TETRA TECH
High Performance Buildings Group

Taking the Heat Out of AI

Sustainable Solutions for Liquid-Cooled AI Data Centers

Green Paper
April 2024
Revision 1

Leading with Science®

Abstract

AI is being widely utilized across many industries and high-powered servers are becoming commonplace in data centers. The next generation of AI servers pushes the bounds of computational power at the cost of increasing power consumption, requiring the use of liquid cooling. Liquid-cooled servers will need to work alongside air-cooled IT equipment, leading to a hybrid environment.

Direct-to-chip and immersion cooling provide great opportunities for increased heat rejection efficiencies and better parameters for heat re-use.

Liquid cooling of AI servers does not require a fundamental change to facility water systems (FWS), but the cooling systems will need to evolve to support both liquid- and air-cooled requirements that will exist in a hybrid environment.

Today, new liquid-cooled data centers are being designed and constructed globally to support the new 2.7 kW NVIDIA Blackwell GB200 GPU and the new 97.2 kW NVL72 liquid-cooled servers.

The increased power demand of AI should also drive the development of renewable and low-carbon technologies to meet decarbonization targets. Data center developments will need to focus on increased usage efficiencies and delivering effective heat reuse.

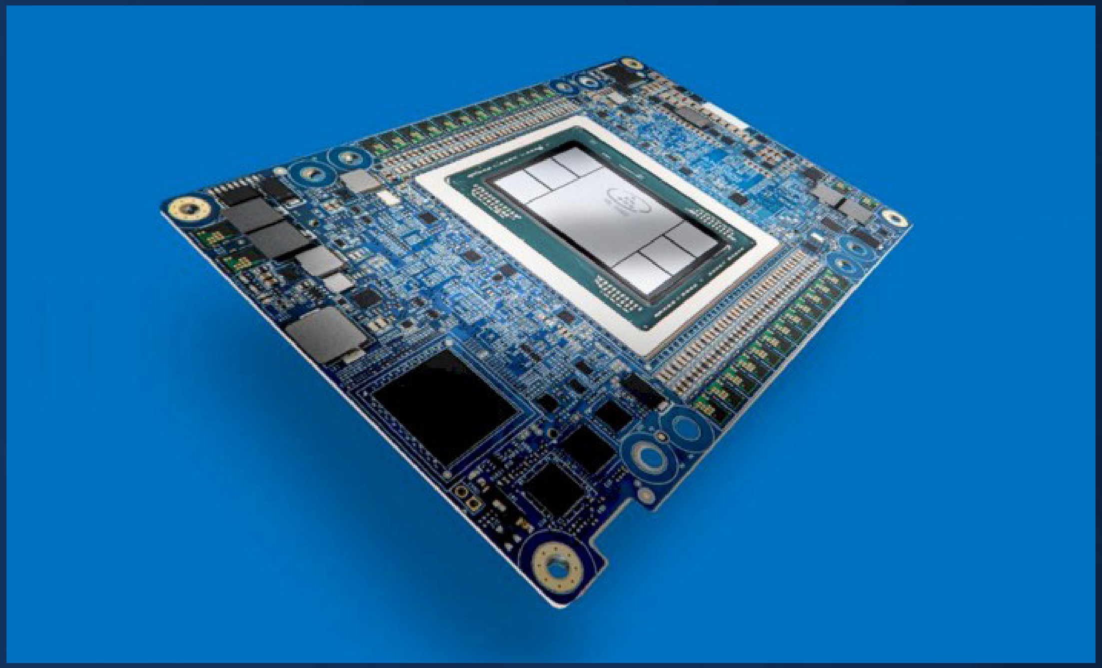

Intel releases Gaudi2 AI GPU with 600W TDP

Introduction
Liquid Cooling is Now a Requirement for AI Data Centers

As the demand for artificial intelligence (AI) and machine learning (ML) applications continues to grow, the need for powerful and efficient computing infrastructure becomes paramount. AI servers generate much more heat than their predecessors, making effective cooling essential to maintain optimal performance, reliability, and longevity of operation. Liquid cooling solutions are now available to deal with these new AI server environments, which will impact the facility cooling infrastructure.

This paper examines the options for server liquid cooling, identifies the impact of the facility water system (FWS) and explores how sustainability can be embraced in the next evolution of data centers.

The Issue

Data centers are energy-intensive facilities, contributing significantly to greenhouse gas (GHG) emissions. Studies have estimated that data centers consume up to 1.5% of global power in 2010. The percentage of global electrical consumption has kept a steady rate below 2% in the last decade. However, this is because servers had been quite efficient and the demand for computationally intensive servers had not grown much. The last few years have already shown a shift in demand as more industries and services now require more available processing power than ever before. As a growing share of the mission-critical market begins to be dominated by hyperscale facilities, more power will be required, and the necessity for efficient IT and cooling systems will be necessary. [29]

Nvidia recently announced the launch of their new Blackwell GPUs in March 2024. The B100 has a TDP of 700W like its predecessor, the Hopper series H100 and H200. However, the B200 GPUs have a projected TDP of 1000W. These GPUs will be offered in a server packaged with the Grace series CPU, the GB200, with a total TDP of 2.7kW. Nvidia will also be offering a full server rack with the GB200 NVL36 and GB200 NVL72, the latter projected to have a max TDP of 120kW. These racks only come as liquid-cooled platforms, as traditional air-cooled racks can no longer transfer the heat at such dense loads. For comparison, the Ampere GPU was launched in 2020 and had a TDP of 400W. In just four years, processing power has increased exponentially as the TDP has more than doubled. This indication concludes that as performance increases and efficiency is optimized, the power consumption and thermal load are still increasing. [26][27][28]

While liquid cooling will improve efficiency, the substantial increase in power will increases the carbon footprint.

The Need for Renewable Energy

Offsetting GHG emissions is crucial to mitigating the environmental impact of liquid-cooled AI data centers.

Disclaimer

The information provided in this paper is an overview of this new technology.
The user is responsible for designing systems that meet the server manufacturer’s specified requirements.

What is an AI Server and What Does it Do?

A Billion Calculations in One Second

The key operating features of an AI server are:

- Intensive computational workloads
- Parallel processing
- Data movement and storage
- Dense configurations
- Real-time processing
- Model complexity

The high-performance levels of AI servers require specialized hardware, including Graphics Processing Units (GPUs) and Tensor Processing Units /Accelerator applications (TPUs).

- GPUs are essential components in AI servers, especially for deep learning tasks. They excel at parallel processing, which is crucial for training and running neural networks.
- Some companies, such as Google, custom design TPUs. They are highly efficient for tensor-based operations commonly found in neural networks.

AI servers may incorporate high-performance Central Processing Units (CPUs) to handle general-purpose computing tasks. However, GPUs and TPUs are often the primary workhorses for AI workloads. They also incorporate high-capacity random access memory (RAM) used especially during training. Having both large and quick storage is a key to the training or inference. It allows one to access and retrieve large data datasets at record speeds.

AI servers also incorporate Parallel Processing Architecture, designed to exploit parallelism, a key characteristic of many AI algorithms. Parallel processing allows for simultaneous execution of multiple computations, enhancing the speed and efficiency of AI tasks.

GB200 NVL72

36 GRACE CPUs
72 Blackwell GPUs
Fully connected NVLink switch rack

Training: 720 PFLOPS
Inference: 1,440 PFLOPS
NVL Model Size: 27T params
Multi-Node Bandwidth: 130TB/s
Multi-Node All-Reduce: 260TB/s

> **Nvidia announced the details of its new Blackwell GB 200 GPU chip and the NVL 72 liquid-cooled server. The GB 200 NVL72 consumes 97.2kW. Liquid cooling is now required for the new generation of GPUs.**

Keeping Them Cool

Thermal Design Power (TDP) is the maximum amount of heat a processor generates. The maximum TDP of a system determines the heat load. Currently, these new high-performance GPUs can generate 700 watts of heat per chip, which is seven times more than conventional CPUs. Thermal design power (TDP) for a traditional high-density rack can be defined as having a TDP of up to 30 kW. Scientific computational servers used in research facilities already far exceed the thermal output of enterprise servers. Based on these numbers, drastic action is required to keep these servers cooled within their design parameters. Lower-powered versions of these servers are being arranged to be air cooled but this is not a long-term solution, especially as the TDP[1] continues to increase. At the current rate of new commercial and scientific processor development, we already see rack TDPs exceeding 50 kW. With the exponential growth in the machine learning industry, we could redefine high-density racks to be upwards of 200 kW TDP or beyond.

Liquid cooling is not a new technology in the telecommunications space. Liquid was used to cool the early mainframe computers. The heat needed to be removed from high-density CPUs, while GPUs require a different treatment scale.

The latest greenfield data centers (and brownfield data centers being retrofitted) are now being deployed using hybrid cooling systems; low-density racks tend to keep using the standard air-cooled solution, while high-density racks are moving to liquid cooling technologies.

