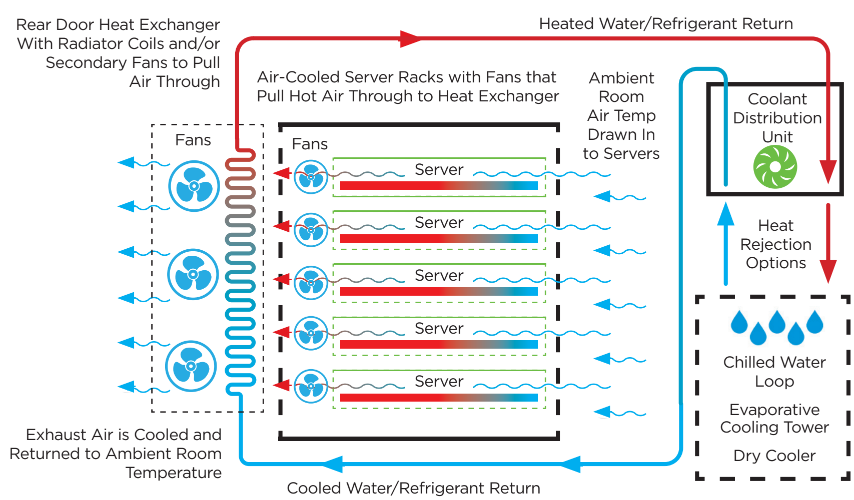

With RDHx, radiator-like doors are attached to the fronts or backs of server racks. Heat produced by IT equipment is removed by air blown into the exchangers via fans. That heated air travels a short distance before transferring to water or refrigerant running inside the radiators. Heated liquid is then carried away from the server rack to a cooling distribution unit (CDU), which in turn is connected to an outside chilled water system. Some RDHx systems have "active doors" with integrated fans, while others rely solely on server fans to eject the air.

Essentially, RDHx is just a better form of air-cooling. Although often mistakenly considered so, it is not a liquid-cooling solution.

Cooling Technologies

Cold Plate Cooling

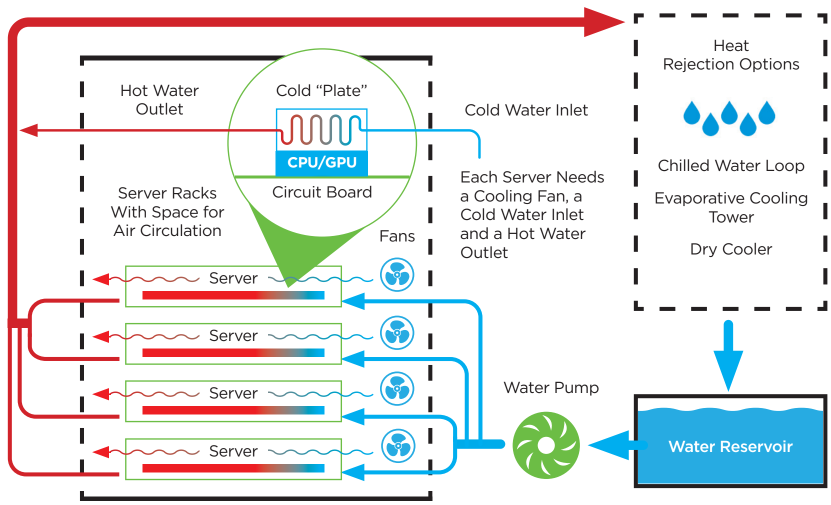

Basically, "touch," "direct-to-chip", or "liquid-to-chip" cold plate cooling replaces older, inefficient air-cooled metal-finned heat sinks with liquid-cooled heat sinks. As the name suggests, it attaches a metal plate atop a CPU or GPU, which transfers heat through a heat-spreading material (such as 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.

Two-Phase Immersion Cooling

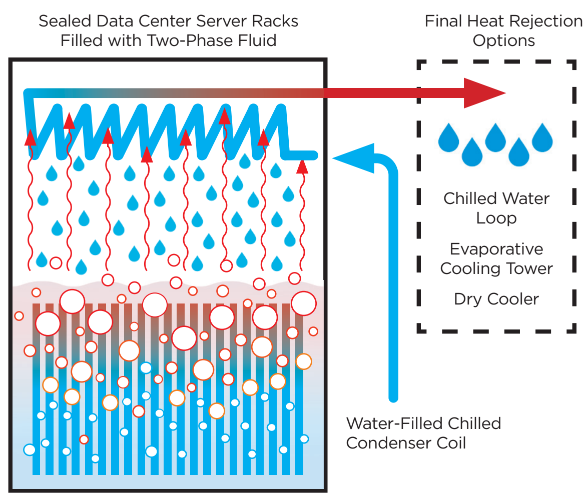

Novec Fluid Boils from Server Heat, Transforms from Liquid to Vapor and Passively Condenses Back to Liquid When It Contacts Water-Cooled Condenser Coils

In a two-phase immersion-cooled system, servers are sealed inside a bath of specially engineered fluorocarbon-based liquid. Because the fluid has a low boiling point (often below 50°C vs. 100°C for water), heat from the servers easily boils the surrounding fluid. The boiling of the liquid causes a phase change (from liquid to gas), which gives two-phase immersion cooling its name. The vapor is then condensed back to the liquid form via water-cooled condenser coils, which are integrated into the top of the sealed racks. The condensed liquid drips back into the bath of fluid to be recycled through the system.

Single-Phase Immersion Cooling

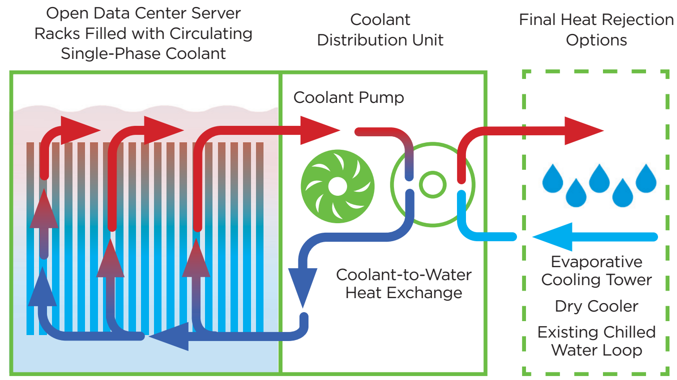

Heated Coolant Exits Top of Rack, Cycles Through the CDU, and Returns to the Rack at a User-Specified Temperature.

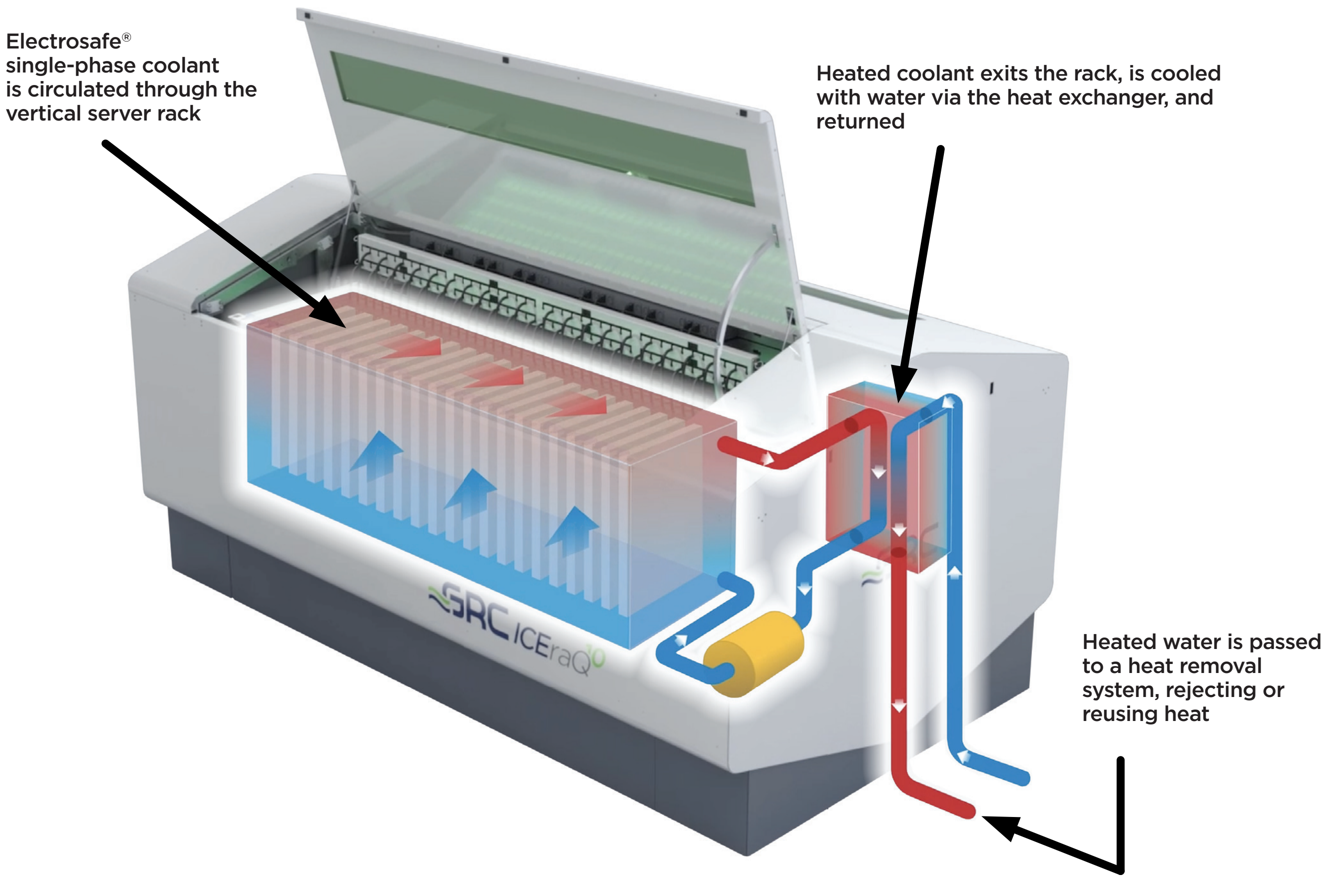

With single-phase immersion cooling, servers are installed vertically in a coolant bath of dielectric fluid. The coolant transfers heat through direct contact with server components. Heated coolant then exits the top of the rack and is circulated through a CDU connected to a warm-water loop. This loop 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.

