WHITE PAPER: PUMPED 2P REFRIGERANT-BASED DIRECT LIQUID COOLING (DLC) TECHNOLOGY FOR NEXT GENERATION AI CLUSTERS WITH HIGH TDP ACCELERATORS

Devdatta Kulkarni, Intel

Russ Tipton, Vertiv

Elvis Leka, Parker Hannifin Sporlan Division

Joe Bankston, Parker Hannifin Hose Division

Qingyang Wang, Accelsius

Jeff King, Intel

Date: February 14, 2025

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Version Table

Date	***Version #***	***Author***	***Description***
02/14/2025	*1.0*	*Dev, Russ, Elvis, Joe, Qingyang, Jeff*	*Original Paper*

Date: February 14, 2025

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Executive Summary

The rise in artificial intelligence (AI) has created an acute demand for thermally dense accelerator chips, servers, and racks. Power and cooling infrastructure in legacy data centers is unable to keep pace, causing a shift toward more highly capable and energy efficient liquid cooling technologies. Open Compute Project (OCP) has been instrumental in standardizing common components to build robust hardware ecosystems and encourage the adoption of liquid cooling at scale. In recent years, pumped two-phase (P2P) direct to chip (DTC) cooling with low-to-medium pressure refrigerants has emerged as a promising technology to address the surging AI demand. Careful implementation of P2P DTC cooling within dense AI hardware applications provides a safe, energy efficient, and highly capable method for improving critical end-user key performance indicators (KPI) such as power usage effectiveness (PUE), water usage effectiveness (WUE), and total cost of ownership (TCO). In addition to its inherent key advantages, P2P cooling utilizes low global warming potential (GWP), chemically inert, highly dielectric working fluids, thus providing low environmental impact and reduced risk of downtime due to corrosion and leaks as compared to single phase (1P) DTC approaches. This paper is intended to provide an overview of P2P DTC cooling technology including key solution ingredients, important performance and reliability advantages, potential risks, and overall supplier ecosystem health.

Date: February 14, 2025

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Table of Contents

1. Introduction 6

2. Compliance with OCP Tenets 8

2.1. Openness 8

2.2. Efficiency 8

2.3. Impact 9

2.4. Scale 9

2.5. Sustainability 9

3 Pumped Two-Phase (P2P) Liquid Cooling Solutions 10

3.1 What is a P2P DLC solution? 10

3.2 P2P DLC Solution Advantages 11

3.3 Path of Heat Flow from Chip to Outdoor Heat Rejection 14

3.4 Cold Plates for 2P Liquid Cooling Solution 15

3.4.1 Different cold plate structures 16

3.4.2 Working conditions and configurations 17

3.4.3 Considerations related to cold plates 18

3.4.4 Thermal performance of P2P cold plates 18

3.5 P2P Coolant Distribution Unit (CDU) 20

3.6 P2P DLC Quick Disconnect Couplings 23

3.7 Refrigerant Grade Hose for 2P Cooling Market 26

3.8 Charge management in the closed P2P refrigerant loop 27

3.9 Flow distribution in a P2P system 28

3.10 Rack liquid to two-phase manifold pressure regulation 29

3.11 Liquid and Vapor Filtration 29

Date: February 14, 2025

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

3.12 Secondary Fluid Network (SFN) 29
3.12.1 Reliability of Secondary Fluid Network 31
4 Coolants/Refrigerants for Pumped Two-Phase Cooling 32
4.1 Medium-Pressure Fluids/Refrigerants 34
4.2 Low-Pressure Fluids/Refrigerants 36
5 Application and Installation Support 36
5.1 Available Documents for Installation 37
6 Conclusions 38
7 Glossary 38
8 References 40
9 License 42
9.1 Creative Commons 42
10 About Open Compute Foundation 43

Date: February 14, 2025

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

1. Introduction

The transformation of the enterprise through generative AI and machine learning is poised to have enormous impact to the data center over the coming years. As technology leaders implement AI solutions to capitalize on available innovation, the need for capable hardware and software deployments correspondingly expands. The chips upon which AI platforms are built - GPUs, accelerators, and CPUs – are being designed to meet the intensive demands of training and inference workloads. These critical hardware building blocks, along with the servers, racks, clusters, and data centers that house them, are by nature becoming more feature rich, more power hungry, and more thermally dense than ever before.

Accelerators present one of the greatest challenges in the AI HW landscape, with individual chips today boasting thermal design power (TDP) up to 1000 Watts. In addition to extreme power dissipation, accelerators are comprised of 2.5D and 3D heterogeneous packages, resulting in complex component integration, challenging thermal resistance stacks, and highly non-uniform power maps. The example shown in Figure 1 illustrates their compound construction, highlighting stacked high bandwidth memory (HBM), a known thermal bottleneck in advanced accelerator chips. State of the art accelerators include 3D HBM stacks of up to eight dynamic random-access memory (DRAM) chips, with future roadmaps that include 12 and 16 high memory stacks.

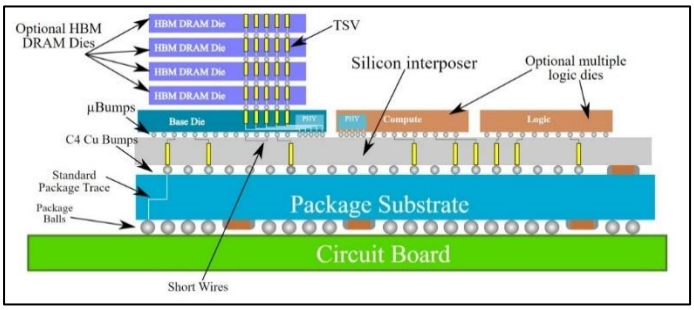

Figure 1. Schematic of a typical accelerator package

The ongoing growth in demand for AI compute is expected to push a continuous increase in accelerator TDP in the coming years. As seen in Figure 2, accelerator TDP exploded by 2.5 times

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

(2.5X) between 2020 and 2024. Assuming this trend continues linearly, accelerator TDP of 2000W by 2030 is not unlikely.

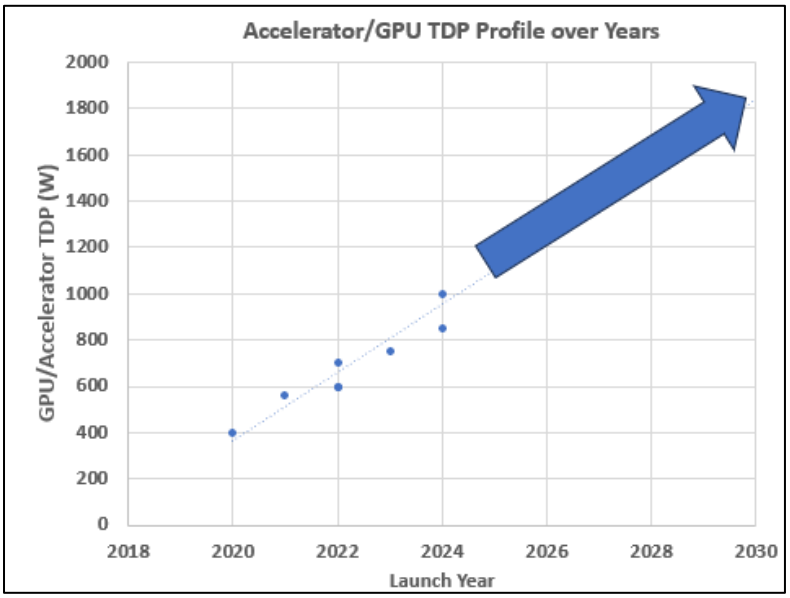

Figure 2. TDP growth of accelerators over the years

Along with raw compute horsepower, AI workloads will continue to demand more and more memory bandwidth and capacity from accelerator chips. This will require integrated packaging with more power-hungry DRAM chips and higher HBM stacks, which translates directly to massive cooling challenges. Meanwhile, compounding the issue, individual component temperatures must remain relatively flat as their requirements for long term reliability in operation do not change.

Today's AI architecture packs up to 72 accelerators into a single compute rack along with accompanying CPUs, switches, and other HW. As such, the AI thermal density problem must be addressed at every level from chip through data center (DC). At the same time, DC operators and governmental agencies continue to focus on energy efficiency and sustainability. For example, in 2023 a United States Department of Energy agency called ARPA-E (Advanced Research Projects Agency – Energy) launched the COOLERCHIPS research program designed to investigate energy-efficient, more sustainable cooling solutions. The agency's intent is to fund initiatives focused on helping DC operators to bring down power usage effectiveness (PUE) and water usage effectiveness (WUE) among other key metrics. In tandem, higher standards for such

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

metrics are being mandated by governmental agencies and discrete municipalities worldwide. Further, heat capture and reuse efforts (i.e. for district, swimming pool, or greenhouse heating) continue to gain interest as data centers become larger energy consumers as a portion of overall grid usage.

Addressing all of these challenges concurrently will force greater adoption of liquid cooling in the data center. At present, the most prominent liquid cooling technology gaining traction is single phase (water-based) direct liquid cooling (DLC). While this approach has significant benefits, refrigerant-based pumped 2-phase (P2P) DLC technology has a few distinct advantages in some usages. This paper describes a P2P DLC cooling solution developed with industry partners Intel, Vertiv, Honeywell, and Parker. Its benefits and basic principles of operation will be introduced, including the use of a specific medium-pressure working fluid, R-515B. Comparisons to other cooling technologies will be given. Note that this paper is not intended to favor any specific P2P cooling vendor, but rather to expose the technology's strengths through an at-scale solution demonstration.

2. Compliance with OCP Tenets

2.1. Openness

This whitepaper is devised to share P2P liquid cooling technology benefits and tradeoffs with the OCP community, and to open doors for future collaboration and overall solution ecosystem enablement.

2.2. Efficiency

The sharing of this whitepaper will provide guidance to potential P2P solution vendors, giving them a head start toward development of next generation cooling solutions and resulting in efficient use of resources and time to market relative to AI HW needs. The outlined technology will provide some end users with operational efficiency along with the enhanced cooling performance provided by the technology's working principles.

Date: February 14, 2025

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

2.3. Impact

The impact of this contribution is significant in that it addresses the cooling demand associated with the power-dense HW required to run heavy workloads such as machine learning (ML) and generative AI. Raising awareness of this technology's working principles and key benefits will help increase available solutions, adoption, and reduce time to market for future customer deployments.

2.4. Scale

This whitepaper is the result of collaboration amongst numerous key component ingredient suppliers and system integrators. This helps articulate the technology's scalability and diverse global reach. The shared content is open for community use as-is and can be tailored to specific use cases, value-add models, or end-to-end solutions as needed within the market.