> **New data centers are being constructed to serve a hybrid solution, liquid-cooled and air-cooled.**

The newer forms of liquid cooling are now being deployed on selected applications and do not use traditional cooling methodology.

The two main types of liquid-cooled server systems are:

- Direct to chip / cold plate
- Immersion

Direct-to-Chip

Direct-to-chip, also known as cold plate, servers have metal – often copper – plates attached to the heat-generating electronic components. As the coolant, usually water or a specialized dielectric fluid with good thermal properties, circulates through channels or passages within the cold plate, the coolant absorbs the heat generated by the electronic components and carries it away from the server via a water-based cooling system.

Figure 1. Direct to chip cooling – ZutaCore waterless, two-phase direct-to-chip liquid cooling

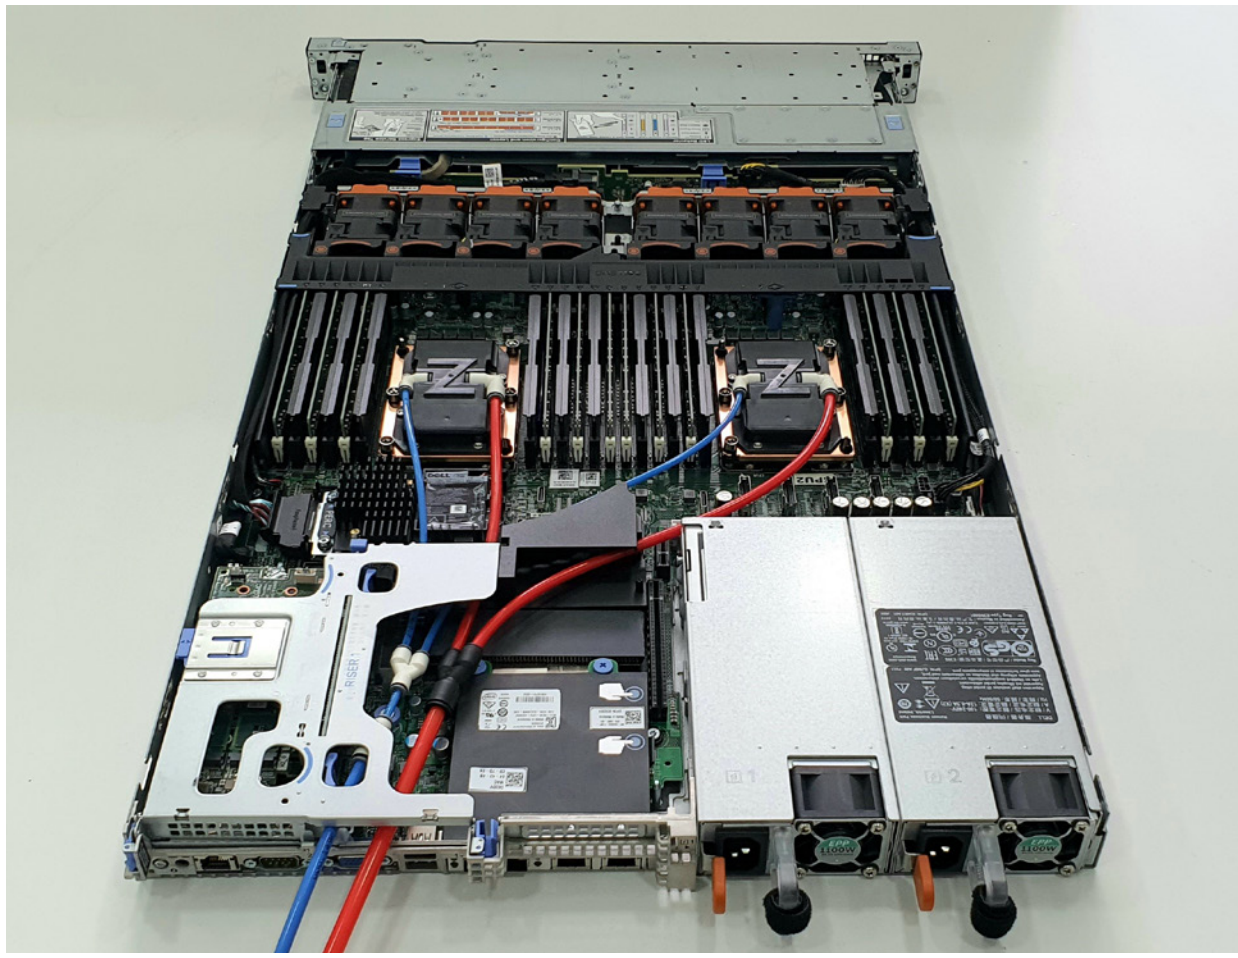

Direct-to-chip cooling can capture up to 80% of the total heat production from the servers, with the remaining 20% needing to be air-cooled.

Immersion cooling comes in two primary forms

Single-Phase Immersion Cooling is a straightforward design where servers are fully immersed in a dielectric fluid. The dielectric fluid absorbs heat from all the components while remaining liquid throughout the cooling process. The heated fluid from the immersion tank is then transferred to the facility water system via a heat exchanger. The main components in a single-phase immersion cooling system include the immersion tank (where the servers and dielectric fluid are located), the Cooling Distribution Unit (CDU), where the heat is transferred to a water circuit, and the water circuit, composed of the pump set, piping, and cooler. This cooling method offers several benefits, including capturing 100% of the heat generated by the server, cost-effectiveness, improved efficiency leading to a Power Usage Effectiveness (PUE) reduction, and the enablement of high-density rack designs with the possibility of heat recovery.

Two-Phase Immersion Cooling involves the transition of the dielectric liquid from liquid to vapor as it absorbs heat. The vapor rises to a condenser unit, where it

---

1. TDP is the maximum amount of heat generated by a computer chip or component that the cooling system in a computer is designed to dissipate under any workload)

Keeping Them Cool

releases heat, condenses back into liquid, and then flows back to the immersion tank. A water-based cooling system removes the heat rejected by the condenser. This can give an enhanced heat dissipation due to phase change and efficient cooling for high-power applications. Note: The refrigerant used in this technology has been determined to be toxic under US guidelines and is not being utilized. This technology is expected to advance as alternate refrigerants mitigate health risk issues. [31]

Figure 2. Immersion cooling by Submer

A benefit of immersion cooling is that it captures 100% of the server heat with minimal air cooling required for the space. The current expectation is that immersion cooling may not be adopted as widely as direct-to-chip.

However, immersion cooling imposes higher structural loads on the building structure, which will need to be considered for multi-story buildings.

An overhead support structure is required for the lifts to remove servers from the immersion tank.

Hybrid cooling
Some hybrid versions of liquid cooling are being deployed in the form of closed coupled air-to-water heat exchangers and rear door coolers, however, these still use air as the heat rejection medium and do not harness the full potential of liquid cooling. These hybrid solutions may well be part of the overall data center cooling solution but AI servers must be fully liquid-cooled.

Engineering the Cooling
Parameters

The liquid-cooled technologies described above rely on a water-based cooling system called the Facility Water System (FWS). The FWS is the chilled or cooling water system operated and controlled by the data center owner, which includes the core of the mechanical system in a data center, including piping, pump set, valves, sensors, and coolers. The type and size of coolers (dry cooler, adiabatic cooler, chiller), pump set and, in general, all the mechanical infrastructure of the Technology Cooling System (TCS) is the secondary (or tertiary) system containing the water that enters the servers or the immersion tank. The TCS is usually separated from the FWS by a cooling distribution unit (CDU) – a heat exchanger and pump unit.

As with current data center design, the FWS can take many forms dependent on several factors listed below. For FWS serving liquid-cooled servers, the TCS requirements become a key consideration.

- Geographic location
- Environmental conditions
- Water temperatures
- System size
- System topology
- System resilience and thermal ride through

Facility Water System (FWS) Temperature

Liquid-cooled servers can use higher water temperatures than air-cooled systems. This can be in the order of 40°C (104°F) vs 20°C (68°F) – liquid- vs air-cooled. This significantly improves heat rejection efficiencies with higher water temperatures allowing mechanical cooling to be reduced or even eliminated in some climatic zones. Higher water temperatures also provide better heat reuse potential.

Based on current deployments, liquid cooling temperatures or TCS are around 30-35°C (86-95°F) on-server temperature. This appears to be due to limited long term test data mainly focused on the following.

- GPU case temperatures must be reduced as the GPUs become more powerful.
- Server efficiency with thermal throttling occurring at peak use.

Depending on the GPU specification, the liquid cooling in the immersion tank can be raised to 50-55°C (112-131°F) in immersion cooling applications. This significantly extends the range of climatic zones where mechanical cooling can be reduced and/or eliminated. Adoption of this will be subject to the server manufacturer’s testing.