Now, let's take a closer look at how these competing technologies compare to single-phase immersion cooling in several key areas.

Complexity and Up-Front Costs

Complexity and Up-Front Costs

Air Cooling vs. Single-Phase Immersion Cooling

A data center utilizing an air-cooling system requires some combination of raised floors, aisle containment strategies, chillers, air handlers, humidity controls, filtration systems, and plenums. Furthermore, to support the above, air-cooled data centers must also operate a comparatively large ancillary infrastructure\—notably backup generators, UPSs, and batteries. All this necessary complication equates to a relatively large capital expenditure (CapEx).

Typical 600 kW Air-Cooled Data Center Infrastructure

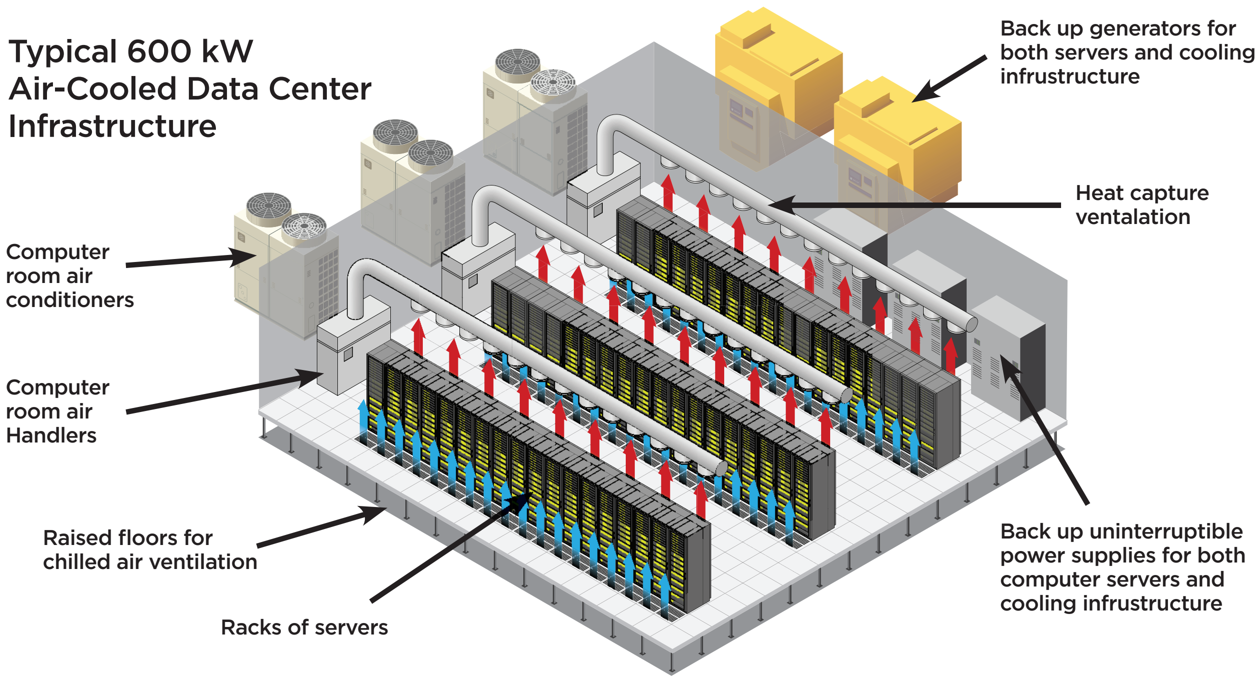

GRC's single-phase immersion cooling solution has just three moving parts; a coolant pump, water pump, and a cooling tower/dry cooling fan. It requires no raised floors nor wasted space through aisle containment. As a result, it can cut data center CapEx up to 50%. What's more, no CFD analysis of air flow is required with immersion because racks can be spaced close together, and even placed on bare concrete floors. Electrical support systems can be downsized as well.