2.5. Sustainability

Use of this technology results in performance and sustainability advantages in critical AI usages. In some instances, the solution allows for warmer incoming facility water than other forms of liquid cooling, resulting in warmer water leaving the facility (for a given IT load) than that of competing technologies. This shift in facility water temperature can be used in some climates to reduce chiller usage, a main contributor to DC energy consumption. Additionally, the warmer return water can be better utilized for waste heat capture and corresponding downstream usages such as district heating. These advantages help reduce the data center's overall environmental impact.

Date: February 14, 2025

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

3 Pumped Two-Phase (P2P) Liquid Cooling Solutions

Liquid cooling of high-performance compute (HPC) information technology equipment (ITE) is not as ubiquitous as air cooling but is gaining traction as AI drives more and more density at all hardware levels. This section discusses a specific type of liquid cooling solution known as pumped 2-phase (P2P) DLC cooling. It will outline the technology's principles of operation, design considerations, inherent pros and cons, component ecosystem, and at-scale viability.

3.1 What is a P2P DLC solution?

Pumped two-phase (P2P) cooling technology is based on the transfer of heat away from power-dense AI HW via phase change of a working fluid (refrigerant) from a liquid state to a vapor state inside of a cold plate directly attached to the hardware. The basic operating principle behind this technology is illustrated in Figure 3 and can be described as follows:

- Refrigerant working fluid is stored in liquid form in a reservoir, typically inside the cooling distribution unit (CDU).
- Refrigerant is conveyed to heat generating ITE via row and rack manifolds downstream of the reservoir via a mechanical pump.
- Liquid refrigerant enters cold plates within the ITE via serviceable quick disconnect fittings.
- The refrigerant inside the cold plates absorbs heat from thermally critical components, converting some percentage of the liquid into vapor. This process is iso-thermal due to latent heat of vaporization, and results in a very high heat transfer coefficient at the wall of the cold plate.
- The liquid / vapor mixture exits the cold plates and passes through return manifolding and into a heat exchanger inside the CDU. (Note: The percentage of vapor in the saturated mixture is also referred to as vapor quality.)
- Heat from the liquid-vapor mixture is transferred to the ambient medium, either air or facility water depending on CDU design choice, via condensation, whereby the refrigerant returns to 100% liquid state.
- Liquid refrigerant is returned to the reservoir, where the cycle begins again.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Figure 3. Schematic of Pumped 2 Phase (P2P) Direct Liquid Cooling (DLC) technology

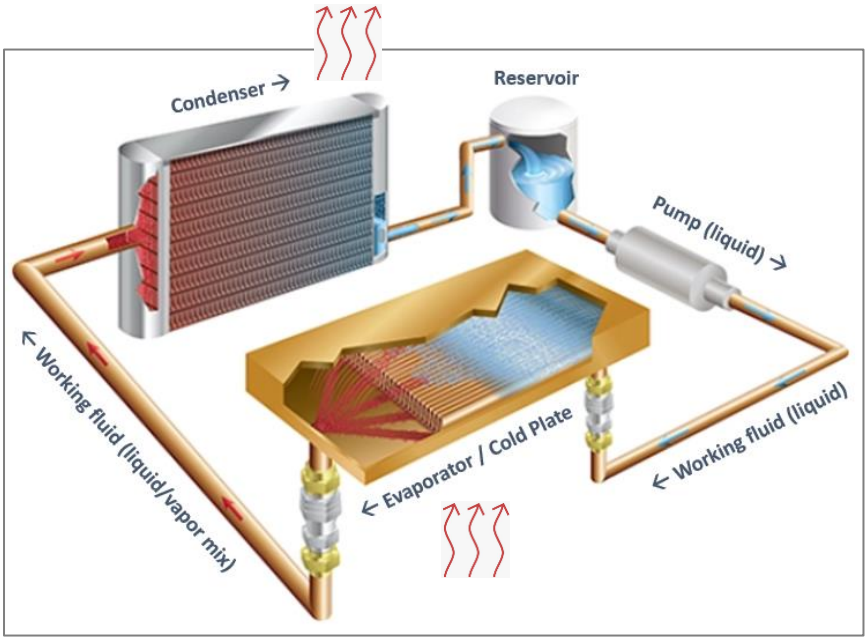

Note that P2P DLC can be achieved with a variety of refrigerant choices depending on use case and end customer goals. Refrigerant choices and corresponding impacts are discussed in a later section of this paper.

3.2 P2P DLC Solution Advantages

This section describes the high-level advantages of 2P DLC cooling. More details behind each claim are further described in subsequent sections of this paper.

- **Performance:** P2P DLC offers best-in-class cold plate wall heat transfer due to boiling of refrigerants at a constant saturated temperature. [11] Cooling of high TDP and high heat flux processors is made possible provided there is sufficient refrigerant mass flow to ensure cold plate leaving refrigerant quality is less than 100%. Technology demonstrations have proven the required mass flow for processors exceeding 2 kW and servers exceeding 10 kW in 1 OU can be easily supported within typical ITE and rack volumetrics. Further details of P2P thermal performance can be found in a later section.
- **Reliability against corrosion & bio growth:** In single phase (1P) DLC cooling technology applications, corrosion of internal fluidic surfaces and bio-organic growth in

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

fluidic elements can be a major problem. The liquid cooling industry has addressed these issues with cooling fluid additives – corrosion inhibitors and biocides – as well as careful and costly fluid chemistry maintenance programs. Refrigerants used in P2P cooling solutions are chemically inert, eliminating internal corrosion concerns, and don't support internal biological growth. Such fluids have been used in closed systems within the HVAC industry for many decades. While P2P DLC systems are closed, the potential for moisture ingress is introduced at serviceable connection points such as quick disconnects. Two-phase connections themselves are designed to prevent moisture ingress, and added procedural steps can be implemented to further prevent it. As an additional precaution, P2P cooling systems typically include filter driers intended to capture solid contamination and absorb entrained moisture. These effective safeguards are commoditized as part of the existing HVAC industry and are designed to be serviceable and swappable in the field.

- **Leaks**: Refrigerants are highly dielectric; refrigerant leaks therefore will not cause damage to sensitive electronic equipment. In addition, typically used refrigerants flash to vapor (i.e. boil) at data center ambient temperatures, thereby leaving no residue at leak sites. Commercially available leak detection solutions can be used to detect rare leak events. These devices can be tied into building management systems to alert maintenance technicians, who can pinpoint and repair leaks using commercially available handheld sniffer devices. In addition, by following industry standard ASHRAE codes and best practices already in use across the HVAC industry today, P2P solutions can be safely installed and implemented.
- **Increased uptime:** Pumped 2-phase DLC cooling systems can be run at elevated fluid temperatures while still providing adequate component cooling. This means components experience smaller temperature deltas or delta-T (i.e. the difference in temperature between component case or junction and working fluid). Lower delta-T may help increase HW reliability by reducing component fatigue related to swings in device temperature. In addition, use of fluids with high dielectric strength in P2P applications reduces downtime associated with leaks, as may be experienced in single phase (water-based) DLC systems.
- **Improved TCO:** Due to the enhanced heat transfer of P2P DLC solutions, they are in some cases able to take advantage of warmer facility chilled water supplies as

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

compared to competing liquid cooling solutions [11]. This helps reduce overall data center power by reducing chiller usage, computer room air conditioner (CRAC) usage, and/or computer room air handler (CRAH) usage (i.e. lower OPEX). In some cases, use of 2P DLC may eliminate the need for chiller installations altogether (lower CAPEX). System BOM cost for P2P systems (CAPEX) is slightly higher than its single-phase counterparts today, but given similar architecture, layout, and component needs between the two, is expected to be at or near cost parity as the technology and ecosystem mature. As discussed above, factors such as elimination of fluid maintenance plans and limited impact of leaks in 2P systems help to bring down their overall cost. Finally, use of warmer facility water also has a "futureproofing" effect on AI data center infrastructure, allowing its extended use and capitalization as hardware density continues to increase with time.

- **Environmental friendliness:** Many commercially available refrigerants used in P2P DLC applications have low global warming potential (GWP) and zero ozone depletion potential (ODP). An emerging concern over per- and polyfluoroalkyl substances (PFAS) may play a key role in the future use of fluorinated refrigerants commonly used in P2P DLC solutions. Related restrictions and regulations could impact many industries including comfort cooling, automotive, and industrial refrigeration. P2P DLC solution providers are expected to respond to future PFAS restrictions in the same way as other industries, as their HFO usage and application are fundamentally identical.
- **Sustainability:** The benefits discussed earlier under "optimized TCO" can also be applied toward enhanced sustainability. The ability to use warmer incoming facility water reduces chiller usage in certain climates, thus reducing a major contributor to data center power usage. This may also allow reduced water usage via replacing evaporative cooling towers with dry coolers. In addition, use of warmer incoming facility water results in a corresponding increase in leaving facility water, making wastewater streams more efficient for waste heat capture applications.
- **Safety**: Several safety factors must be considered when choosing refrigerants and designing systems for pumped 2-phase cooling applications. These are described in detail in Section 4 of this white paper. To summarize, many suitable P2P candidate refrigerants are categorized as having low toxicity and low to zero flammability per ASHRAE 34 standards. Pumped 2-phase DLC systems themselves include standard

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

safety features such as pressure relief valves on closed circuits, reducing the risk of catastrophic events. Further aspects of safety can be managed at the data center level under applicable, industry standard building codes such as ASHRAE Standard 15. In addition, safe handling of refrigerants is performed by certified technicians in the HVAC industry and can be similarly performed by data center technicians through training in certain regions (e.g. as provided in the United States by the Environmental Protection Agency (EPA) under Section 608 of the Clean Air Act.) Combining all these considerations leads to the safe application of P2P cooling in the data center.

- **Availability:** Many of the components used within P2P DLC cooling solutions are directly leveraged from the HVAC industry. Examples include flexible refrigerant tubing, valves, fittings, reservoirs, condensers, and brazed plate heat exchangers. Given their ubiquitous global usage, a rich ecosystem of parts for use in P2P cooling systems already exists. Additional work has been done to tailor and optimize certain components, like fluidic quick disconnects, for data center cooling applications.
- **Ease of adoption**: As with single-phase DTC cooling, pumped 2-phase cooling can be adopted in existing, air-cooled Brownfield data centers using air assist liquid cooling (AALC), also known as side cart cooling. In these applications, heat is rejected into the air stream by circulating refrigerant through condenser coils in a CDU adjacent to ITE racks with integrated DLC cold plates. In this way, the performance of P2P can be harnessed to cool dense AI HW in legacy environments.