Engineering the Cooling

Both methods of server liquid cooling (Direct to Chip and Single-Phase Immersion Cooling) employ a cooling distribution unit (CDU) to isolate the server cooling from the primary water system hydraulically. This separation via a water-to-water, or dielectric fluid-to-water, heat exchanger adds inefficiency to the system and elevates the on-server cooling water temperature by about 5°C (41°F). This results in a FWS flow temperature of 25-30°C (77-86°F), not much higher than a modern air-cooled server chilled water system.

It is most efficient to drive these temperatures as high as possible. However, server manufacturers have advised that server performance decreases with higher cooling fluid temperatures, so the cooling engineers may have little or no say in this. There may well be options for seasonal peak increases in FWS flow temperature to avoid the need for full mechanical cooling, but this will be user/site-specific.

Based on the above for water-cooled applications, the FWS should be designed for a 25°C (77°F) flow temperature, but with the flexibility to increase this, as server technology evolves to accept higher temperatures without impacting performance.

Additional limitations and considerations for FWS temperatures must account for future IT loads. With higher TDPs, even water temperatures serving the TCS may need to be lower, 32°C (89.6°F) or even down to 27°C (80.6°F). Performance comes with the cost of less efficient systems. [32]

Another restriction on the FWS is that many facilities will be hybrid with both air-cooled and liquid-cooled servers. This will lead to decisions regarding the use of common or separate systems. This has been considered more below.

> **Driving FWS return temperatures as high as 60°C (140°F) provides options for heat reclaim and using dry coolers instead of adiabatic cooling or refrigeration.**

Heat Rejection

The elevated water temperatures required by both direct-to-chip and immersion liquid cooling may allow most of the heat to be rejected by dry coolers in temperate climates. In geographies with higher ambient temperatures, mechanical or adiabatic /evaporative cooling may still be required, but this will be reduced from the current chilled water load for air-cooled environments. The deployment of hybrid IT cooling (air and water) will, however, need to be considered and this may well dictate the heat rejection requirements.

There are four main types of heat rejection that can be adopted for large scale data centers.

1. Air cooled chiller with free-cooling coils
2. Dry air coolers with water-cooled chillers
3. Cooling towers / hybrid coolers with water-cooled chillers
4. Cooling towers / hybrid coolers with water-cooled chiller assist

See below for an assessment of where different types of heat rejection can be applied.

Table 1. Water quality guidelines for the FWS and TCS

Climate zone	ASHRAE 2021, n=20 Max temperatures DB/WB (°C)	Heat rejection for 25°C FWS flow temperature	Heat rejection for 35°C FWS flow temperature
TROPICAL (EQUATORIAL)	37.9 / 30.8	ACC	ACC
DRY	46.6 / 26.9	ACC	DAC or WCCA
TEMPERATE	41.9 / 27.7	ACC or WCCA	HC/CT
POLAR	10.4 / 8.8	DAC	DAC

Key for above:
- ACC – air cooled chiller with water at 20°C (68°F) mixed to provide 25°C (77°F) floor to the CDU
- DAC + WWC – dry air cooler with water cooled chiller (where water stress applies)
- HC/CT – dry cooler, hybrid cooler cooling tower (closed circuit)
- WCCA – water cooled chilled assist

Heat rejection options are based on the following:
- 5°C (9°F) approach at CDU heat exchanger
- 5°C (9°F) approach for cooling tower or hybrid cooler (wet bulb to water outlet temp)
- Location not under water stress

Engineering the Cooling

System Size

Data centers can have a vast range of cooling requirements from 1 MW for a communications switch site to more than 500 MW for a hyperscale campus. A water-based cooling system needs to be sized and designed to suit the readily available (and commercially viable) parts focused on the heat rejection and distribution system.
A current typical FWS for a large data center would be between 10-15 MW, with multiple systems used as required. Systems may be interlinked for phased deployment or redundancy, but the overall size will usually be capped by the viability of both the manufacture and installation of the distribution system. A 15 MW FWS would comprise 8-10 heat rejection units, utilize 600 mm diameter header pipework with 250 mm diameter concurrently maintainable ring mains.

System sizes may increase with the advent of AI data centers, but the products and skills needed to deliver large scale FWS may come at a premium.

System Topology, Redundancy, and Thermal Ride Through

The topology of the FWS will depend on the scale of the installation and the required redundancy. Most data centers are being constructed to be concurrently maintainable, which provides a good base framework. The framework can then be developed to suit any modular / phasing requirements. Additional component or system redundancy can be added where client specifications require it.

A key aspect of liquid cooling that differs from air cooling is the direct impact on the water-cooled server resulting from a deviation from the set point. In an air-cooled data hall, the flooded room and the thermal mass of the hall dampen fluctuation in the supply air temperatures, usually resulting in a limited impact on the server air inlet temperatures. If the temperature deviates by 2-3°C (4-5°F), which is normally within the service level agreement (SLA) range, the server fans ramp up to reject more heat. In a water-cooled application, the change and rate of change to the water temperature or flow rate can dramatically impact the server due to the much larger quantity of heat being rejected. For liquid-cooled servers, thermal stability and thermal ride through need to be considered very carefully. Providing stable cooling water flow at a suitable temperature is essential for the continued correct operation of the servers.

In the event of a mains power failure (or a transfer to emergency power generation for maintenance purposes), the central heat rejection plant will normally shut down due to the high-power demand. Other cooling system equipment with smaller demands, i.e., pumps, CRAHs, and CDUs, would be fed by an uninterrupted power supply (UPS) and continue to operate. Thermal storage built into the FWS must provide the ride-through until the heat rejection plant is back online. The time for power to be reinstated to the plant will vary based on the electrical system topology but will typically be 2-3 minutes for medium voltage (MV) generation and around 30 seconds for low voltage (LV) generation. Once power is restored, the heat rejection plant can take 3-4 minutes to get to full capacity and this is the period that will need to be addressed. (Designers should consider the additional UPS load for operating CDU Pumps and controllers.)

Thermal storage is often via water stored in pressurized, stratified tanks rather than with vented tanks, which are more common for use in demand shifting applications. Systems are configured so that the storage tanks are continuously charged when not in use and will discharge with little or no need to reposition valves to redirect flow. Circuiting and controls to direct flow from the tank to the return side of the system to be cooled are necessary after the tank has been utilized to prevent the warmer water that has accumulated in the tank from being sent to the load. The volume of stored water required for emergency thermal storage duty depends on the heat rejection load, minimum average temperature difference during discharge, and the discharge duration. The thermal storage tank requires careful analysis. The tank must be properly sized to remove excess heat for potentially extended periods of time. The system must be piped to operate in a failproof mode.

Engineering the Cooling

Direct-to-Chip Cooling

The FWS removes heat from the CDU, which includes a heat exchanger and pumping system. Treated fluids are often PG25 (25% glycol and 75% treated water). [30] Redundant piping and pumps are not shown for clarity.