Equivalent 600 kW Immersion-Cooled Data Center Infrastructure

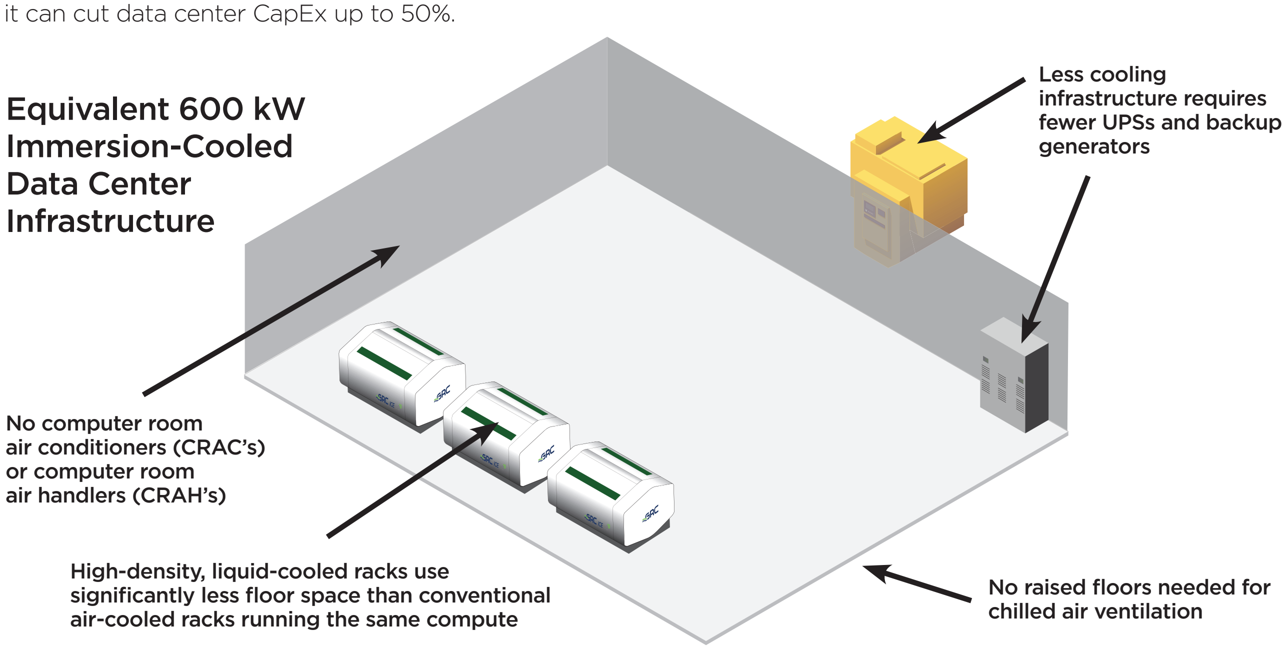

Complexity and Up-Front Costs

RDHx Cooling vs. Single-Phase Immersion Cooling

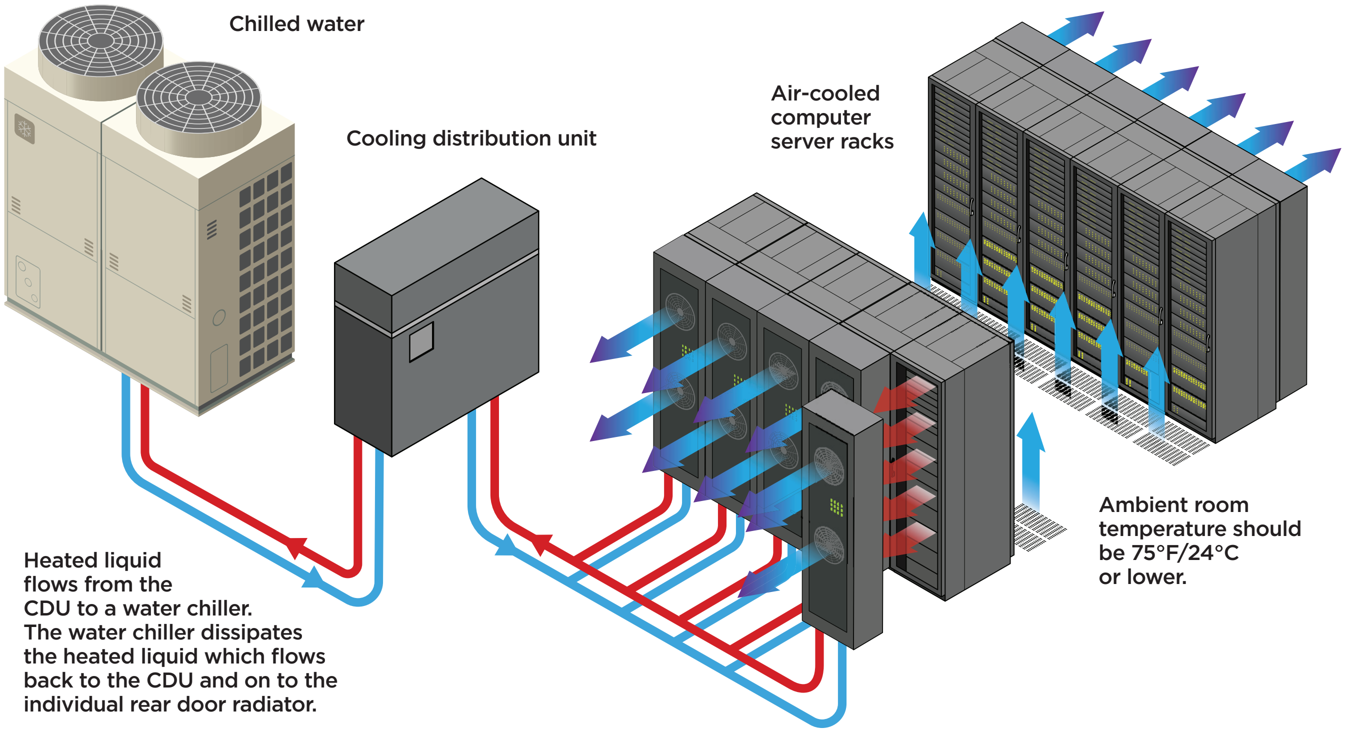

Rear Door Heat Exchangers replace standard rear doors on IT rack enclosures and by using the rack-mount devices draw cool supply air through the chassis, the heated exhaust air passes through a liquid-filled coil, transferring heat to the liquid with cool neutralized air flowing back into the data center.

RDHx is simpler than traditional air cooling but since it still uses air as a medium of cooling, it requires compressor or refrigerant-based cooling at the CDU level, or a chilled water loop. It also requires specialized building design and engineering such as CFD analysis to ensure appropriate layouts and cooling.

Further, environmental and humidity control systems are also needed. While RDHx does not require raised floors, aisle containment, false ceilings, and plenums may still be needed. All these systems increase complexity and add costs. Like any air-cooled solution, RDHx is more complex, and hence more costly, than single-phase immersion cooling.

Complexity and Up-Front Costs

Cold Plate Cooling vs Single-Phase Immersion 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. Additionally, it often requires a chiller plus specialized engineering depending on the design of the servers. This all adds up to significantly greater upfront costs.

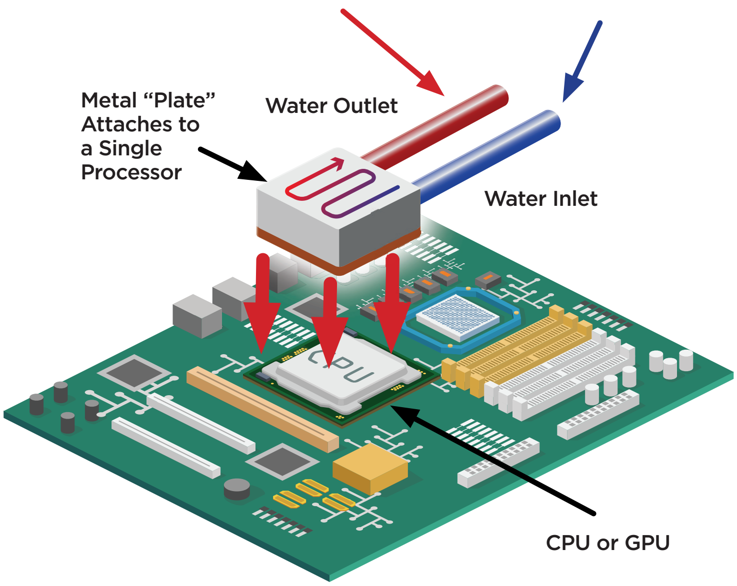

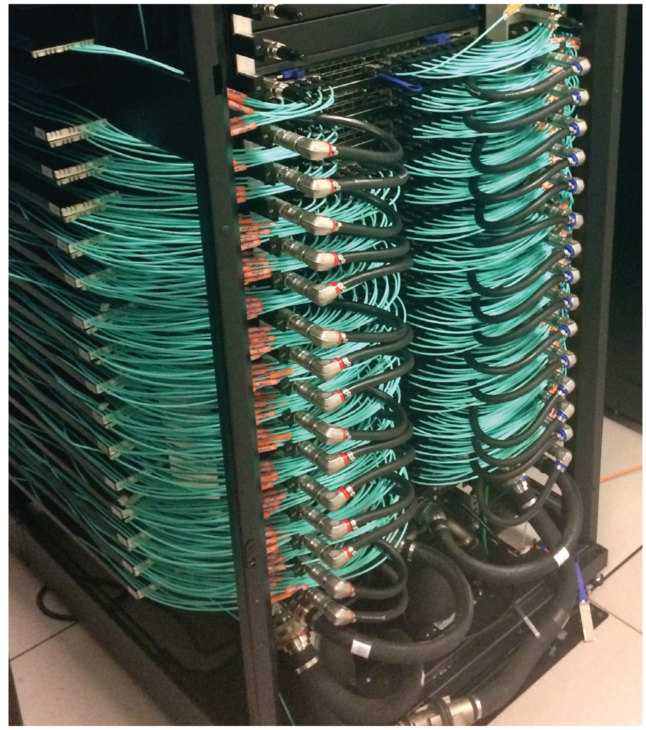

Cold Plate Plumbing

To further complicate matters, each heat source (CPU/GPU) typically requires its own cold plate 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.

The simple design of a single-phase liquid immersion cooling system can reduce data center CapEx up to 30% over cold plate cooling.

Two-Phase Cooling vs. Single-Phase Immersion Cooling

Two-phase immersion cooling systems are a bit more complex in their engineering than single-phase cooling and present higher up-front costs. While pumps are not needed to circulate fluid around the servers, sealed racks with integrated condenser coils are required to make it all work.

The biggest cost difference between single and two-phase relates to the coolant. Fluorocarbon-based, two-phase coolants typically cost an order of magnitude more than hydrocarbon-based ones such as our ElectroSafe coolants, which are validated for use with GRC systems through our ElectroSafe Fluid Partner Program. When you're using over 300 gallons of the coolant in a single rack, those additional costs can be significant.

Although it requires a pump to keep the coolant circulating through the system, single-phase immersion cooling beats two-phase cooling in the complexity category.