3.3 Path of Heat Flow from Chip to Outdoor Heat Rejection

Figure 4 illustrates the flow of heat from server components (chips) to the outside ambient environment in a P2P system. The P2P thermal transport solution consists of coolant distribution units (CDU) shown in blue boxes, and secondary fluid networks (SFN), shown with black arrows. The DLC thermal solution transports heat to condensers where the heat is transferred from the refrigerant to the primary fluid, thus isolating primary and secondary fluid streams. The primary fluid in chiller-equipped data centers is the facility cooling system liquid. To enable liquid cooling deployment in an air-cooled data center, the conditioned air of the white space acts as the primary fluid. Only one CDU/SFN is necessary. Figure 4 shows the various configurations available.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

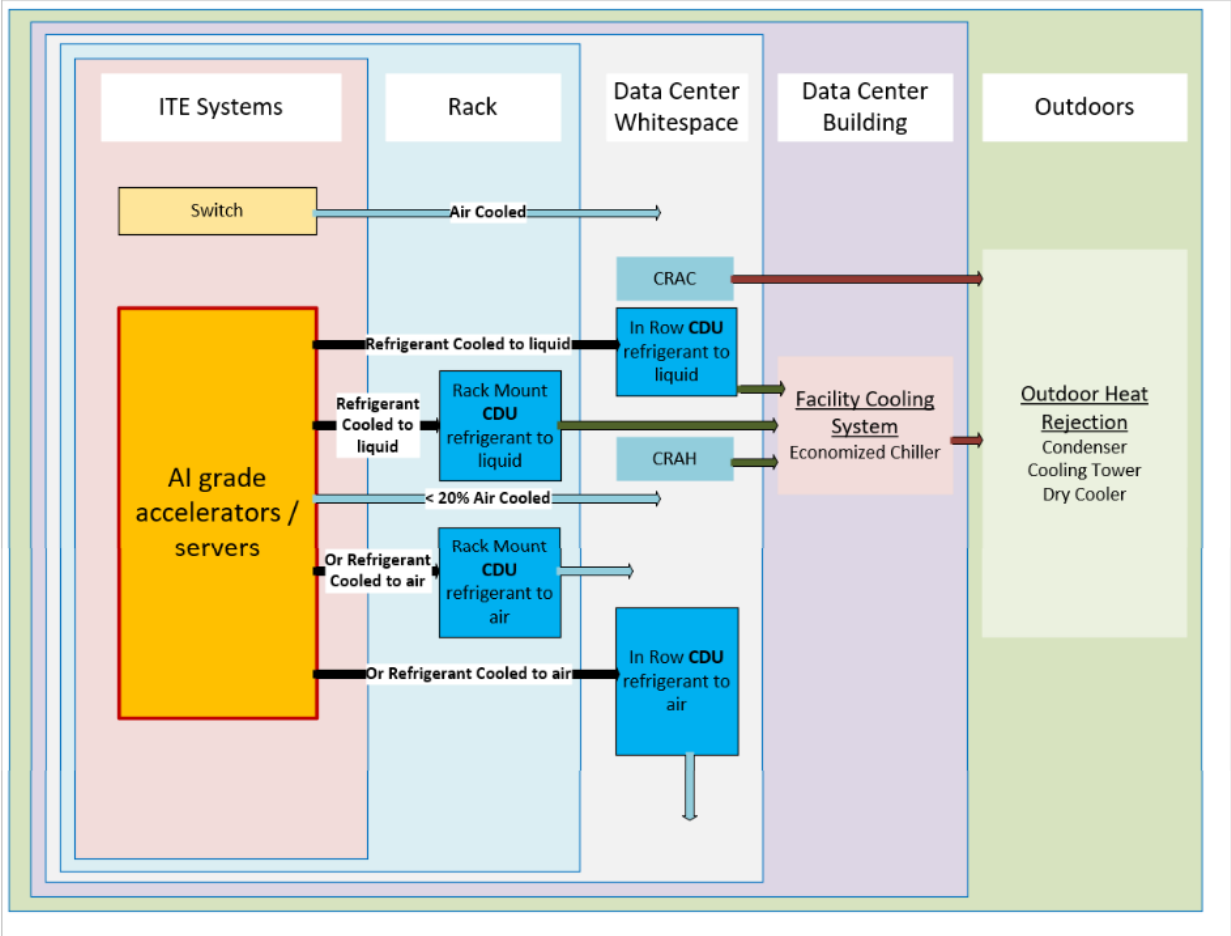

Figure 4. Pumped Two-Phase Thermal Transport Path

3.4 Cold Plates for 2P Liquid Cooling Solution

In P2P systems, cold plates are attached directly to heat-generating components. Refrigerant is delivered to the cold plates in liquid or mixed vapor / liquid form, where added heat generates more vapor. Many manufacturing techniques can be used to create cold plates used in P2P systems; the final design chosen can be tailored to the application in terms of cooling demand and cost. This section will discuss various cold plate types and their associated operating conditions and configurations, design trade-offs, and thermal performance results.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

3.4.1 Different cold plate structures

Due to the high heat transfer coefficient offered by phase change heat transfer, it is likely that in P2P DLC cold plates, the convective/phase-change thermal resistance is negligible. Consequently, internal geometries used within traditional single-phase cold plates can also provide sufficient performance in corresponding two-phase cold plate designs; this provides a good starting point for design and deployment of 2-phase cold plates, provided the increased pressures associated with 2P cold plates are accommodated. Microchannels formed by parallel straight fins with sub-mm feature size (fin thickness) and spacing are commonly used for P2P cold plates due to their relatively mature application in single-phase cooling and ease of manufacturing. Skiving on copper is the most widely employed way to make microchannel cold plates. The form factor of the cold plate can be designed to fit the form factor of the processor being cooled, as shown in Figure 5. By varying the geometry of the cold plate top cover, the direction and characteristic length of the internal refrigerant flow can be dictated, and accordingly, the particular two-phase heat transfer mode can be controlled to be pool boiling, flow boiling, jet impingement, thin film evaporation, or a combination of multiple mechanisms, depending on the cooling heat flux/capacity needed and the thermal resistance required for different components.

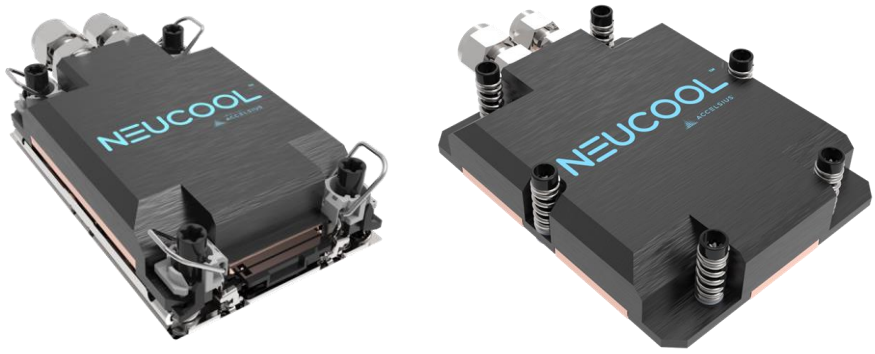

Figure 5. Cold plates with form factors matching different chip packages: Intel Sapphire Rapids and Emerald Rapids (left); AMD Genoa (right) [Courtesy: Accelsius]

While skived fins are easily and cheaply produced, alternative manufacturing techniques can also be used to create more complex 2P cold plates at higher cost. For example, computer numerical control (CNC) machining can produce sophisticated internal structures such as segmented microchannels or micropillar/pin fin arrays [1, 2]. Sintering metallic particles, sometimes with 3D

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

structured molds or dies, can realize hierarchical porous structures with desired porosity and pore size for nucleation, capillary feed, and fluid transport [3]. Additive manufacturing can offer significant structural freedom and allow the manufacturing of internal manifold or even topology-optimized structures for more efficient flow and heat transfer [4, 5]. These manufacturing methods have been employed for R&D purposes both in academia and in the industry, but large-scale application is not yet realized and is likely only possible when newer generations of processors with more stringent cooling requirements warrant their higher cost.

3.4.2 Working conditions and configurations

For P2P systems, flow control and regulation are of vital importance to ensuring sufficient cooling of the IT equipment, and for avoiding disastrous thermal runaway. When the liquid refrigerant flow into a cold plate is completely or partially vaporized by the heat generated on the processor, the exit refrigerant flow is composed entirely of vapor or is a mixture of liquid and vapor. If at any point the exit quality is 100% (i.e. all vapor), any additional heat generated from the processor must be dissipated by vapor convection inside the cold plate. This is extremely inefficient and will likely cause the processor to overheat or shut down. This condition is known as cold plate “dry out” and should be avoided. This means that P2P systems must be designed such that, for a given thermal load, sufficient mass flow is provided to avoid dry-out. On the other hand, when exit vapor quality is too low, subcooled single-phase liquid convective contribution becomes significant, and the thermal resistance is higher than fully developed boiling conditions. For these reasons, P2P systems should be designed to target vapor qualities of ~70%.

In most cases, multiple processors and other heat sources exist in one server, thus driving a potential need for multiple cold plates. In P2P systems, this array of cold plates can be arranged in parallel or series. For a parallel arrangement, each flow path is equivalent, and under ideal conditions the same thermal performance can be expected from each cold plate. For a serial arrangement, due to the pressure drop along the flow, downstream cold plates will have a lower saturation temperature compared to upstream cold plates, which tends to lower the case temperature downstream. Meanwhile, the average vapor quality of the downstream cold plate is higher than that of the upstream cold plate. Depending on the operating vapor quality range, the downstream cold plate could show worse thermal performance if the quality is large (approaching

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

dry-out) and better thermal performance if the quality is small (less contribution from subcooled single-phase convection). These two effects can cause case small temperature differences among processors cooled by serial cold plates (although much smaller than single-phase cold plates arranged in series due to the caloric temperature rise of water with added heat). Parallel cold plate arrangements require server-level flow balancing to eliminate flow maldistribution amongst flow paths under non-uniform heating conditions. Such flow balancing requires the use of flow regulators [6] or flow restrictors [7]

3.4.3 Considerations related to cold plates

As the cold plates are directly attached to the chip packages, higher internal operating pressure may cause bulging of the cold plate and result in excessive stress on the chip. Hence during the design of cold plates, structural analysis is of utmost importance. Regarding material selection, a cold plate assembly typically includes a base plate with heat transfer enhancement structures, and a lid with geometrical designs to direct and route the flow. These components are preferably made of metal, which mitigates potential working fluid leaks. The base plate is usually made of copper due to its high thermal conductivity. Certain elastomers can also be used as seals or gaskets as long as they are compatible with working fluids (i.e. do not shrink or change shape with exposure and/or heat.) Compatible non-metallic materials can even be used as the lid material for low-pressure working fluids.