Figure 3. Typical direct-to-chip cooling schematic

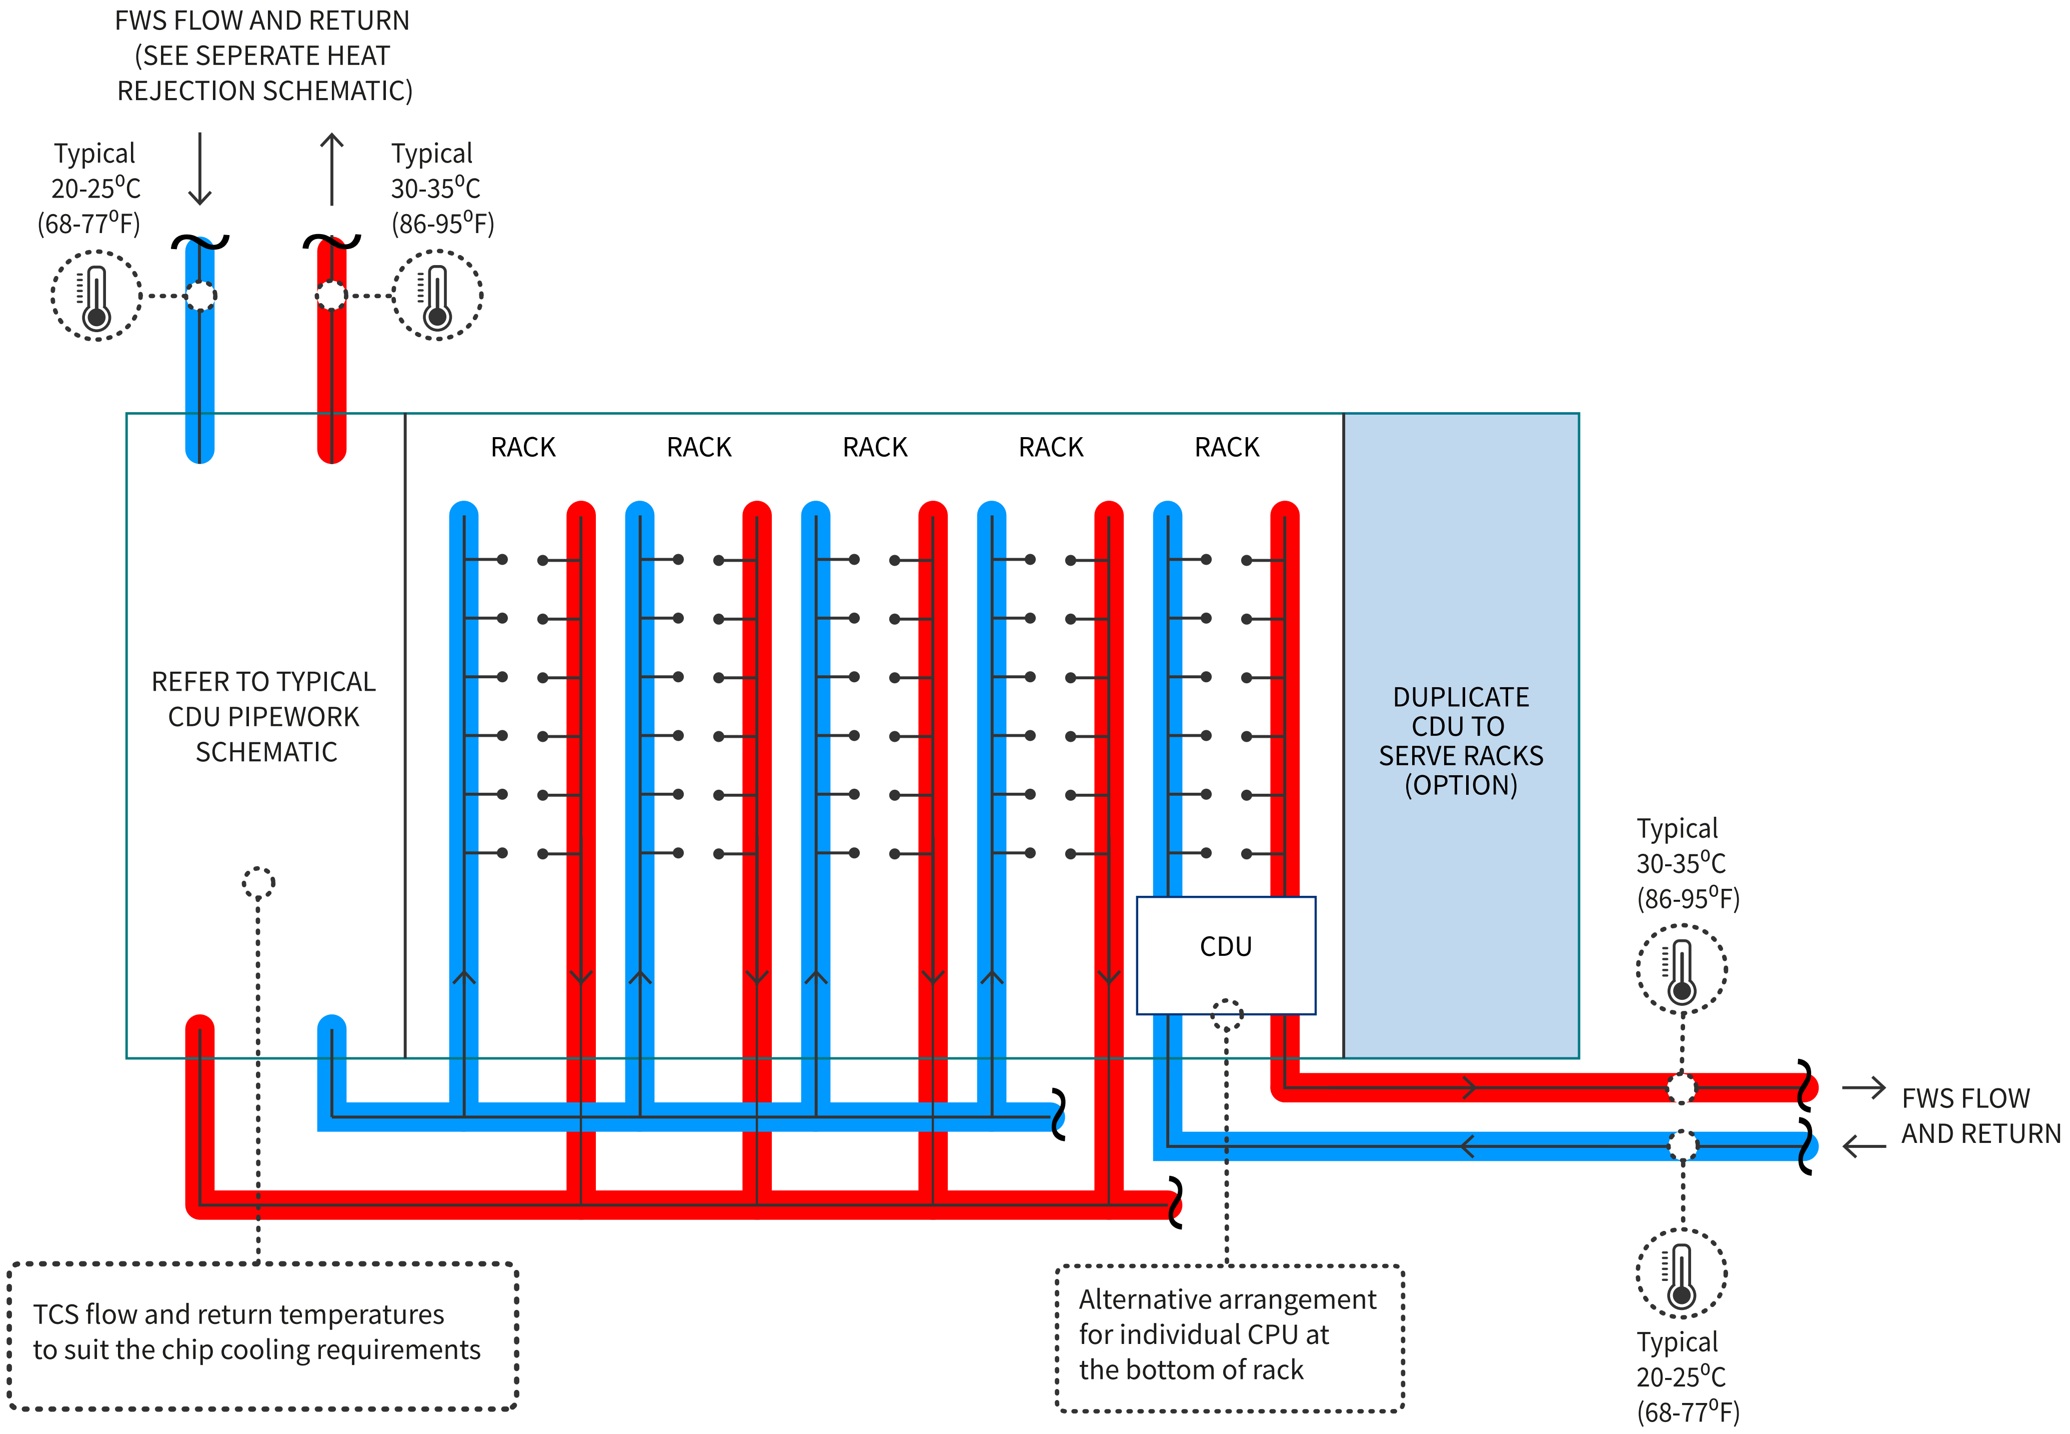

Engineering the Cooling

Immersion Cooling

The FWS removes heat from the Immersion tank heat exchanger. The dielectric fluid in the immersion tank cools the servers and chips directly. Redundant piping and pumps are not shown for clarity.