Efficiency & Operating Expenses

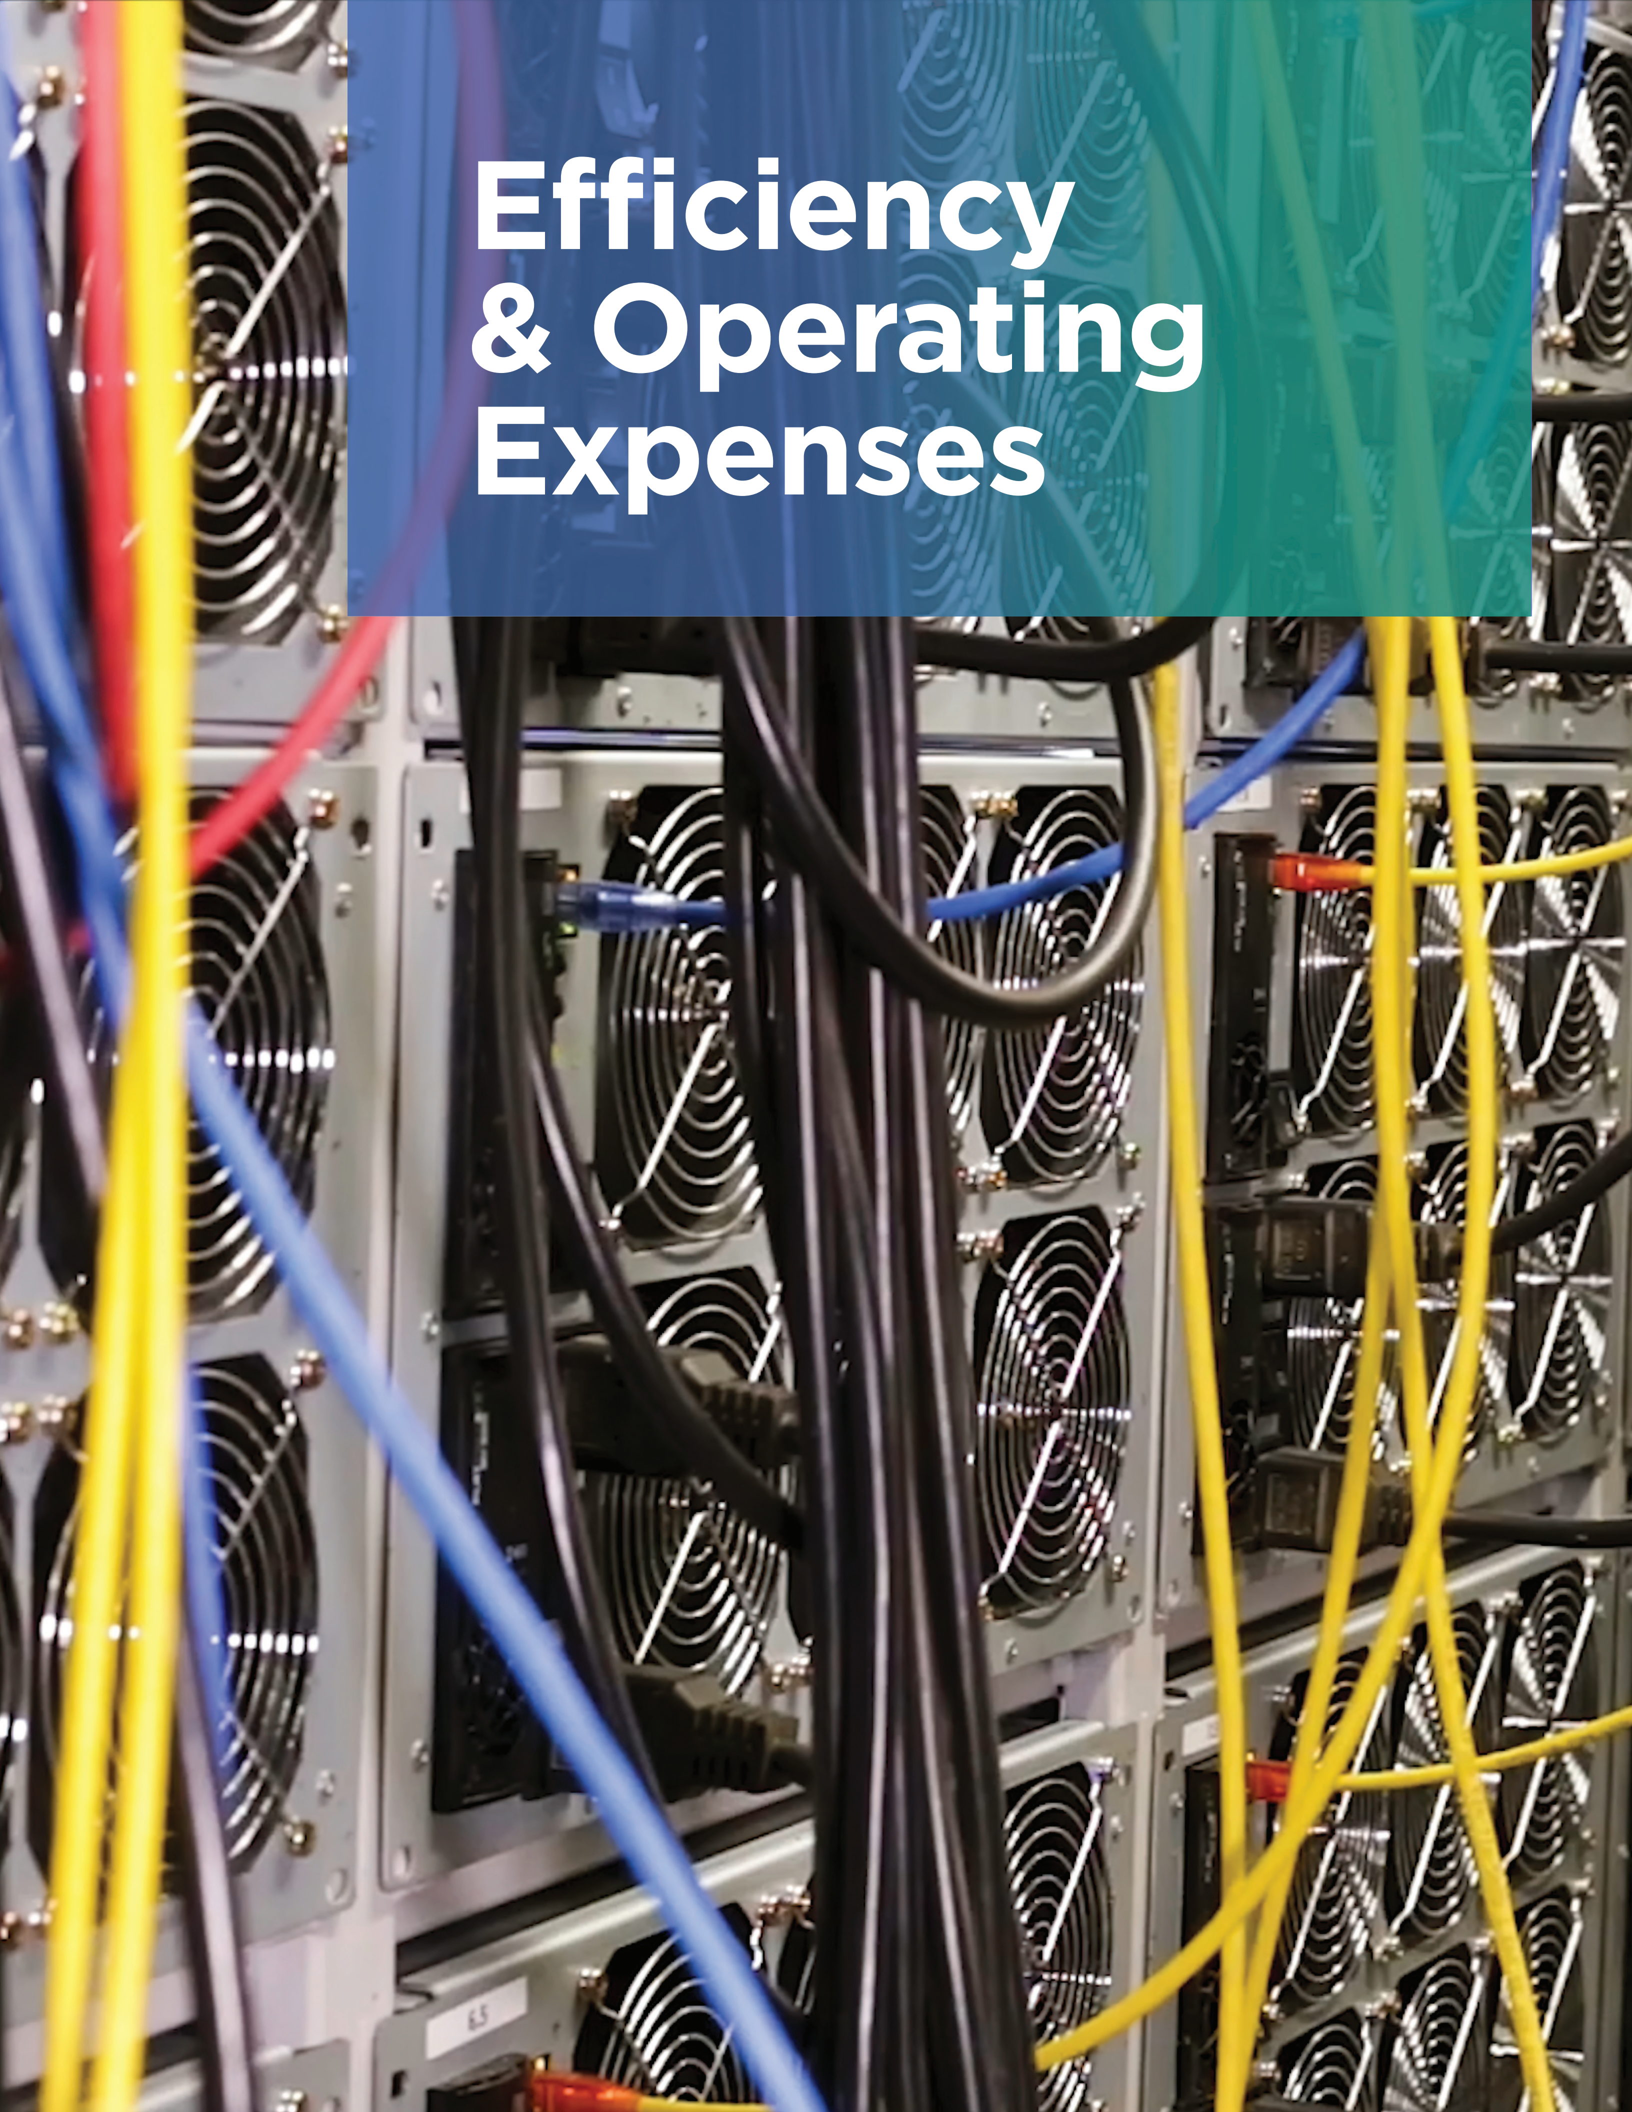

Efficiency and Operating Expenses

Air Cooling vs. Single-Phase Immersion Cooling

Air is 1200X less effective as a conductor of heat than liquid. This not only makes air-cooled data centers intrinsically less efficient, but it creates ripple effects that have a serious impact on operating expenses.

First off, fans account for approximately 20% of server power consumption. To bolster the effectiveness of air cooling, energy-sapping refrigeration components like chillers and air handlers are needed as well. Those in turn impact the size of power infrastructure.

Given all the above, basic air cooling requires the highest operating expenses of all major data center technologies while delivering a PUE of approximately 1.35 to 1.69.

Single-phase immersion cooling delivers up to a 90% reduction in cooling energy and 50% cut in total data center energy usage over air cooling. As a result, operators can realize an immersion cooling PUE of <1.03.

In a typical air-cooled data center, a large percentage of power is used specifically to cool heat generated by the servers.

Air handlers and electrical systems for the HVAC in a data center.

Where Does Power Go in an Air-Cooled Data Center?