3.4.4 Thermal performance of P2P cold plates

Investigation of two-phase heat transfer has been fruitful in the research community. For example, hybrid two-phase cooling, which couples pumped two-phase cooling with capillary-driven two-phase cooling, was demonstrated with a heat flux of 970 W/cm^2 using R245fa as the working fluid [3]. P2P cold plates can be enhanced with sintering powders [3] or using additive manufacturing methods [4]. These results show remarkable performance of two-phase cooling and the potential to meet the AI HW cooling demands.

Recent P2P DLC solutions have been shown to have cooling capability of up to 2200W per socket using thermal test vehicles (TTV) in Intel Sapphire Rapids and Nvidia H100 form factors [8]. Such solutions used refrigerant R1233zd as the working fluid. Figure 6 shows performance results for

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

this solution with different refrigerant flow rates. Note that the demonstrated performance variation at certain fixed power and heat flux levels within this data are the result of amount of single vs. two phase phenomena present (e.g. lower performance below 1000W in the presence of higher flow rate) or varying effects of localized dry out with lower mass flow and increased power and/or heat flux (e.g. above 1000W and/or 120 W/cm^2). A maximum dissipated heat flux of ~250 W/cm^2 was demonstrated. These cold plates were later tested in real servers with 350W TDP CPUs, and the recorded junction temperature was well below the throttling point with 35 °C facility water [8].

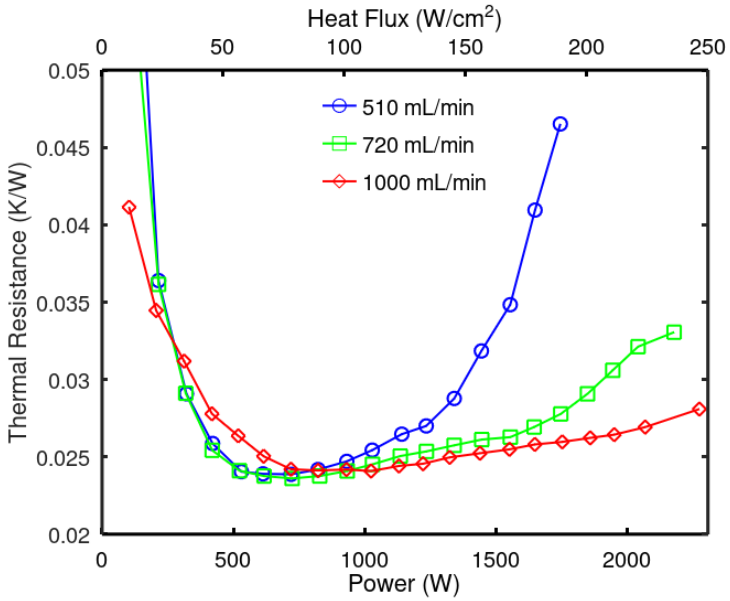

Figure 6. Cooling performance of the skived-fin cold plate on a thermal test vehicle with Nvidia H100 form factor

Kulkarni et al. [11] recently published a paper illustrating thermal performance of a cold plate that using two different working fluids: PG25 (single phase) and R-515B (2-phase). Heat flux and TDP were varied using a TTV. Figure 7 shows the direct performance comparison of these two working fluids as applied to the same cold plate. Note that thermal performance for 1P DLC is derived using fluid inlet temperature as its reference, while thermal performance for 2P DLC references the saturated temperature of the liquid/vapor mixture exiting the cold plate. As seen in the figure, thermal performance of the 1P cold plate is dependent on the fluid flow rate passing through the cold plate. For the 2P cold plate, refrigerant flow rate has no impact on thermal

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

performance provided the liquid remains in saturated condition throughout the test (i.e. the cold plate does not experience dry-out.) This helps illustrate the iso-thermal performance of 2P liquid flows in cold plates, which in turn helps maintain consistent temperatures across heterogeneous packages. (e.g. Each high bandwidth memory, or HBM, stack within a package sees the same fluid temperature despite its location within the flow path. This may become critical for next gen AI processors with many integrated HBM stacks.) Kulkarni et al. also showed the impact of heat flux on thermal performance for each design [11].

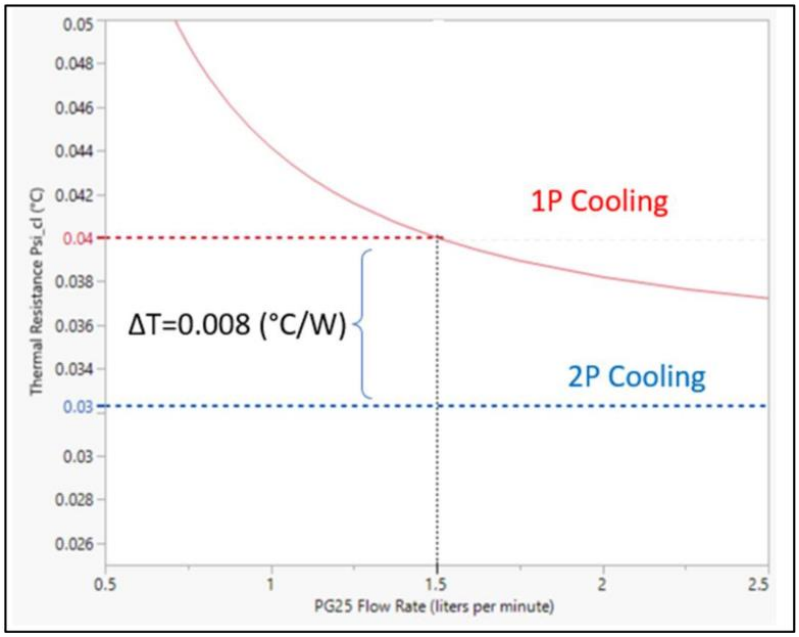

Figure 7. Thermal performance comparison between 1P and 2P cold plate on a uniform heat flux Thermal Test Vehicle (TTV) [11]

3.5 P2P Coolant Distribution Unit (CDU)

Four P2P coolant distribution unit (CDU) configurations are described in this section. These are as follows:

1. **In-row refrigerant-to-liquid**: This CDU type is designed to handle multiple racks. Current designs have a rated cooling capacity of 400 kW with 45°C incoming facility water temperature. These CDUs are typically installed adjacent to the ITE racks they cool and

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

remain exposed to cold aisle air. The racks support an overhead secondary fluid network (SFN) that is installed in-situ at system commissioning using industry standard no-braze fittings as shown in Figure 9. Racks are part of the hot aisle / cold aisle containment system. A rendering of this style of CDU is shown in Figure 8.

Figure 8. In-Row Refrigerant-to-Liquid CDU with Racks and SFN

Figure 9. CDU to SFN No-Braze Connector

2. **In-row refrigerant-to-air:** This type of CDU handles fewer racks and has a typical cooling capacity of 180 kW with 35°C ambient air. They are installed immediately adjacent to the IT racks they cool and pull cold aisle air through the front of the CDU, then reject the heated air into the hot aisle. The racks support an overhead secondary fluid network

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

(SFN) that is installed in-situ at system commissioning using industry standard no-braze fittings as shown in Figure 9.

3. **In-rack refrigerant-to-liquid:** This CDU type is dedicated to cooling the rack in which it is installed. They have a typical cooling capacity of 100 kW. They are typically installed at the bottom of the rack, below the cooled ITE. Facility side connections are at the rear of the CDU and provide primary liquid flows for heat rejection. Renderings of this style of CDU are shown in Figures 10 and 11.

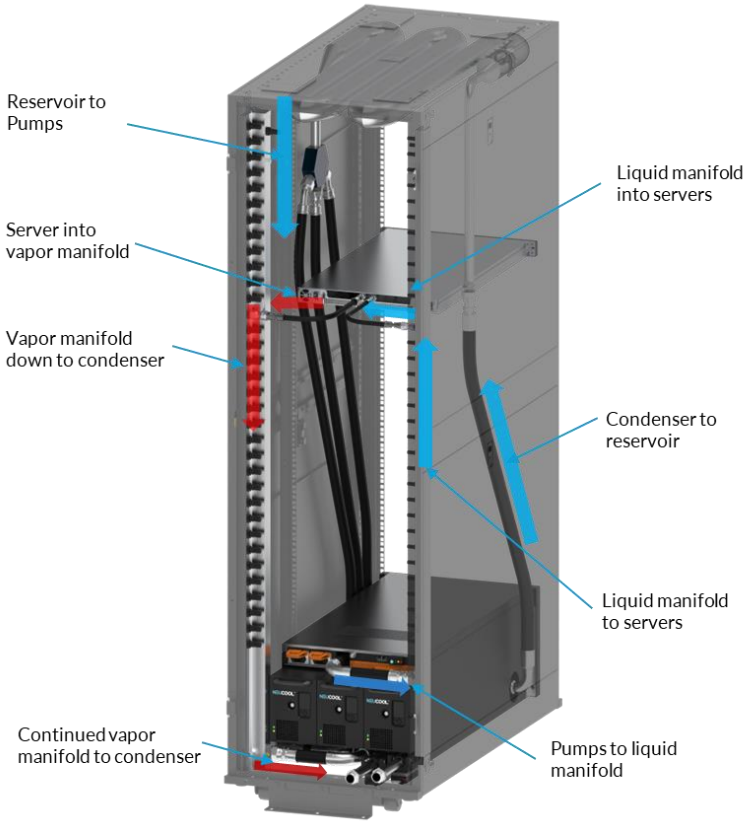

Figure 10. Example of In Rack Liquid to Refrigerant CDU (Courtesy: Accelsius)

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

Figure 11. The front (left) and rear (right) view of the in-rack CDU system.

4. **In-rack refrigerant to air**: This CDU type is dedicated to cooling the rack in which it is installed. They have a typical cooling capacity of 40 kW. They typically include a pumping unit installed at the bottom of the rack, below the cooled ITE, and a condenser and fans in a rear door (RDHX). Air from the cold aisle is pulled through the rack and condenser coils in the rear door with heated air rejected to the hot aisle.