Figure 4. Immersion cooling schematic

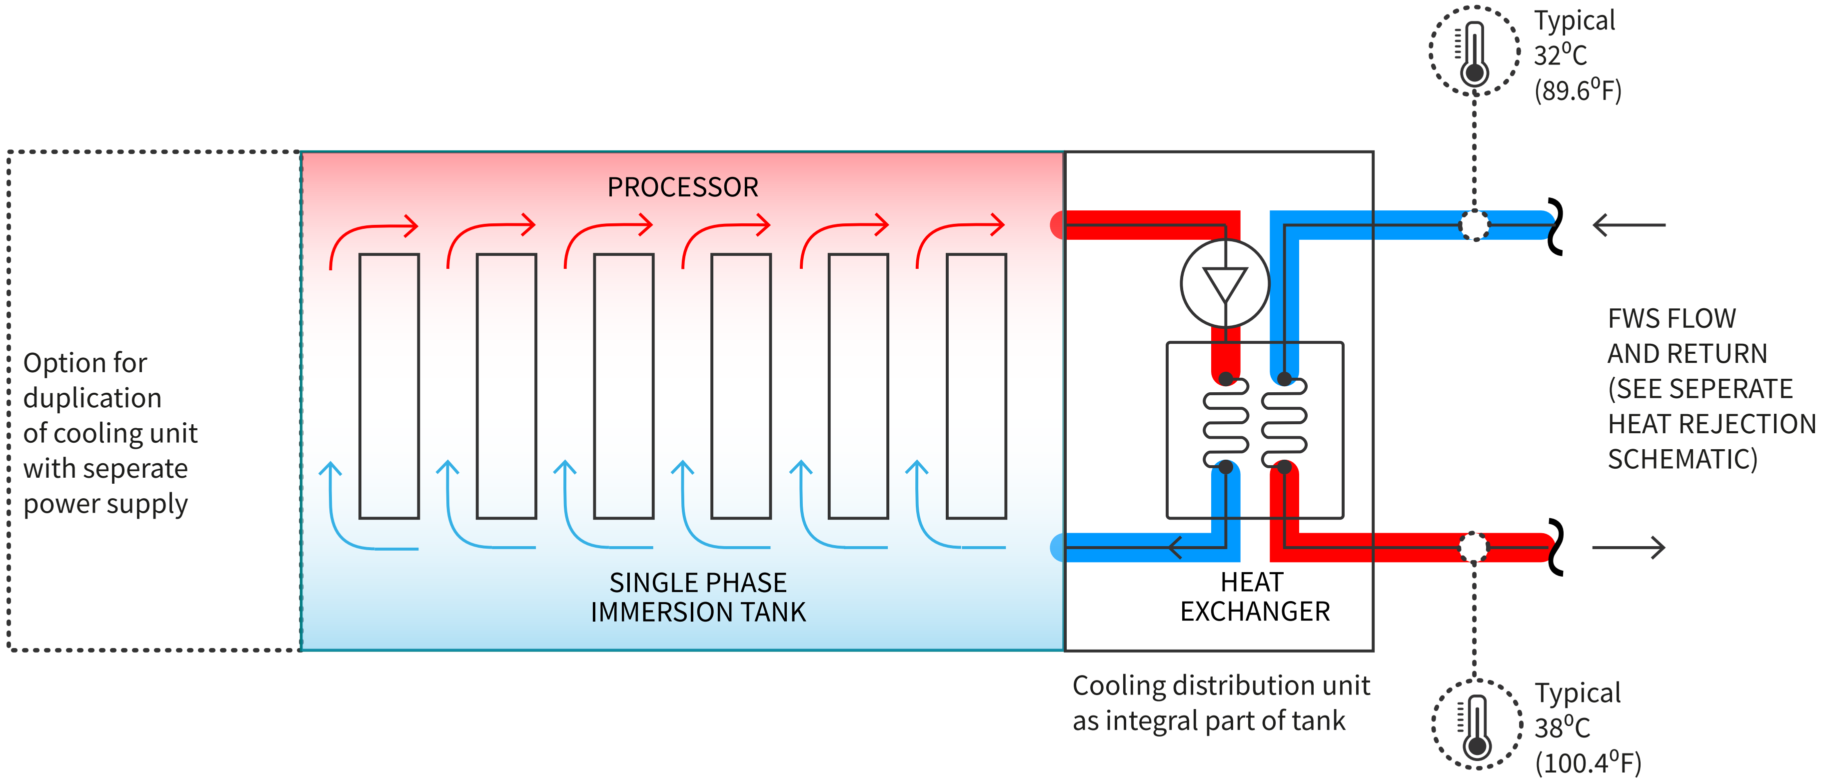

Engineering the Cooling

Typical CDU

The FWS removes heat from the CDU, which includes redundant heat exchangers and pumping systems. Treated fluids are often PG25 (25% glycol and 75% treated water). Redundant piping and valves are not shown for clarity.