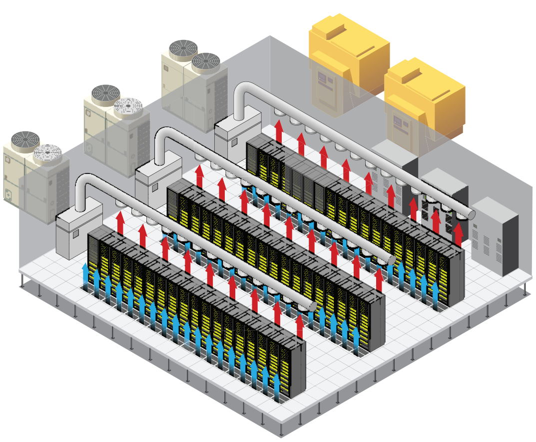

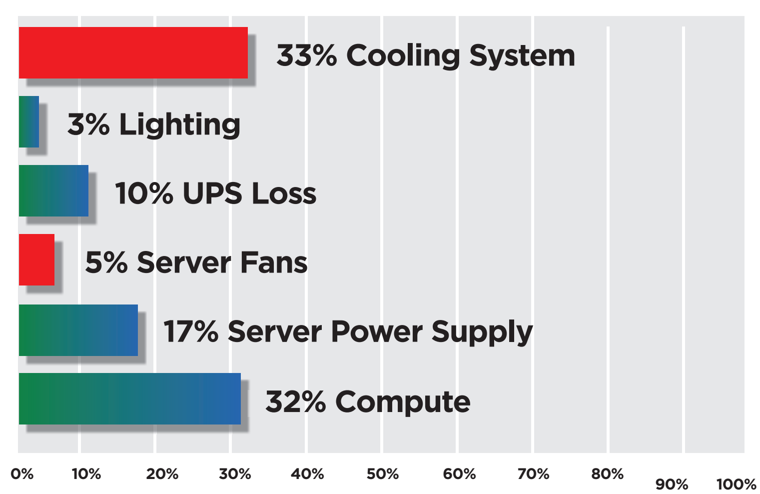

Source: Improving Energy and Power Efficiency In The Data Center (semiengineering.com)

Efficiency and Operating Expenses

RDHx Cooling vs. Single-Phase Immersion

Because RDHx uses air, it requires a lot of air flow and a huge temperature delta between the incoming air and the heat source to maintain optimal core temperatures. The higher delta means that air must be cooled to temperatures that mandate the use of compressors or chiller plants to cool the air directly. Either way, when you consider the need for compressors, plus the number of fans necessary to create the required air flow, RDHx consumes significant amounts of energy, much more than single-phase cooling systems. Additionally, unlike immersion cooling, RDHx solutions still require server fans and do not reduce the server power consumption. Given the combined electrical load of fans, chillers, CDUs, and more, RDHx delivers a PUE of between 1.2 and 1.3.

It's also important to note that, given tightening worldwide environmental regulations, many refrigerants are likely to get harder and costlier to acquire in the long run. This should be given serious consideration when building a data center that requires compressor-based cooling.

With single-phase immersion cooling systems, 100% of the heat is picked up by the coolant. In addition, the coolant's comparatively higher density and superior thermal conductivity means that it can maintain optimal core temperatures with a much lower delta between the coolant and the heat source. Add the disabling or removal of server fans and a 10-30% reduction of server power is possible.

Cold Plate Cooling vs. Single-Phase Immersion Cooling

Because it needs expensive and inefficient air conditioning to be 100% effective, cold plate cooling is intrinsically less efficient and more expensive to maintain 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 of about 1.15.

Single-phase cooling systems are 80% more energy-efficient than cold plate methods, delivering a PUE of 1.03. That delta between cold plate and single-phase immersion PUE is amplified by the 10-30% server power reduction immersion cooling enables through fan removal. The result equates to a total data center energy cut of up to 35%.

Two-Phase Immersion Cooling vs. Single-Phase Immersion Cooling

Both single and two-phase immersion cooling systems blow the doors off air-cooled data centers when it comes to cooling efficiency with two-phase delivering a slightly lower PUE of 1.01-1.02.

However, unlike ElectroSafe fluids, fluorocarbon-based coolants must be regularly replenished, again at hundreds of dollars per gallon. That's enough to raise overall operating expenses and offset any small power efficiency gains, to say nothing of total cost of ownership (TCO). ElectroSafe is far less expensive upfront, does not evaporate, and will easily last 15 or more years.

While two-phase systems are energy-efficient, performing server maintenance on immersed ITE can be complicated. To access the servers, operators must unseal the enclosure while trying to minimize coolant loss and vapor inhalation. With single-phase, you simply open the lid and lift out the equipment.

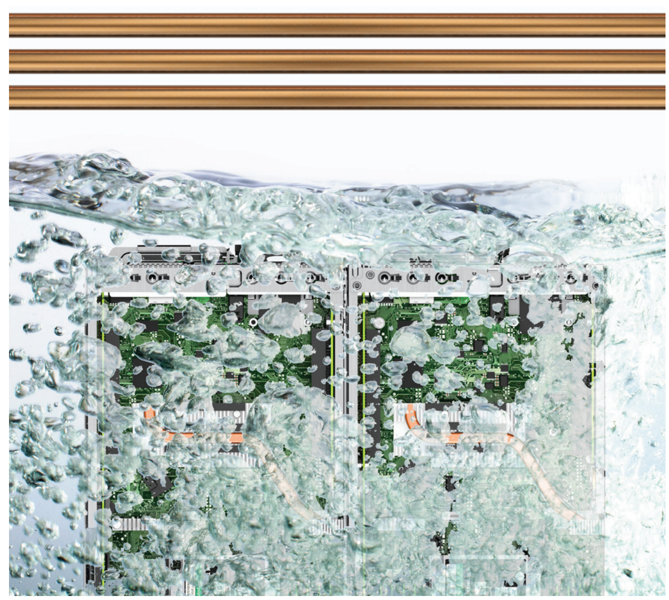

Two-phase immsersion cooling "boils" off heat. Vaporized fluid condenses on a chilled pipe and drips back down into the pool of fluid — all within a sealed container.