3.6 P2P DLC Quick Disconnect Couplings

Quick disconnects have been widely manufactured and used in the HVAC industry for many decades. One example of such a coupling is the Parker Sporlan 5400 series connector, a dry break, self-sealing coupling for use with refrigerant. These can be found in products spanning nearly every cooling technology within a data center: CRAC, CRAH, AHU, chiller, etc. As data centers shift to P2P DLC, this component will play another critical role in allowing for service and hot swap of servers in an ITE rack. Building upon this established usage, Parker Sporlan is developing a connector more tailored to the data center use case – see Figure 12. This next-gen coupling enables toolless service at operating pressure, reducing install times by up to 70%. This connector provides all performance features of the current generation product including minimal air inclusion, minimal fluid loss, high vacuum capability, end connection adapter compatibility, corrosion resistance, wide temperature range, and high-pressure robustness. The next

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

generation coupling utilizes special neoprene seals that are compatible with all common refrigerants including R134a, R454B, R-515B, R1233zd, R1234yf, and R1234ze.

One of the key characteristics of a coupling is the ratio of flow rate performance to coupling size. The desired outcome is that the coupling takes up a minimal amount of space with maximum flow performance. Coupling size can be defined by two measurements: diameter and length. The coupling diameter is pre-defined by the coupling's nomenclature of -4, -6, -8, -12, etc. The next-gen refrigerant coupling's length was designed as compact as possible. Its volumetric flow rate performance can be found in Figure 13 with R-515B and Figure 14 with R-1233zd(E).

As next generation couplings are based on standard 5400 series QDs, production capacity has been maintained. The next gen coupling will use existing assembly lines to guarantee scalability when the demand for P2P DLC increases.

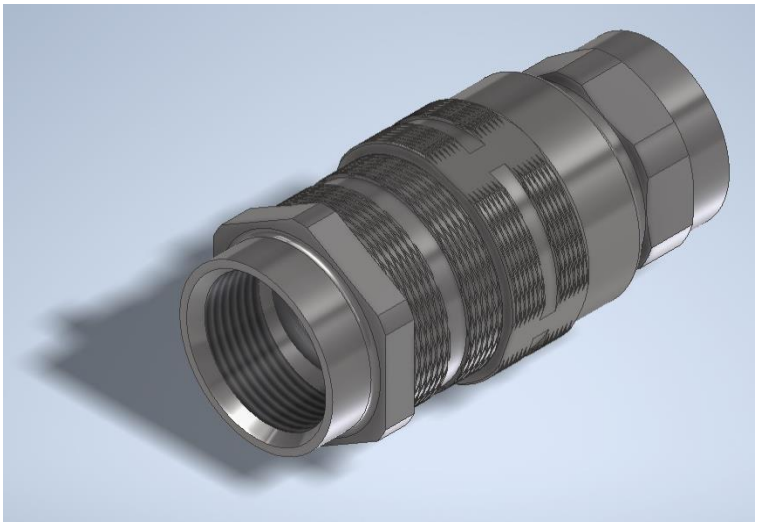

Figure 12: Example of a Next Gen Data Center Quick Disconnect for P2P Technology

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

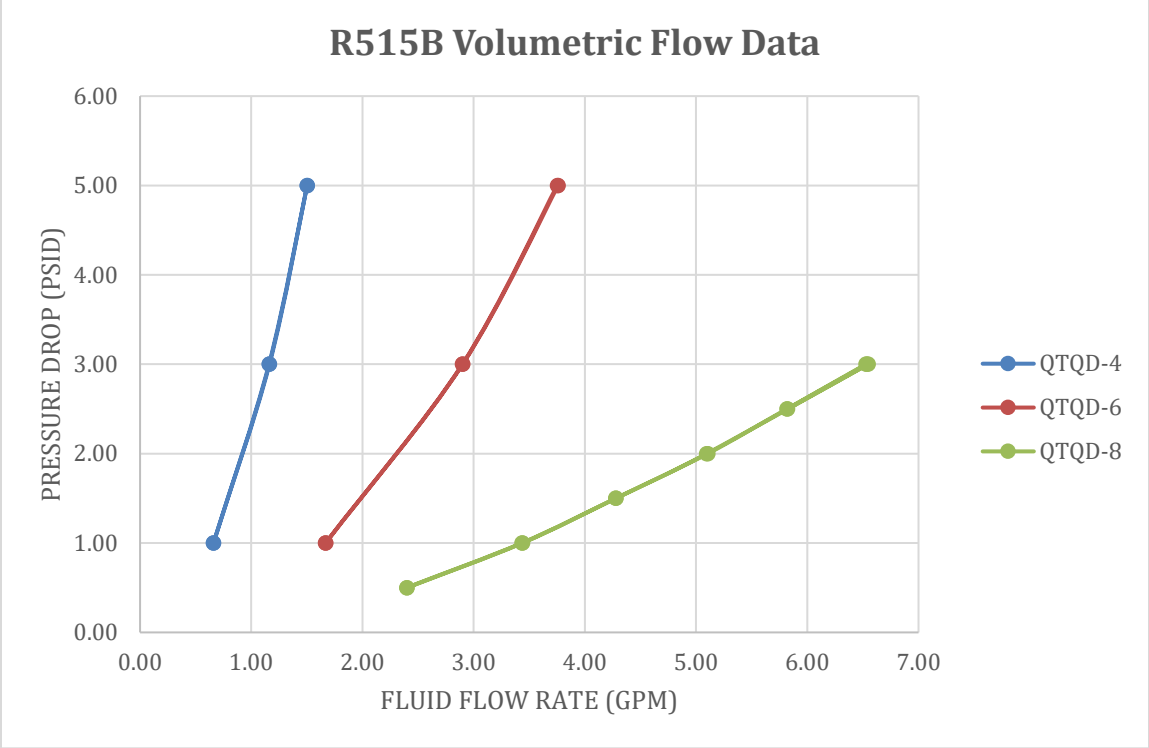

Figure 13. Pressure drops of different sizes of QDs for R-515B refrigerant.

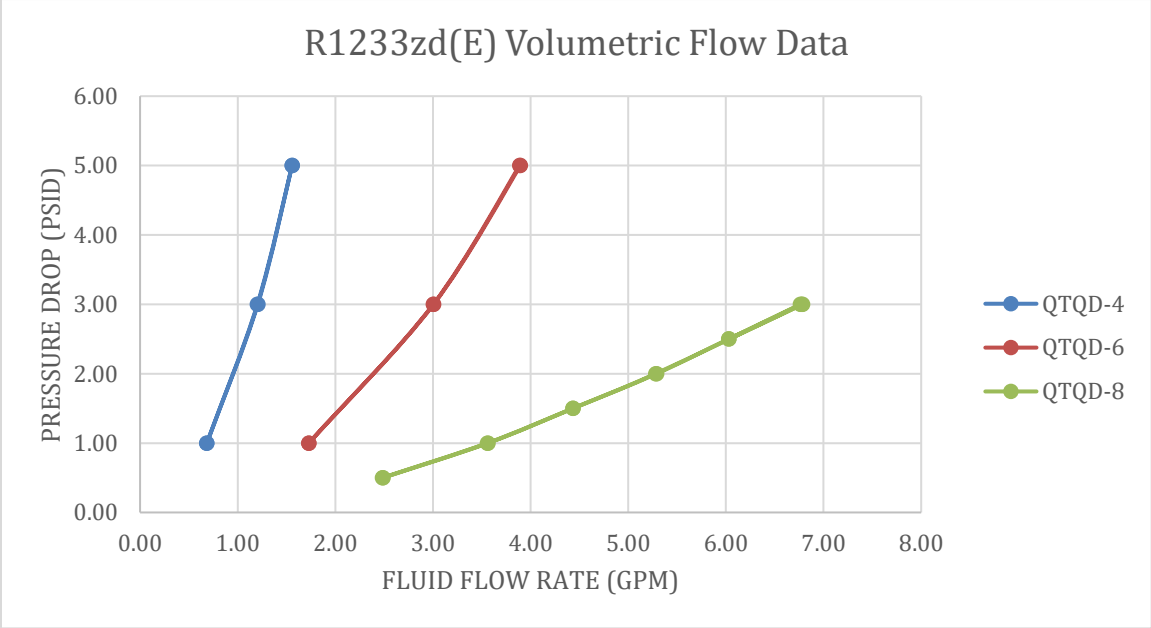

Figure 14. Pressure drops of different sizes of QDs for R1233zd(E) refrigerant

Date: February 14, 2025

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

3.7 Refrigerant Grade Hose for 2P Cooling Market

The use of AC barrier hose for refrigerant cooling has been around for many years. With ever-evolving advancements in server technology, improving server cooling requirements is essential for today's high-power IT gear. Rubber barrier hoses are the perfect fit for this need, providing a means to convey the refrigerant through the system and providing flexibility to enable tight routing and ease of installation.

AC barrier hoses have been in use for more than 30 years and have been utilized by various customers and applications including the HVAC, automotive, truck, and agricultural markets. These industries see a wide variety of environments and climates. For these products to be approved in these markets, they must conform to the SAE J3062 standard. This standard sets the performance requirements of the barrier hose construction so that it can be used in a wide range of air-conditioning system environments, functioning between -30 °C to 125 °C. With various requirements around flexibility, flow, and pressure drop, AC barrier hose is offered in sizes ranging from 5 mm to 25 mm ID. This allows customers to design the cooling loop for optimal heat transfer.

An AC barrier hose is a multi-layered textile-reinforced hose consisting of a rubber tube, a thermoplastic barrier, a rubber jacket layer, reinforcement, and a rubber cover. The tube consists of a rubber material that is compatible with the media being used in the application. An extraction test is performed per SAE J3062 to ensure the refrigerant does not pull an unacceptable amount of extractable from the tube material. Excess extractable material could foul heat exchanger surfaces or restrict narrow flow paths. The tube also provides a robust seal to prevent leakage from the hose barb interface.

Following the tube layer, a thin thermoplastic barrier layer is applied. The barrier layer is required to provide permeation resistance for the refrigerant that is being used. Refrigerants will naturally permeate through the layers of the hose. The purpose of the barrier is to slow this rate down to an acceptable level as required. Limiting refrigerant loss due to permeation is important for closed-loop cooling designs. Although many of the refrigerants being considered for direct-to-chip cooling loops operate at a lower pressure than traditional refrigerants deployed in industrial vapor-

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

compression applications, the permeation resistance of the hose should still be considered. This is because the permeation rate of a refrigerant through a polymer barrier is strongly correlated with the refrigerant molecule size and the temperature of the application – not simply the system pressure. Figure 15 shows the permeation limits for each hose type.