Figure 5. Typical CDU schematic pipework

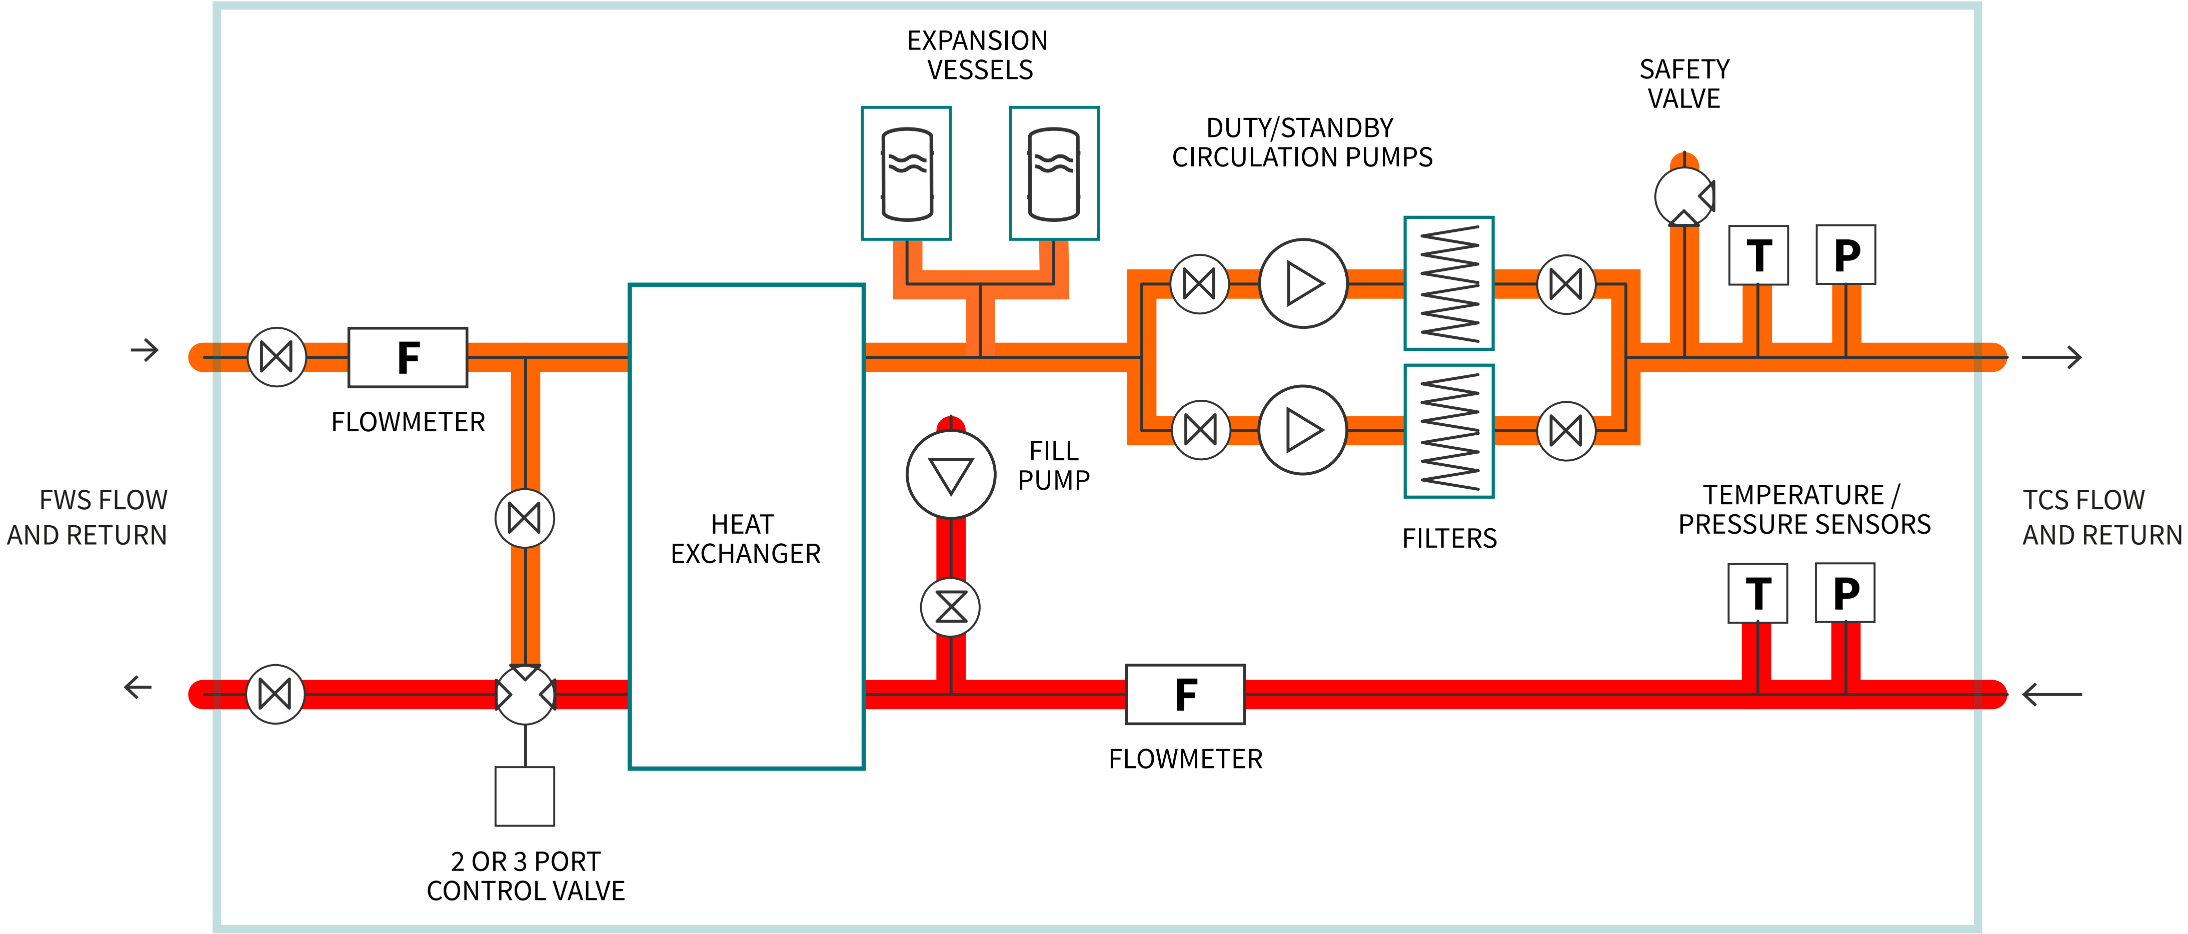

Engineering the Cooling

Typical Heat Rejection

A typical heat rejection system would utilize air-cooled chillers or dry/wetted coolers, depending on the FWS flow temperature requirement. Some applications may have heat rejection plants in series or use injection circuits (as shown later – chiller assist). The heat rejection element of the FWS does not significantly differ from a current data center, but the inclusion on heat exchangers for heat re-use will become much more widespread.

Figure 6. Typical air-cooled chiller heat rejection schematic

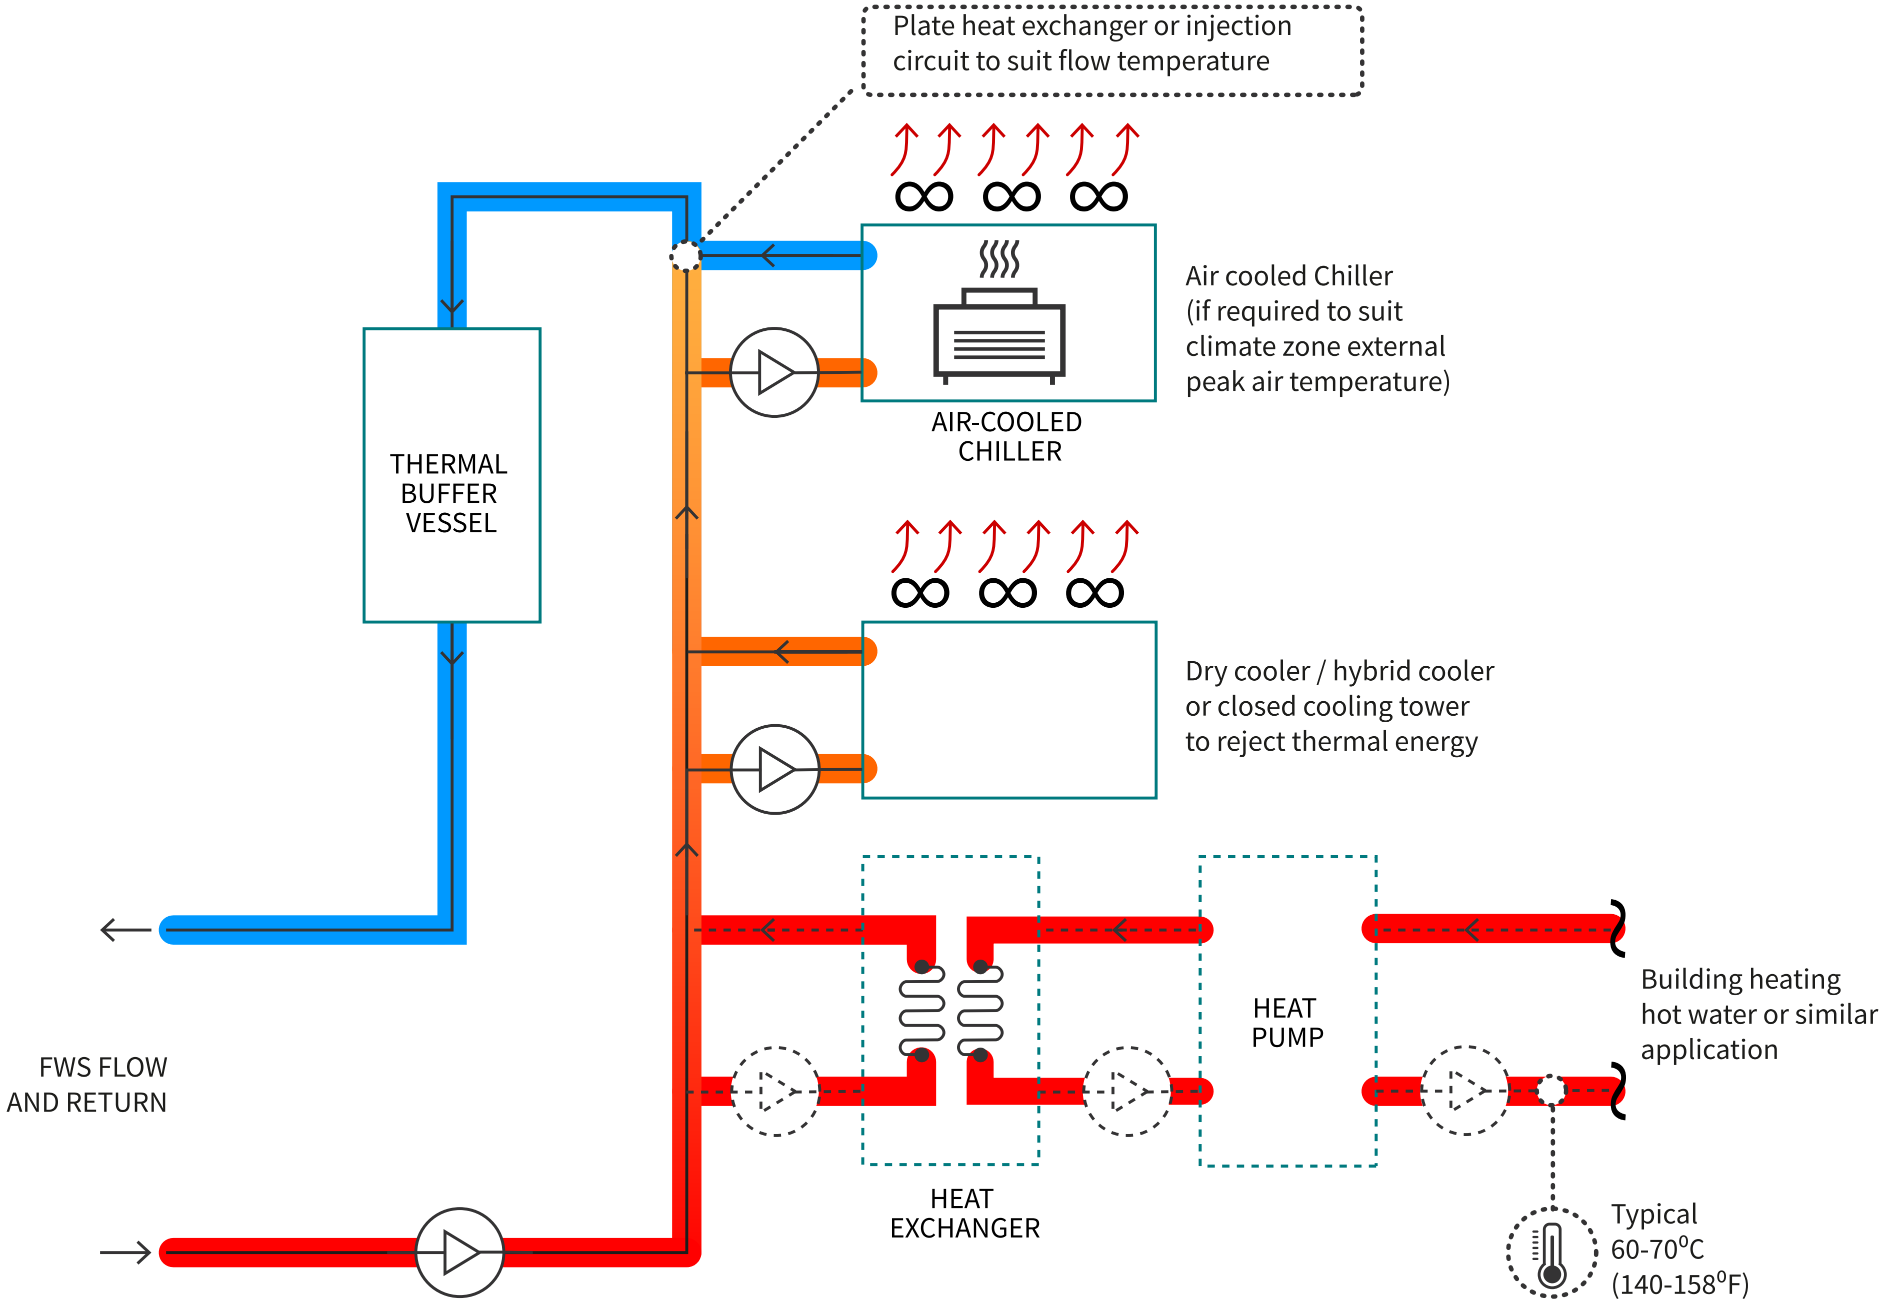

Engineering the Cooling

Hybrid Solutions

The first deployment of AI data centers is a hybrid solution of air- and water-cooled server technologies. Air is still the primary source for most server cooling, and it is likely that non-AI servers will remain this way for many years to come. This will result in a hybrid data center where both air and liquid cooling will be required.