Cooling Capacity and High-Density Performance

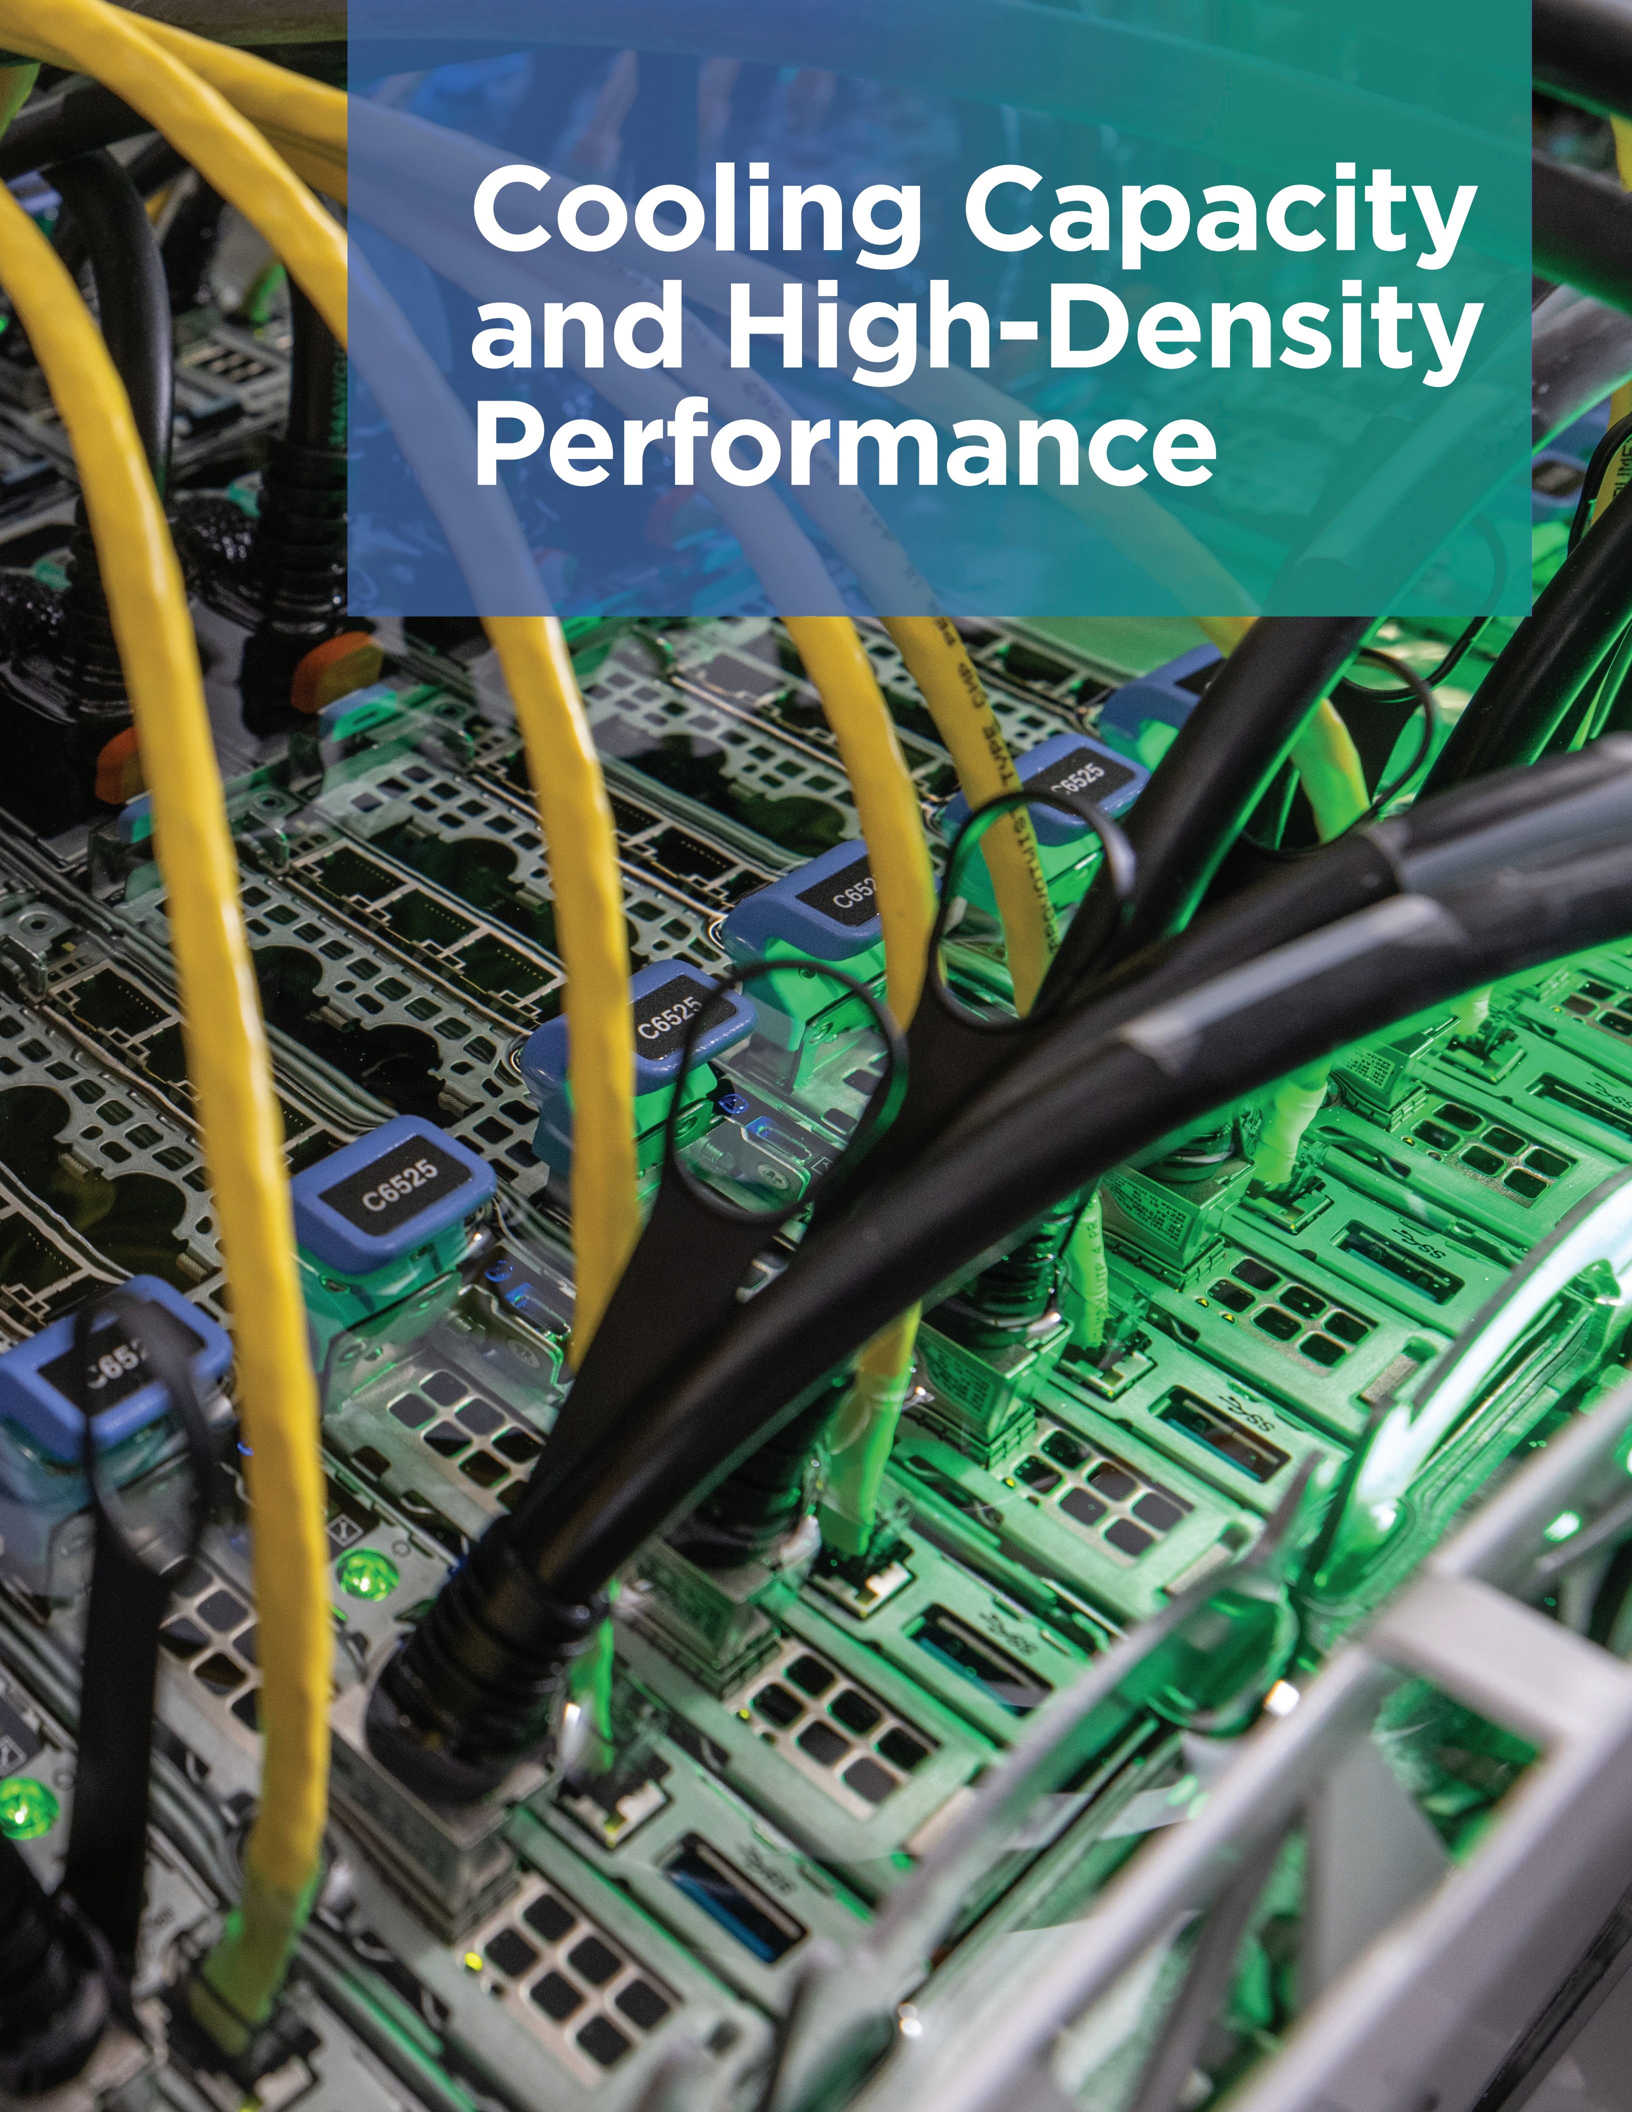

Cooling Capacity and High-Density Performance

Air Cooling vs. Single-Phase Immersion Cooling

It's true that some air-cooled data centers are capable of cooling upwards of 30-35 kW/rack. However, air-cooled data centers become very inefficient above 15 kW/rack. Industry trends are making the situation worse. Electricity-hogging GPUs are moving in to tackle HPC applications like IoT and AI.

To keep up with demands, data center operators may be inclined to create mixed-density racks. Where air cooling is concerned, this inevitably leads to hot spots, which can lead to hardware failure. This evolution in hardware will create a real moment of reckoning for operators of air-cooled data centers at their next hardware refresh.

GRC ICEraQ® and ICEtank® systems are ideally engineered to break through the heat barrier and can easily cool up to 100 kW/rack; theoretically up to 200 kW when used with a chilled-water system. Also, hot spots are not an issue with these systems.

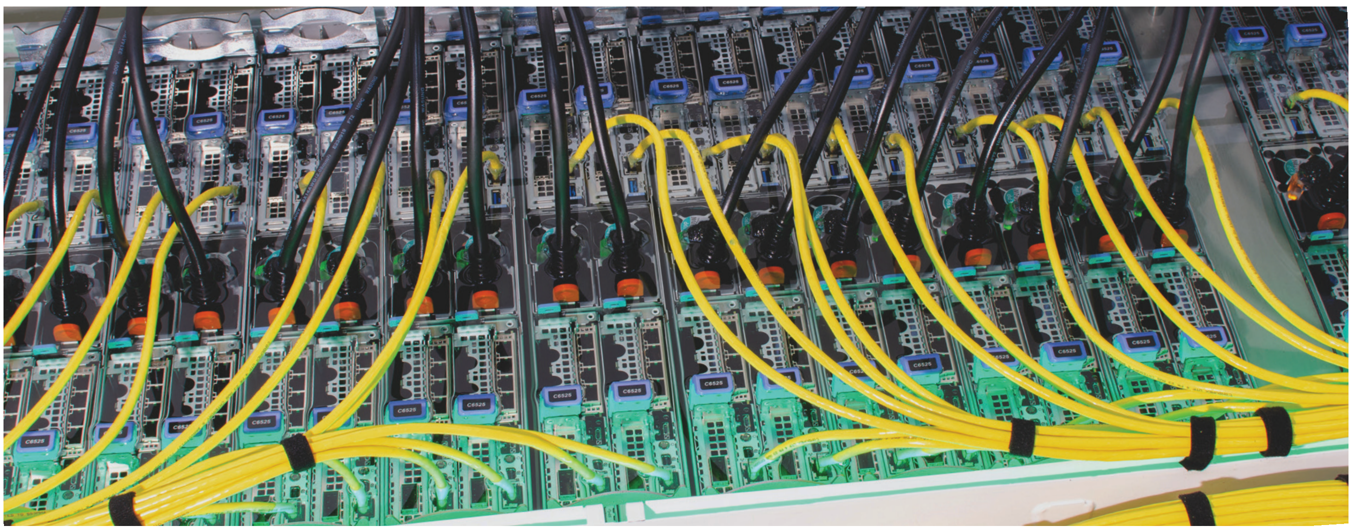

GRC Single-phase immersion cooling systems can densely pack high-performance computer servers with plenty of cooling to spare