Parker AC barrier hose is a Type C since it is a barrier, textile-reinforced construction. However, the AC barrier hose performance of 0.34 kg/m^2/yr (-6 size) with R134a puts it in the Ultra Low Permeation range resulting in a Type Cu permeation rating. R-515B refrigerant performs to a permeation rate of 0.39 kg/m^2/yr which also categorizes it as Type Cu. The reinforcement layer provides strength for the hose. Typically, in AC hoses this is a textile reinforcement that is designed to provide burst performance and to ensure structural integrity of the product during flexing/bending. This allows the hose to keep its form while being able to route through the desired path in the cooling loop.

The rubber cover layer completes the construction of the hose. It helps protect the reinforcing layer and provides resistance to abrasion, corrosion, and other elements as required. The cover material is also compounded to provide a protective barrier that prevents moisture ingress into the cooling loop.

	Refrigerant	
Hose Type	R-134a	R-1234yf/R-152a
A, B	15 kg/m^2/year	18 kg/m^2/year
C, D, E, F, G	5 kg/m^2/year	5 kg/m^2/year
C_u, D_u, E_u, F_u, G_u	1.5 kg/m^2/year	1.5 kg/m^2/year

$_u$ Designates ultra low permeation limit.

Figure 15: Permeation Limits

AC Barrier hose is an ideal fit for the emerging 2P cooling market. This hose will be more than capable of meeting and exceeding performance requirements for the operating conditions and environment in which it will be utilized.

3.8 Charge management in the closed P2P refrigerant loop

For P2P, the refrigerant is pumped continuously through the system. When there is no IT load, no vapor is generated, and the recirculation system is filled with liquid refrigerant. At full IT load,

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

the cold plates, all return lines to the condenser, and half the condenser are filled with vapor. The displaced liquid is stored in a make-up chamber; typically, refrigerant receivers are used for this purpose. Sizing this receiver is dependent upon the internal volume of the cold plate, return hoses and manifolds, and ½ the refrigerant side of the condenser. Whether CDU and SFN are from a single supplier or multiple, a coordinated integrated design is necessary for proper operation.

3.9 Flow distribution in a P2P system

P2P refrigerant pumps provide a constant pressure head at a constant flow. An IT cooling loop includes an inlet quick disconnect (QD), flow regulator, cold plate array, interconnecting tubing, and outlet QD. At a constant mass flow rate, the pressure drop (dP) across the cold plate varies with the IT load. To maintain the design mass flow rate as the IT load varies, a flow regulator providing constant volume flow for a range of differential pressure devices [1] is applied at the entrance of each cold plate array. The flow regulator establishes a consistent pressure differential across each circuit at varying IT loads enabling proper flow distribution. Figure 16 shows the typical pressure flow response of a flow regulator.

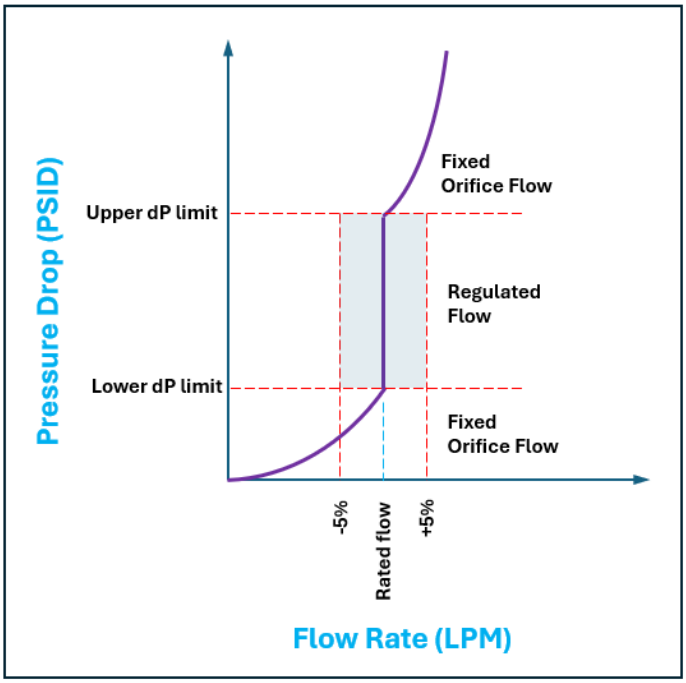

Figure 16. Typical Flow Regulator Pressure Flow Curve

Date: February 14, 2025

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

3.10 Rack liquid to two-phase manifold pressure regulation

A constant refrigerant pressure differential across the array of cooling loops is achieved using a pressure-regulating bypass valve installed between the liquid supply and two-phase return row manifold lines.

3.11 Liquid and Vapor Filtration

A desiccant drier is included in the flow path to capture trace moisture not removed during the evacuation process. Geometries of cold plate heat transfer surfaces approach 100 µm. A 25 µm filter is included in the flow path. Figure 17 shows a typical installation of a filter drier assembly. The canister and bypass piping assembly enable changing the filter / drier media without interrupting system operation.

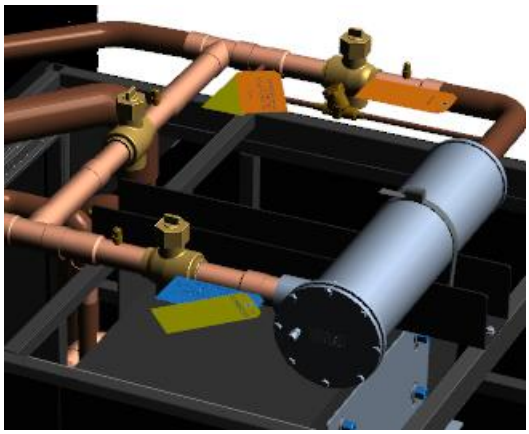

Figure 17. Replaceable Core Filter Drier Assembly

3.12 Secondary Fluid Network (SFN)

The SFN consists of a liquid supply row manifold, a liquid supply rack manifolds, liquid supply hoses, mating QDs, two-phase return mating QDs and hoses, two-phase return rack manifolds, two-phase return row manifold, and row-to-rack-manifold connecting tubing. The SFN row manifolds reside atop the racks and below coms and power distribution tracks. A typical row manifold assembly is shown in Figure 18.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

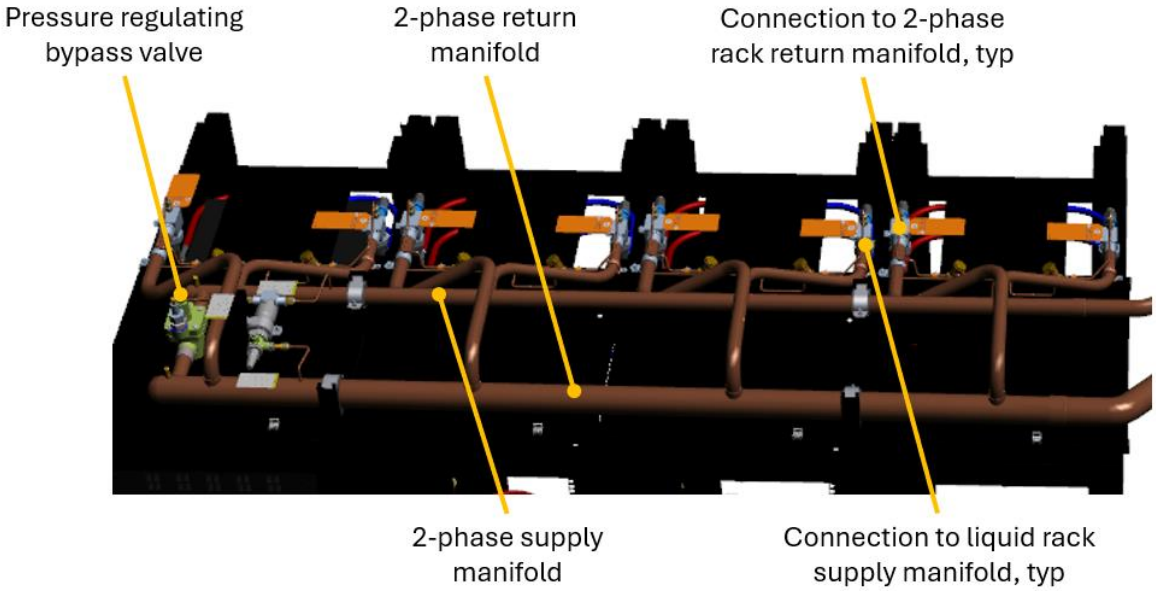

Figure 18. Row Manifold Assembly

At the rack level, the row manifolds are connected to vertical manifolds on the rack as shown in Figure 19. The number of connections and type of connections are defined based on the server-side connections and which type of servers are being installed. The row and rack-level manifolds are usually brazed joints with refrigerant-grade hoses and quick connects.

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

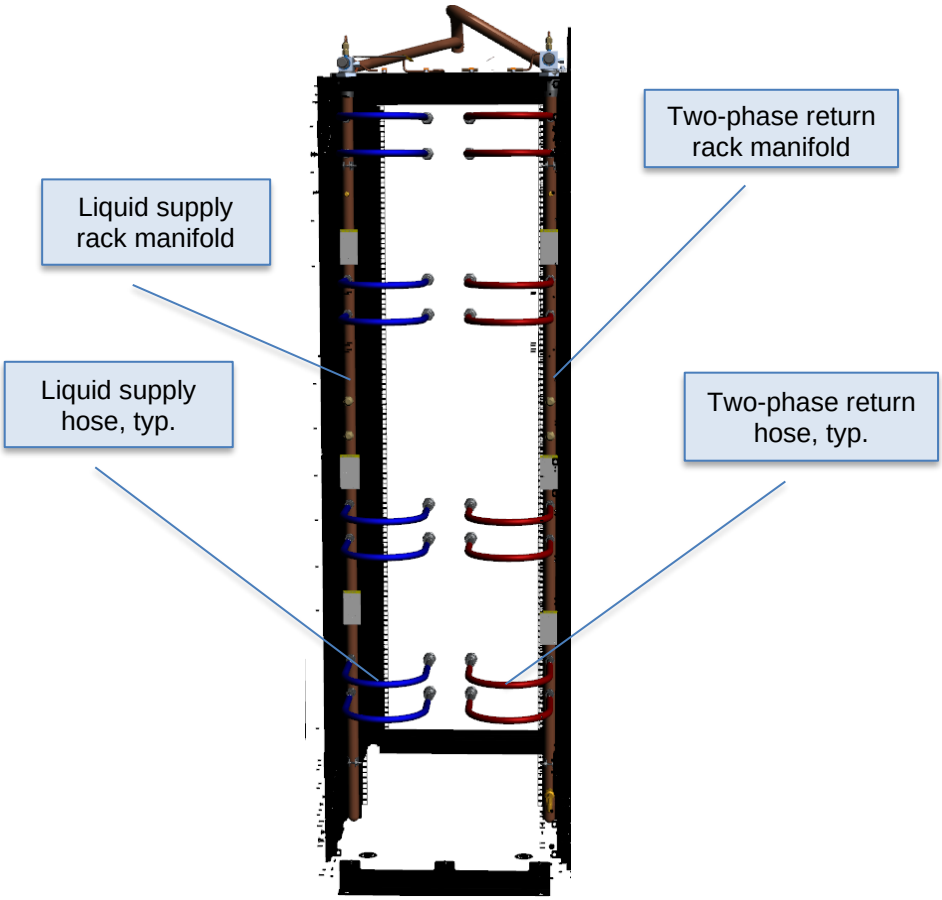

Figure 19. Rack Manifold Assembly

3.12.1 Reliability of Secondary Fluid Network

P2P components have been developed and widely used for similar parallel applications as shown in the following list:

- Refrigerant-to-fluid condensers are brazed plate heat exchangers while refrigerant-to-air condensers are micro-channel coils.
    - Both are state-of-the-art in the HVAC/R industry.
- The pump is a regenerative turbine design utilized for pumped refrigerant economizers in data center facility cooling equipment.
- Receiver, piping, isolation and check valves, pressure relief valves, and high-pressure controls are used in the HVAC/R industry.
- Hoses and QDs are used in split system HVAC equipment.