CRAH/FWU within existing data centers will have been installed to suit peak cooling loads of 10-15kW per rack where the upper threshold is around 50kW per rack. With the addition of AI servers, this will increase well above the current design loads. As a result of this fewer racks can be deployed and these will need to be suitably spaced out to avoid hot spots

In this case when AI servers are introduced gradually, some CRAH/FWU could be replaced with CDU to cool the AI servers directly. However, the overall cooling capacity of a data hall/data center will be limited to installed infrastructure, which will correlate to the site's power availability. Without significant changes to the infrastructure (both power and cooling), deploying more AI servers will leave vacant space within the data halls.

Where new data centers are being designed and constructed to accommodate AI servers, the use of return door coolers (RDC) should be considered instead of CRAH/FWU. The RDC would be attached to the back of the racks and incorporate a heat exchanger connected to the FWS and fans (in its active version). This provides more flexibility as rack loads increase over time, above what can be supported by CRAH/FWU, and uses less fan energy than the latter. When direct-to-chip processors are introduced, the RDC can work in conjunction with the CDUs. Depending on the liquid-cooled server requirements, the FWS could be connected in series between the CRAHs and CDU to enhance the efficiency of the heat rejection system.

Hybrid Direct-to-Chip and Rear Door Cooling

Figure 7. Direct-to-chip CDU cooling schematic

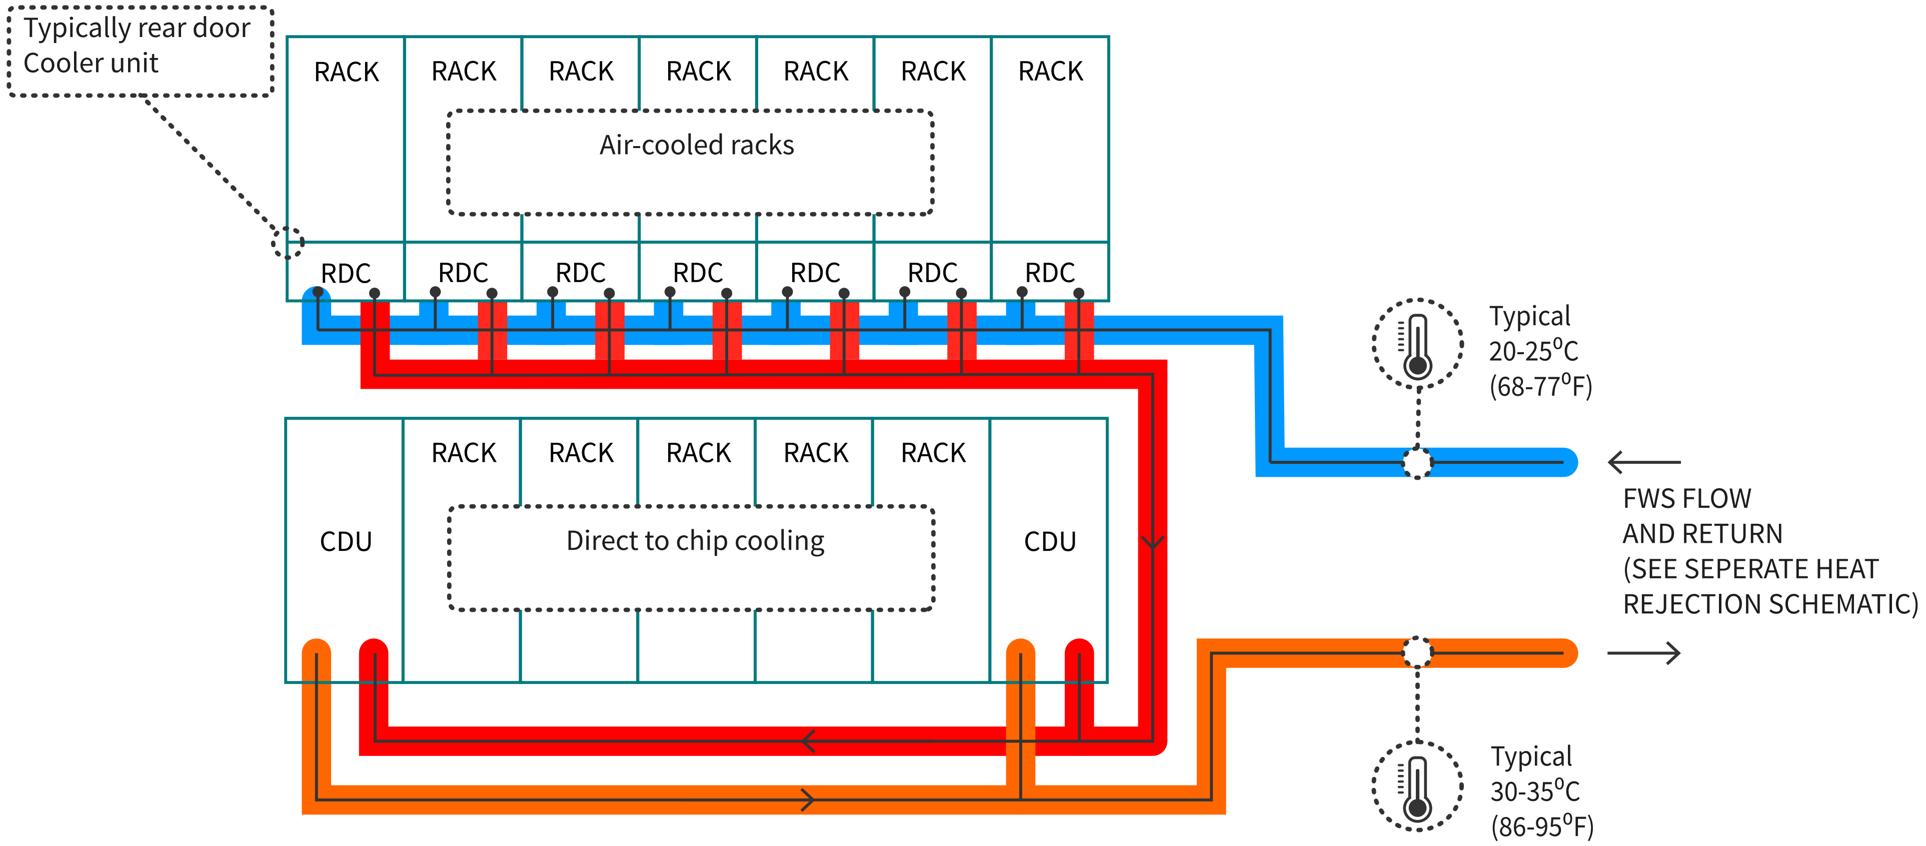

Figure 8. Typical hybrid air-cooled RDC

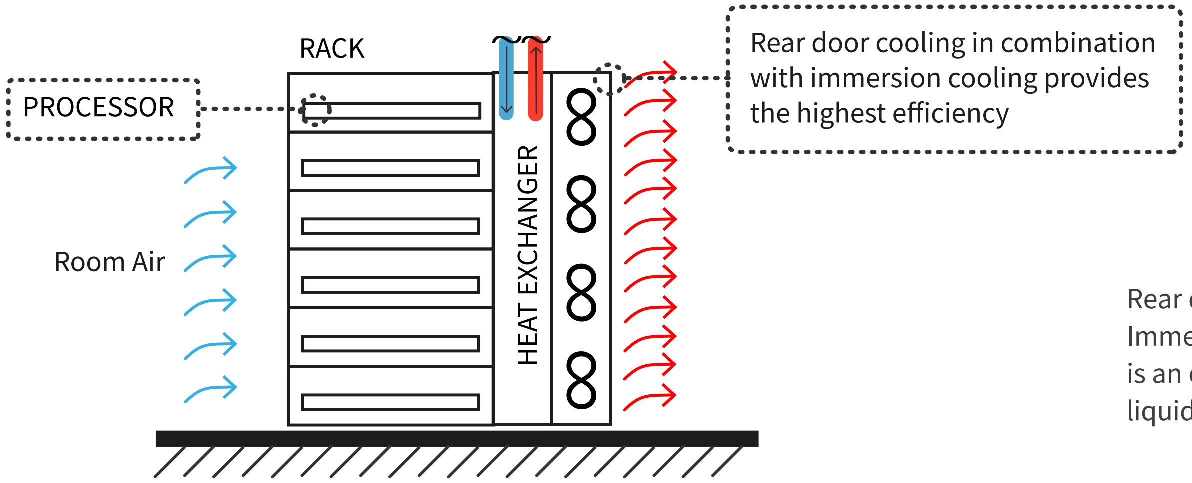

Rear door cooling in combination with Immersion or Cold Plate Liquid Cooling is an optimum hybrid solution for air and liquid cooled servers

Engineering the Cooling

Hybrid Direct-to-Chip and CRAH/Fan Wall with Series FWS flow

The FWS removes heat from the CDU. The CDU includes a heat exchanger and pumping system. Treated fluids, often PG25 (25% glycol and 75% treated water). This same FWS serves CRAC (Computer Room Air Conditioning) units or AHUs (Air Handling Units. Redundant piping and pumps are not shown for clarity.