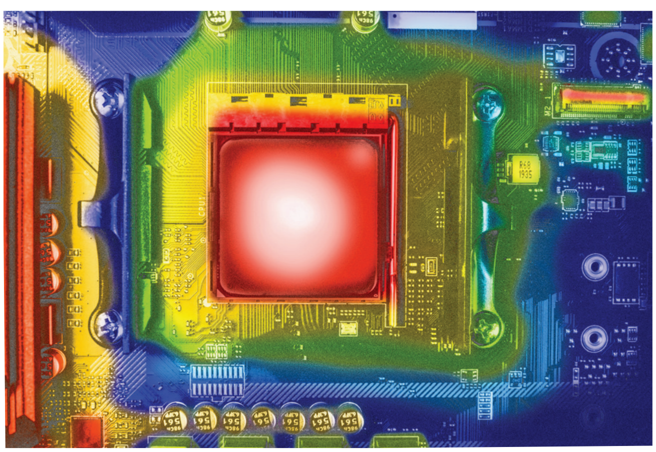

Thermal imaging shows just how intense heat can get for CPUs and GPUs as well as other electrical components.

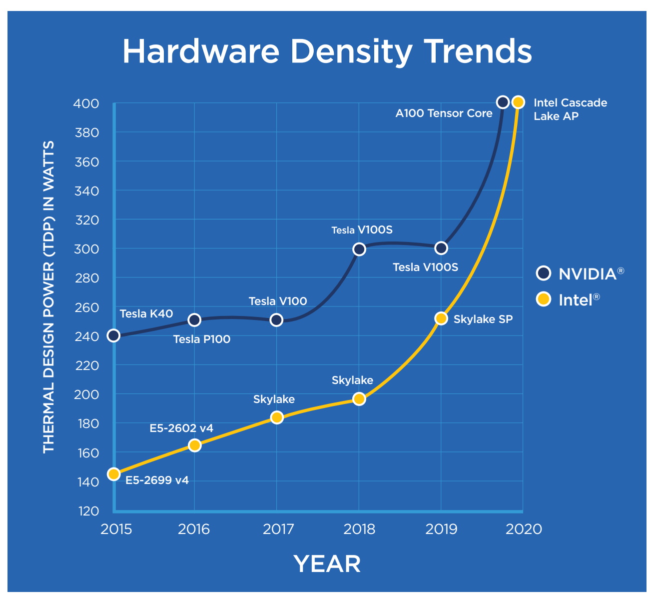

CPUs and GPUs are getting hotter each year, and showing no signs of slowing down.

Cooling Capacity and High-Density Performance

RDHx Cooling vs. Single-Phase Immersion

RDHx systems can utilize either water or refrigerant. The former comes with major safety and reliability caveats to be discussed later in the Reliability and Location Flexibility section. Although RDHx can be cost-competitive for 25 to 30 kW rack densities, there are limits to those benefits. In fact, its advantages drop sharply as densities rise beyond 15 kW per rack. Given the ever-increasing density of chips and hardware, operators may soon start to hit the 30 kW per rack wall with hardware refreshes.

GRC ICEraQ and ICEtank systems can easily cool up to 100 kW per rack (200 kw with chilled water). And while the financial case for GRC's immersion cooling makes sense at rack densities as low as 10 kW (less in some cases), the higher capacity future-proofs the infrastructure against rising hardware density and the ever-evolving technology landscape.

Most air-cooled racks require blanking plates for extra air space between servers

Cold Plate Cooling vs. Single-Phase Immersion Cooling

Cold plate cooling systems typically employ water or a water/glycol mix. Since the 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. Thus, even though they are great for cooling hot spots, they don't perform 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.

The coolants used in cold plate cooling offer higher thermal conductivity than dielectric coolants, but are also good conductors of electricity. This 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.

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 must 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.

The dielectric coolant used in single-phase cooling can directly contact all components within a whole rack of servers to capture 100% of the heat and cool each chip effectively. Although GRC's ElectroSafe coolants do not have the heat-carrying capacity of water, they are still 1,200X more efficient than air.

Two-Phase Cooling vs. Single-Phase Immersion Cooling

Two-phase cooling can typically support extreme rack densities, with some providers claiming to support up to 250 kW in a single rack! While the claim sounds extremely impressive, it's hard to find commercially available hardware that will get you to even 100 kW.

GRC's ICEraQ is capable of handling 200 kW per rack when attached to a chilled water system. Given the anticipated long-term "heat trend" for incorporating next-gen applications like AI, IoT, facial recognition, and others, that's more than enough cooling performance for the foreseeable technology evolution.

Reliability and Location Flexibility

Reliability and Location Flexibility

Air Cooling vs. Single-Phase Immersion Cooling

Any cooling technology that draws air from the outside is destined to create hardware reliability issues because it exposes IT assets to potentially harmful airborne contaminants as well as corrosion and oxidation resulting from the air itself. That risk increases depending on the air quality and natural humidity levels. Clearly, locales with a high degree of humidity, air pollution, or windblown particulates can wreak havoc on ill-equipped data centers. These concerns increase as the need for remote edge deployments increases.

The inherent complexity and larger infrastructure requirements of air cooling also present significant hurdles as to where computing power can be placed. Plus, air cooling creates hot spots which can lead to hardware failure.

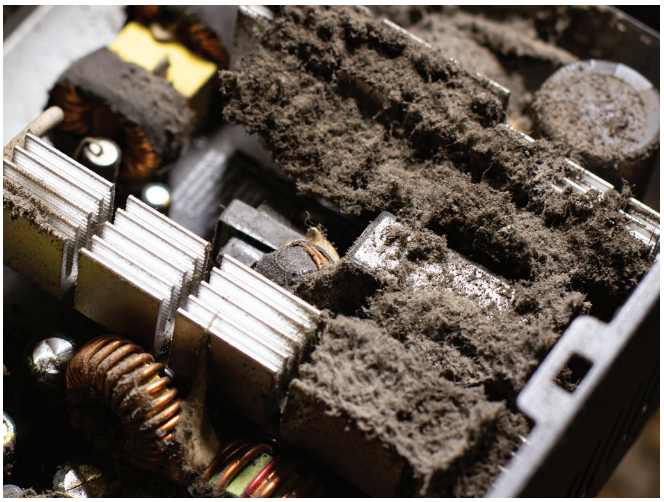

All air-cooled data centers have airborne particulates, dust, moisture, and pollutants which will eventually harm components and disrupt cooling.

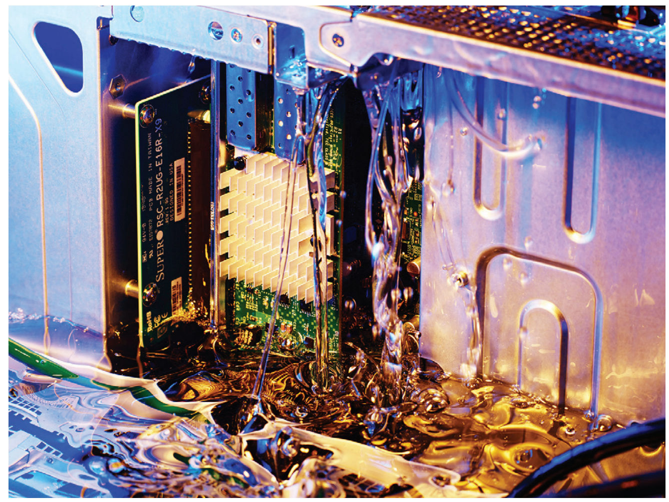

Liquid immersion is environmentally resilient, insulating IT equipment from all forms of particulates and temperature variations.

With single-phase immersion cooling, there are fewer things to go wrong: no chillers, air handlers, humidity controls, etc.; and no server fans to create vibrations that can (and do) increase mean-time between failures (MTBF). Additionally, critical IT assets are completely sealed off from the outside air, negating any environmental issues. Finally, there are no hot spots in the data center.

Close to, or far away from civilization, GRC's ICEtank provides a powerful standalone data center that can be placed virtually anywhere power, water and internet is accessible.

Reliability and Location Flexibility

RDHx Cooling vs. Single-Phase Immersion Cooling

Feeding the radiator-like doors of RDHx systems are fluid lines containing either water or refrigerant. Depending on which, and whether they're routed above or below critical IT equipment, leaks can be a major problem. Additionally, RDHx refrigerants have a global warming potential (GWP), putting them under the constant scrutiny of the EPA and EEA. In fact, some formulations are in the process of being banned, causing more than a few operators to be concerned about supply disruptions.