Date: February 14, 2025

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

- Flow regulators are used in commercial building heating and cooling.

4 Coolants/Refrigerants for Pumped Two-Phase Cooling

Certain fluorinated fluids, referred to as refrigerants, are dielectric and have suitable thermophysical properties for pumped two-phase cooling. However, pressurized leaks of refrigerants can produce large amounts of vapor. Depending on the type of refrigerant used, such leaks can be asphyxiants, toxic, or even cause fires. In addition, legacy refrigerants such as HFCs and HFEs may have high global warming potential (GWP); their release is harmful to the environment. By law, refrigerant manufacturers must provide robust safety data sheets (SDS) along with fluid usage and handling guidelines. Please refer to the specific refrigerant manufacturer for these documents prior to designing, installing, operating, or maintaining pumped 2-phase cooling systems to ensure safety to users, operators, building infrastructure, and the environment.

ASHRAE standard 34 classifies refrigerants based on their level of toxicity and flammability. Classes A and B signify refrigerants with lower and higher toxicities, respectively, based on the Occupational Exposure Limit (OEL) which is a prescribed measure of chronic (8-hr time-weighted average value) toxicity. Flammability class 1 indicates refrigerants that do not show flame propagation in moist air at 60 °C, 8.81 g of moisture /kg of dry air, and 101.3 kPa when assessed per ASHRAE 34 2022. Other classes signify refrigerants with flammability; the distinction in class for flammable refrigerants depends on the burning velocity, lower flammability limit, and the heat of combustion. As compared to class 2 and 3 refrigerants, the class 2L refrigerants are much more difficult to ignite and have very low flammability potential and tend to display lower values for heat of combustion, limiting the risk of building, component, and employee harm. They also tend to display lower burning velocities if an ignition occurs. ASHRAE Standard 34 has also established the maximum refrigerant concentration limit (RCL) in the air to avoid escape-impairing effects such as asphyxiation, acute toxicity, and flammability in normally occupied enclosed spaces. The RCL is established based on the most conservative limit between asphyxiation, acute toxicity, and flammability. This safety group classification is referred to in the installation and maintenance standards for HVAC equipment (e.g., ASHRAE 15, ISO 5149, and EN 378) and building codes to determine the allowable refrigerant charge in the equipment installed inside a building. The standards allow for larger amounts of refrigerant in systems that use non-flammable

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

refrigerants compared to flammable ones. Class 2L refrigerants are more difficult to ignite and pose less risk if an ignition event occurs as compared to Class 2 and 3 refrigerants. Product safety standards for refrigeration and AC equipment, such as UL 60335-2-40 and UL 60335-2-89, have been updated for the safe use of flammable refrigerants. Together, these standards have provided specific restrictions based on the risk profile of the refrigerants. While specific product standards are yet to be developed for liquid cooling in data centers, these standards can guide the industry when designing two-phase systems using class 2L refrigerants. These product and installation standards specify the implementation of occupational safety measures to reduce the risks associated with toxicity and flammability. Refrigerant leak detection sensors should be installed in areas where a leak would gather and trigger an alarm to start one or more safety measures, such as dilution by circulation, mechanical ventilation, and de-energizing potential internal ignition sources. Release mitigation controls, such as automatic shut-off valves, can be used to limit the amount of released refrigerant in the event of a leak. Additionally, handheld refrigerant leak detectors with a long, flexible probe can be used to locate low ppm level leaks for servicing and maintenance. Refrigerant must comply with globally recognized AHRI standard 700 specifications for refrigerant purity and blend composition tolerance.

The Kigali Amendment of the Montreal Protocol went into force on January 1, 2019. It requires a global gradual reduction of HFCs measured in CO2 equivalence that are available in the market. Each country or region has its schedule for reducing HFC availability, such as the F-gas regulation in the EU and the AIM Act in the US. These regulations are already putting significant pressure on HFC refrigerant availability, and only a limited amount of HFCs will be available for servicing and maintenance of existing HVAC equipment by the end of the decade. Detecting leaks is crucial for conserving refrigerants, protecting equipment, improving performance, reducing emissions, and ensuring occupational safety. The EU's F-gas regulation has introduced the requirement for regular leak checks of equipment containing HFCs to prevent their release into the atmosphere and the frequency of leak checks depends on the CO2e of the fluid quantity in the system, so systems with higher GWP refrigerants will require more frequent leak checks. So far, GWP limits and periodic leak checks have only been mandated for refrigeration and air conditioning, and there are no such restrictions for the use of refrigerants for pumped liquid cooling of electronics. However, it is recommended to choose HFC alternatives with environmentally preferable properties for two-phase cooling systems considering the aggressive HFC phase-down schedule

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

in the EU and US. Please reference both applicable industry standards and building codes for your specific site when considering this technology to ensure compliance with local regulations.

4.1 Medium-Pressure Fluids/Refrigerants

Higher GWP, HFC based refrigerant, R-134a is used today for the thermal management of electronics. There is an increasing focus on a new group of lower GWP fluorinated fluids known as hydrofluoroolefins (HFOs) and hydrochlorofluoroolefins (HCFOs). HFOs such as R-1234yf are used heavily in automotive today for new vehicles. These refrigerants are low to ultra-low GWP and will be present in most air side equipment such as residential and commercial AC by 2025 in the US. HCFOs are a class of refrigerants that are used predominantly in chillers today. These refrigerants are low to ultra-low GWP and have a minimal amount of ozone depletion potential. Many end user operations may be using products form this class today if using chillers in their installation. The main refrigerants in these groups are currently R1234yf and R1234zd(E), as well as blends between these and non-flammable HFCs. These HFOs and HCFOs and certain blends have replaced HFCs in several HVAC applications. Table 1 below summarizes the key properties of lower GWP alternatives to R134a. R1234ze and R1234yf have the lowest GWP among R134a replacement options. Both options are mildly flammable per US codes, so additional risk mitigation may be necessary as compared to the non-flammable options, particularly in larger cooling capacity centralized systems.

Table 1. Medium-pressure fluids for pumped two-phase cooling

Refrigerant	AR4 GWP	ASHRAE safety classification	Normal boiling point, °C	Pressure at 50 °C, kPa	Volumetric capacity at 50 °C, MJ/m^3	RCL, g/m3
R-134a	1410	A1	-26	1318	167	210
R-513A	573	A1	-29	1377	135	320
R-515B	293	A1	-19	995	153	290
R-1234ze(E)	1 (AR5)	A2L	-19	997	156	76
R-1234yf	<1 (AR5)	A2L	-29	1302	121	75

R-515B is a non-flammable azeotropic blend of HFO R-1234ze(E) and HFC R-227ea. It is the lowest GWP non-flammable replacement for R134a and is now widely used in a wide range of HVAC systems. Its refrigeration concentration limit (RCL) value is almost 4 times that of R1234ze(E) and R1234yf, which is particularly advantageous in high-density data centers. Its similar product design to R1234ze(E) provides a long-term solution that is suitable for both

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

medium temperature A1 and A2L refrigerants. The thermal resistance of a microchannel cold plate supplied with different medium-pressure refrigerants has been reported by Karwa [9].

Table 2 summarizes the results of material compatibility evaluations of thermoplastics and elastomers exposed to a 50/50 blend of R-515B and POE32 (a lubricant commonly used in refrigerant systems for improved valve and fitting performance) after exposure at 90 °C for 21 days. Each material was evaluated for changes in hardness, volume, and mass. The key consideration was net reduction in % volume, or shrinkage, as this could indicate potential system leaks where these materials contact refrigerant / oil mixtures. All thermoplastics except phenolic were shown to be compatible with the R-515B / POE32 blend as measured by lack of shrinkage. Of the tested elastomers, hydrogenated nitrile butadiene rubber (HNBR) did not experience shrinkage and was shown to be compatible with the R-515B / POE32 blend, while the remaining elastomers were deemed incompatible. Of note, the compatibility results of thermoplastics and elastomers with R-515B were similar to mixtures containing POE32 and other ubiquitous refrigerants such as R-1234ze(E) and R-134a. Note that polymer compatibility with system refrigerant may differ from the results in Table 2 due to differences in operating conditions, lubricants and greases used, and the formulation of the polymer itself. The acceptance criterion may differ between static and dynamic seals, especially in quick disconnect couplings. It is therefore recommended to select polymers based on tests under actual operating conditions and based on component-specific acceptance criteria.

Table 2. Compatibility results of thermoplastics and polymer with R-515B, R-1234ze(E), and R-134a [10]

Polymer material	R-515B with POE32			R-1234ze(E) with POE32			R-134a with POE32		
	Mass change%	Volume change %	Hardness change %	Mass change%	Volume change %	Hardness change %	Mass change %	Volume change %	Hardness change %
Thermoplastics									
Nylon	<1	<1	4	<1	<1	7	<1	<1	3
PAI	<1	<1	0	−1	<1	4	<1	<1	1
PET	1	1	3	<1	<1	3	2	1	0
Polyimide	4	3	0	−18	−27	−14	1	−1	−2
Polyacetal	1	1	0	<1	<1	<1	2	1	−1
ABS	2	1	1	<1	<1	<1	7	5	−25
PEI	<1	<1	5	<1	<1	6	<1	<1	4

Date: February 14, 2025

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

PP	5	4	3	4	3	<1	2	2	0
PTFE	3	4	0	2	3	−7	3	4	2
PPS	<1	<1	0	<1	<1	1	<1	<1	−1
Phenolic	−2	−2	3	−2	−1	1	−2	−2	3
PEEK	<1	<1	5	<1	<1	−1	<1	<1	1
PU	1	<1	−15	<1	<1	−22	2	1	2
PVDF	3	4	0	3	3	−4	2	3	−3
Elastomers									
Neoprene	−9	−10	7	−11	−15	3	−8	−10	1
HNBR	6	7	4	11	12	3	15	16	<1
NBR	−5	−6	8	−3	−3	6	−2	−3	1
EPDM	−17	−23	10	−19	−25	22	−19	−25	6
Silicone	−2	−2	−8	<1	1	−6	1	2	0
Butyl Rubber	−16	−20	11	−6	−9	18	−14	−18	2

4.2 Low-Pressure Fluids/Refrigerants

Refrigerants with lower working pressures than those described in the previous section are also used for the thermal management of electronics in limited, specific DLC applications. These fluids may also be used in shorter cooling loops such as in-rack cooling systems. Table 3 summarizes the key properties of these fluids that include legacy HFE Novec 7000 and R1233zd, an HFO alternative with a much lower GWP and an ASHRAE A1 safety classification.

Table 3. Low-pressure fluids for pumped two-phase cooling

Refrigerant	AR4 GWP	ASHRAE safety classification	Normal boiling point, °C	Pressure at 50 °C, kPa	Volumetric capacity at 50 °C, MJ/m^3	RCL, g/m3
Novec 7000 (HFE-347mcc3)	575	NA/ non-flammable, OEL: 250 ppm)	34.2	174	170	Not Established
R-245fa	1030	B1	15.1	344	221	190
R-1233zd(E)	1	A1	18.2	293	212	85

5 Application and Installation Support

End-to-end pumped 2-phase solutions require collaboration between ITE suppliers, data center architects, and P2P solutions providers. Customers interested in deploying P2P solutions can

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

capture intended use cases and reach out to industry partners of choice. The outline below is a framework highlighting key work breakdown structure elements required for successful install of P2P infrastructure and ITE:

- Deployment
  - Site assessment
  - Design
  - Project management
- Commissioning
  - Installation
  - Startup
  - Testing
- Maintenance
  - Preventive maintenance
  - Fluid maintenance
  - Troubleshooting

5.1 Available Documents for Installation

The list below provides an example of the documentation accompanying an end-to-end P2P DLC installation. While some are common to any cooling installation, others are unique - these are designated by (*) in the list. A careful review of this list helps the reader define the requirements, design considerations, and trade-offs associated with P2P LC systems.

- Submittals for CDU, SFN, and ITE racks
- CDU host management interface (HMI) user guide
- System installation and start-up guide
- System installation, floor plan, and overhead clearance drawings
- Rack manifold interface control drawing (ICD) with 2P QD types designated (*)
- ITE interface control drawing (ICD) defining refrigerant connections (*)
- P2P system commissioning plan and checklist (*)
- Preparing servers for P2P operation (i.e. leak check, evacuate, and charge SOP) (*)
- 2P server install and removal SOP (*)
- Replaceable core filter drier SOP (*)

Date: February 14, 2025

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

- System vacuum and charge SOP (*)
- System and ITE charge recovery SOP (*)

6 Conclusions

Due to explosive growth in AI, a major focus on investments in capable ITE, software, and data center infrastructure – cooling, power delivery, and networking – is in full swing. As chips, racks, and cluster power densities increase, liquid cooling deployments gain increasing importance and adoption. These solutions must not only keep pace with growing power levels but must also help address customer sustainability targets and efficiency goals. Pumped 2-phase DLC provides a viable, safe, and effective path toward meeting these needs through proven performance at scale, a growing component ecosystem based largely on the long-established, commoditized HVAC industry, and a number of safe, commercially available, inert working fluids. In addition, P2P DLC offers decreased system downtime by minimizing the risk of equipment damage due to working fluid leaks along with mitigating risks and costs associated with corrosion and fluid contamination. Further, this technology allows use of warmer facility water, thereby reducing chiller usage and thus data center TCO, while providing a path to more efficient waste heat capture for downstream usage. Further exploration and adoption of this technology is expected to help address the many challenges that accompany the numerous benefits of AI advancement.

7 Glossary

2P	Two Phases (usually liquid and vapor/gas)
AI	Artificial Intelligence
CDU	Coolant Distribution Unit
DLC	Direct Liquid Cooling
GWP	Global Warming Potential

Date: February 14, 2025

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

HVAC	Heating, Ventilation and Air Conditioning
P2P	Pumped Two Phase Liquid Cooling
PUE	Power Usage Effectiveness
QTQD	Quick Thread Quick Disconnect
SFN	Secondary Fluid Network
TCO	Total Cost of Ownership
TDP	Thermal Design Power
WUE	Water Usage Effectiveness
ML	Machine Learning
KPI	Key Performance Indicator
HBM	High Bandwidth Memory
ASIC	Application Specific Integrated Circuit
DRAM	Dynamic Random Access Memory
DC	Datacenter
ITE	Information Technology Equipment
HPC	High Performance Computing

Date: February 14, 2025

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

ODP	Ozone Depletion Potential
HFC	Hydrofluorocarbon
HFE	Hydrofluoroether
HFO	Hydrofluoroolefin
RCL	Refrigerant Concentration Limit
OEL	Occupational Exposure Limit
QD	Quick Disconnect

8 References

[1] Aranya Chauhan, and Satish G. Kandlikar. "Characterization of a dual taper thermosiphon loop for CPU cooling in data centers." Applied Thermal Engineering 146 (2019): 450-458.

[2] Morteza Alipanah, and Saeed Moghaddam. "Ultra-low pressure drop membrane-based heat sink with 1000 W/cm2 cooling capacity and 100% exit vapor quality." International Journal of Heat and Mass Transfer 161 (2020): 120312.

[3] Mohammad Reza Shaeri, et al. "Thin hybrid capillary two-phase cooling system." International Communications in Heat and Mass Transfer 112 (2020): 104490.

[4] Mohammad Reza Shaeri et al. “Additively Manufactured Hybrid Two-Phase Cold Plate.” International Electronic Packaging Technical Conference and Exhibition. Vol. 87516. American Society of Mechanical Engineers, 2023.

Date: February 14, 2025

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

[5] Serdar Ozguc, et al. "An approach for topology optimization of heat sinks for two-phase flow boiling: Part 2–Model calibration and experimental validation.” Applied Thermal Engineering 249 (2024): 123338.

[6] Ali Heydari, et al. "Direct-to-Chip Two-Phase Cooling for High Heat Flux Processors." International Electronic Packaging Technical Conference and Exhibition. Vol. 86557. American Society of Mechanical Engineers, 2022.

[7] Qingyang Wang, et al. “A Server-Level Test System for Direct-To-Chip Two-Phase Cooling of Data Centers Using a Low Global Warming Potential Fluid.” 2024 23rd IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm). IEEE, 2024.

[8] Akshith Narayanan, et al. “Investigation of Server Level Direct-to-Chip Two Phase Cooling Solution for High Power GPU and CPUs.” International Electronic Packaging Technical Conference and Exhibition. (Accepted) American Society of Mechanical Engineers, 2024

[9] Nitin Karwa, “Ultra-Low Global Warming Potential Heat Transfer Fluids for Pumped Two-Phase Cooling in HPC Data Centers.” in 2019 19th IEEE ITherm, Orlando, FL, USA, (2020): 283-290.

[10] Henna Tangri, et al., “Solubility, miscibility and compatibility studies of low GWP non-flammable refrigerants and lubricants for refrigeration and air conditioning applications,” International Journal of Refrigeration, 148(2023): 45-63.

[11] Kulkarni, Devdatta, Gulick, J., King, J., Jarrett B., George, P., Feldman, Y., “Thermal Performance of Common Cold Plate for Pumped Single- and Two-Phase Direct Liquid Cooling for Next generation High Power Server Processors”, IPACK2024-141367, V001T02A013-10 pages.

Date: February 14, 2025

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

9 License

9.1 Creative Commons

OCP encourages participants to share their proposals, specifications and designs with the community. This is to promote openness and encourage continuous and open feedback. It is important to remember that by providing feedback for any such documents, whether in written or verbal form, that the contributor or the contributor's organization grants OCP and its members irrevocable right to use this feedback for any purpose without any further obligation.

It is acknowledged that any such documentation and any ancillary materials that are provided to OCP in connection with this document, including without limitation any white papers, articles, photographs, studies, diagrams, contact information (together, “Materials”) are made available under the Creative Commons Attribution-ShareAlike 4.0 International License found here: https://creativecommons.org/licenses/by-sa/4.0/, or any later version, and without limiting the foregoing, OCP may make the Materials available under such terms.

As a contributor to this document, all members represent that they have the authority to grant the rights and licenses herein. They further represent and warrant that the Materials do not and will not violate the copyrights or misappropriate the trade secret rights of any third party, including without limitation rights in intellectual property. The contributor(s) also represent that, to the extent the Materials include materials protected by copyright or trade secret rights that are owned or created by any third-party, they have obtained permission for its use consistent with the foregoing. They will provide OCP evidence of such permission upon OCP’s request. This document and any "Materials" are published on the respective project's wiki page and are open to the public in accordance with OCP's Bylaws and IP Policy. This can be found at http://www.opencompute.org/participate/legal-documents/. If you have any questions, please contact OCP.

Date: February 14, 2025

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

10 About Open Compute Foundation

The Open Compute Project (OCP) is a collaborative Community of hyperscale data center operators, telecom, colocation providers and enterprise IT users, working with the product and solution vendor ecosystem to develop open innovations deployable from the cloud to the edge. The OCP Foundation is responsible for fostering and serving the OCP Community to meet the market and shape the future, taking hyperscale-led innovations to everyone. Meeting the market is accomplished through addressing challenging market obstacles with open specifications, designs and emerging market programs that showcase OCP-recognized IT equipment and data center facility best practices. Shaping the future includes investing in strategic initiatives and programs that prepare the IT ecosystem for major technology changes, such as AI & ML, optics, advanced cooling techniques, composable memory and silicon. OCP Community-developed open innovations strive to benefit all, optimized through the lens of impact, efficiency, scale and sustainability. Learn more at www.opencompute.org.

Date: February 14, 2025

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.