RDHx systems also expose IT assets to airborne contaminants, moisture, and increased levels of vibrations. Thus, they require additional systems and equipment to protect them from harsher environments. Plus, RDHx needs significant amounts of power and supporting infrastructure, which can limit location choices, especially for edge applications.

GRC ICEraQ and ICEtank systems, on the other hand, are the epitome of simplicity, and can be placed virtually anywhere. In fact, our fully containerized ICEtanks can literally be deployed in a parking lot or loading dock.

Rear door heat exchangers pull hot air through a radiator, cooling the air in the process, then expell the cooled air back into the data center.

Cold Plate Cooling vs. Single-Phase Immersion 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, water leaks can be catastrophic. 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.

As for location flexibility, cold plate 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 cooling 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.

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.

Reliability and Location Flexibility

Two-Phase Cooling vs. Single-phase Immersion Cooling

Both single and two-phase immersion cooling systems on the market today give information and communications technology (ICT) leaders tremendous flexibility in locating their data centers. This is because neither requires the extensive infrastructure traditional air-cooled installations do.

However, studies have shown that fluorocarbon-based fluids used in two-phase immersion cooling present several safety and reliability issues GRC's ElectroSafe coolants don't. For one, fluorocarbon-based fluids have a high Global Warming Potential (GWP). Some experts have also suggested that the boiling action on which two-phase cooling depends may harm IT assets over time through a process known as cavitation. Seals for the coolant bath can also fail. Plus, the expensive coolant evaporates quickly, leaving the rack empty and without cooling. Finally, it should be noted that the condenser that's placed directly over the servers in two-phase systems is water-cooled. While the two are unlikely to come in contact, there are widespread concerns about having a conductive liquid directly above the IT rack.

GRC's single-phase immersion cooling systems pose virtually no maintenance issues and can consistently cool 200 kW/rack using one of our ElecroSafe® dielectric coolants, which are totally inert, non-conductive, non-flammable, non-corrosive, and do not need replacement over the life of a typical data center.

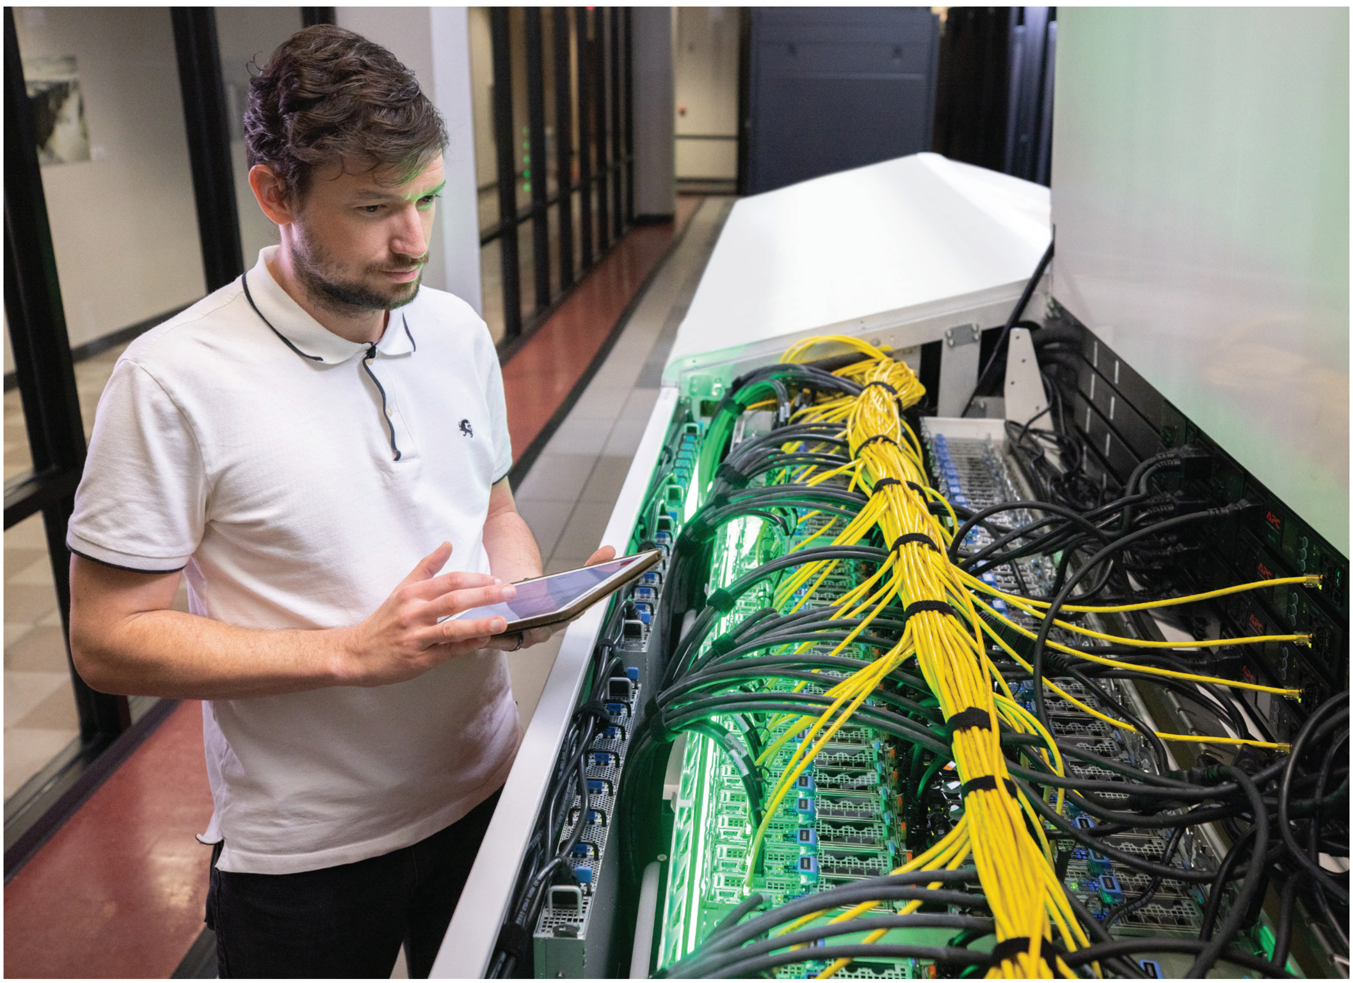

Maintenance on our cooling systems is simple. Servers are easily accessable, and ElectroSafe single-phase immersion coolants do not evaporate like two-phase coolants.

Sustainability

Sustainability

Traditional data centers leave a significant carbon footprint. Increasing pressure from the public and industry to reduce their environmental impact is forcing operators to seek more sustainable options. **There are four key factors impacting data center sustainability that must be addressed:**

1: Power Consumption

As stated earlier, fans account for up to 20% of server power consumption in air-cooled data centers. The additional infrastructure required for air-based cooling, including power-hungry refrigeration components like chillers and air handlers, increases power consumption drastically. Since both cold plate cooling and RDHx systems still rely on air cooling to some extent, power consumption is not reduced by using either of these options. With both single- and two-phase cooling, power consumption is reduced another 35-40% with the removal of all the unnecessary infrastructure.

2: Water Use

Water use is also a significant sustainability factor for traditional, air-cooled data centers and is primarily driven by two factors, power generation (indirect) and chilling water for cooling purposes (direct). Again, because of their reliance on air cooling, both cold plate cooling and RDHx systems have the same issues with water use. Immersion-cooled data centers can significantly reduce the water consumption due to the minimal power requirements of the systems themselves, as well as their ability to cool high-density racks with chiller-free water. However, two-phased cooling systems do employ water-cooled condensers which makes them a less attractive option than single-phase immersion cooling in this regard.

3: E-Waste

Given the amount of equipment and infrastructure required for air cooling, RDHx, and particularly cold plate cooling systems, e-waste is an inevitable and unavoidable result. At the end of their life cycle, there is an enormous amount of material that will end up in landfills. While both single- and two-phase immersion cooling systems result in far less physical e-waste, the extremely expensive fluids employed in two-phase cooling have a high GWP and are thus, far more harmful to the environment than the recyclable fluids used in single-phase immersion cooling. It is also worth noting that with OEM manufacturers set to bring immersion-ready servers to market, this will mean even less e-waste for data centers employing single-phase immersion cooling systems.

4: Heat Reuse

Air cooling only captures 30% of the heat generated by the servers providing little opportunity for heat reuse. The same can be said for cold plate cooling and RDHx systems. Because in an immersion cooling system the liquid circulates around all the components, the fluid captures 100% of the server heat, making it more efficient and cost effective.

Conclusion

Single-phase immersion cooling is the obvious solution for bringing data centers into the future. With its simple design and unequaled location flexibility, it has been proven cost-effective and highly reliable by GRC over the course of more than a decade. Our ICEraQ and ICEtank solutions are helping some of the world's largest cloud, enterprise, government, education, and telecom organizations future-proof and grow their data centers.

While each of the other cooling solutions examined may offer some attractive features, they simply cannot match single-phase immersion cooling in the five critical performance categories discussed here.

Take the Next Step in Discovering How Single-Phase Immersion Cooling Can Meet Your Data Center's Demands — Contact a GRC Expert!

Call: **+1.512.692.8003** Email: **ContactUs@grcooling.com** Visit: **grcooling.com**

GRC
THE IMMERSION COOLING AUTHORITY®

11525 Stonehollow Dr. Suite A-135 • Austin, TX 78758 • United States • +1.512.692.8003 • ContactUs@grcooling.com