Figure 9. Immersion tank schematic

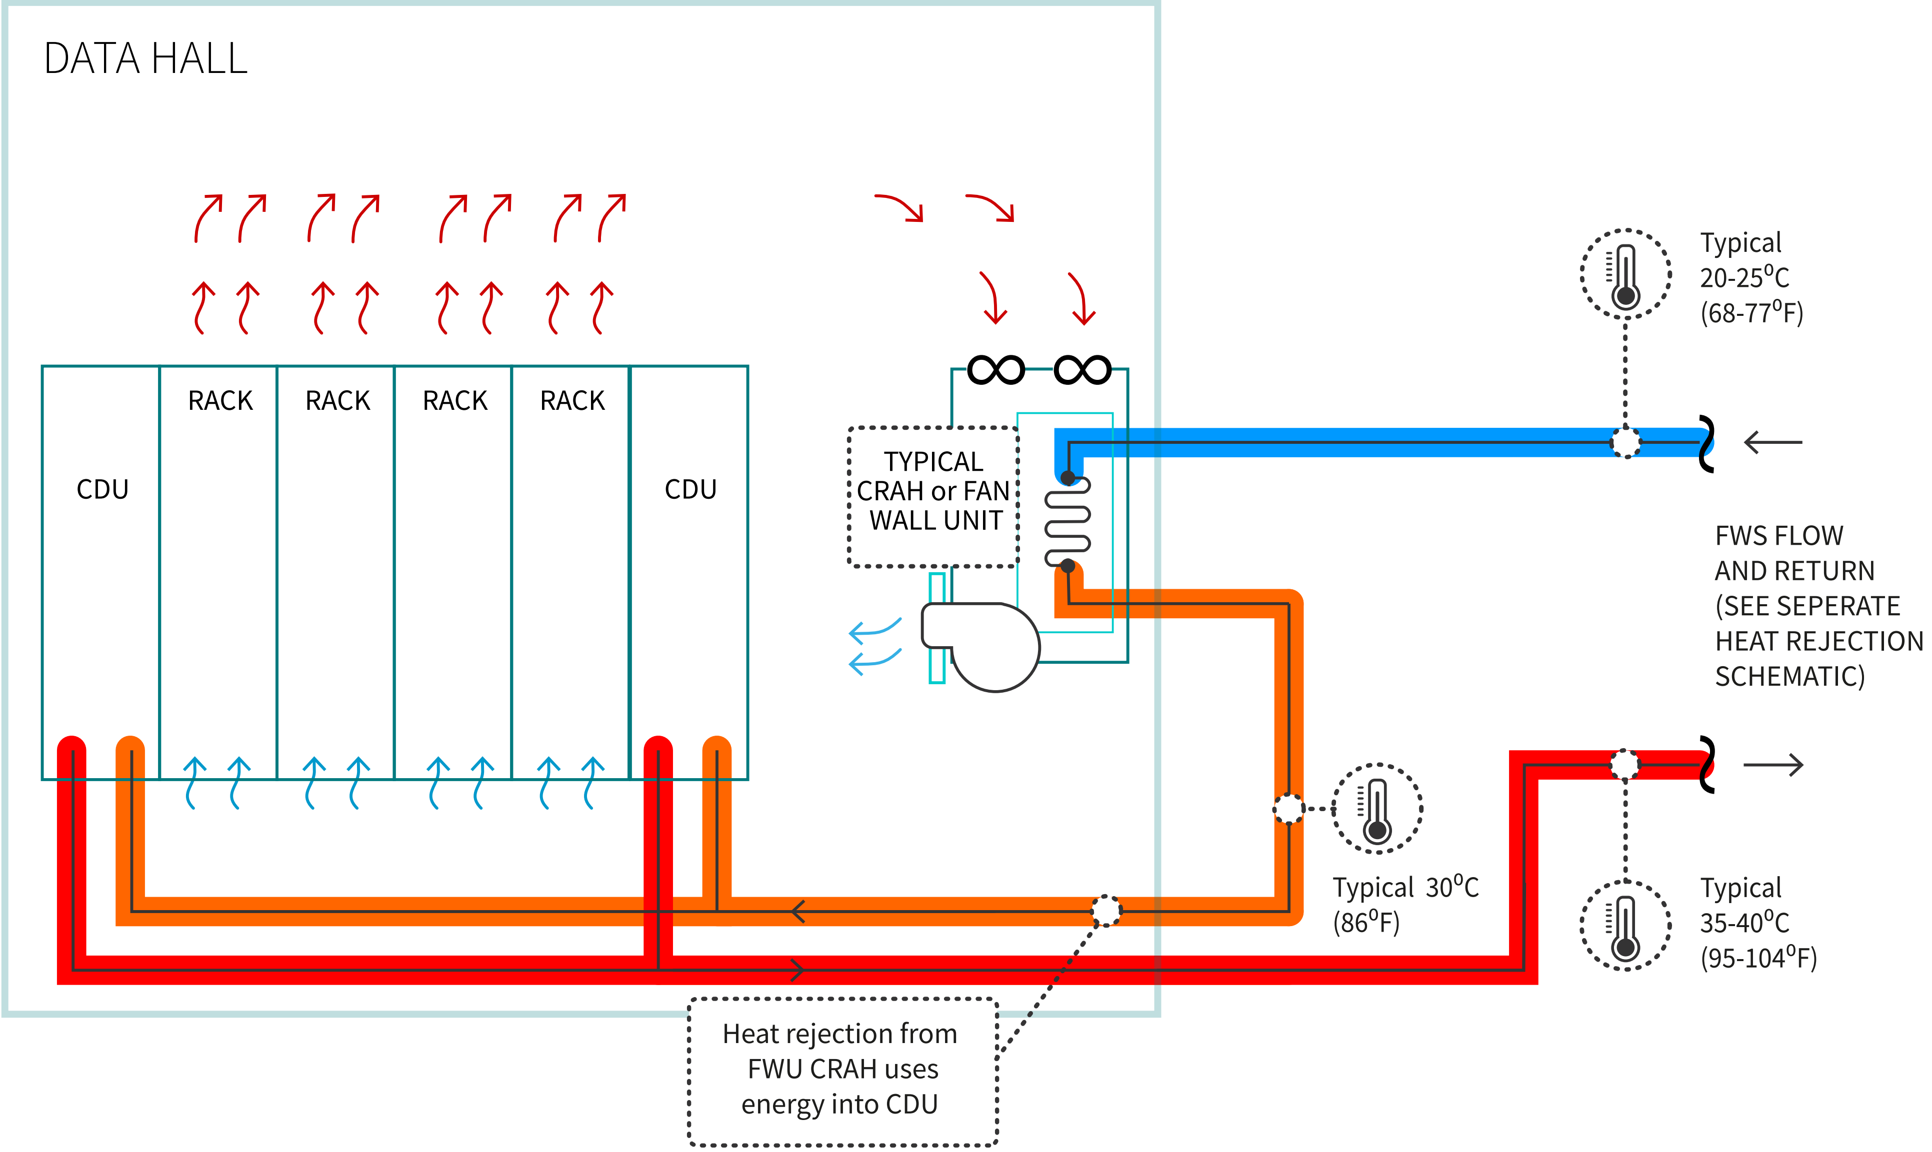

Engineering the Cooling

Hybrid Heat Rejection with Chiller Assist

A hybrid system comprises both air and water cooling to servers. This can lend itself to series cooling from the air to water cooling devices, reducing water flow and the capacity of the associated distribution services. For heat rejection, a chiller assist arrangement can be used to minimize the installed capacity of mechanical cooling. Using cooling towers or adiabatic coolers will increase the level of free cooling.

Figure 10. Water-cooled chiller schematic

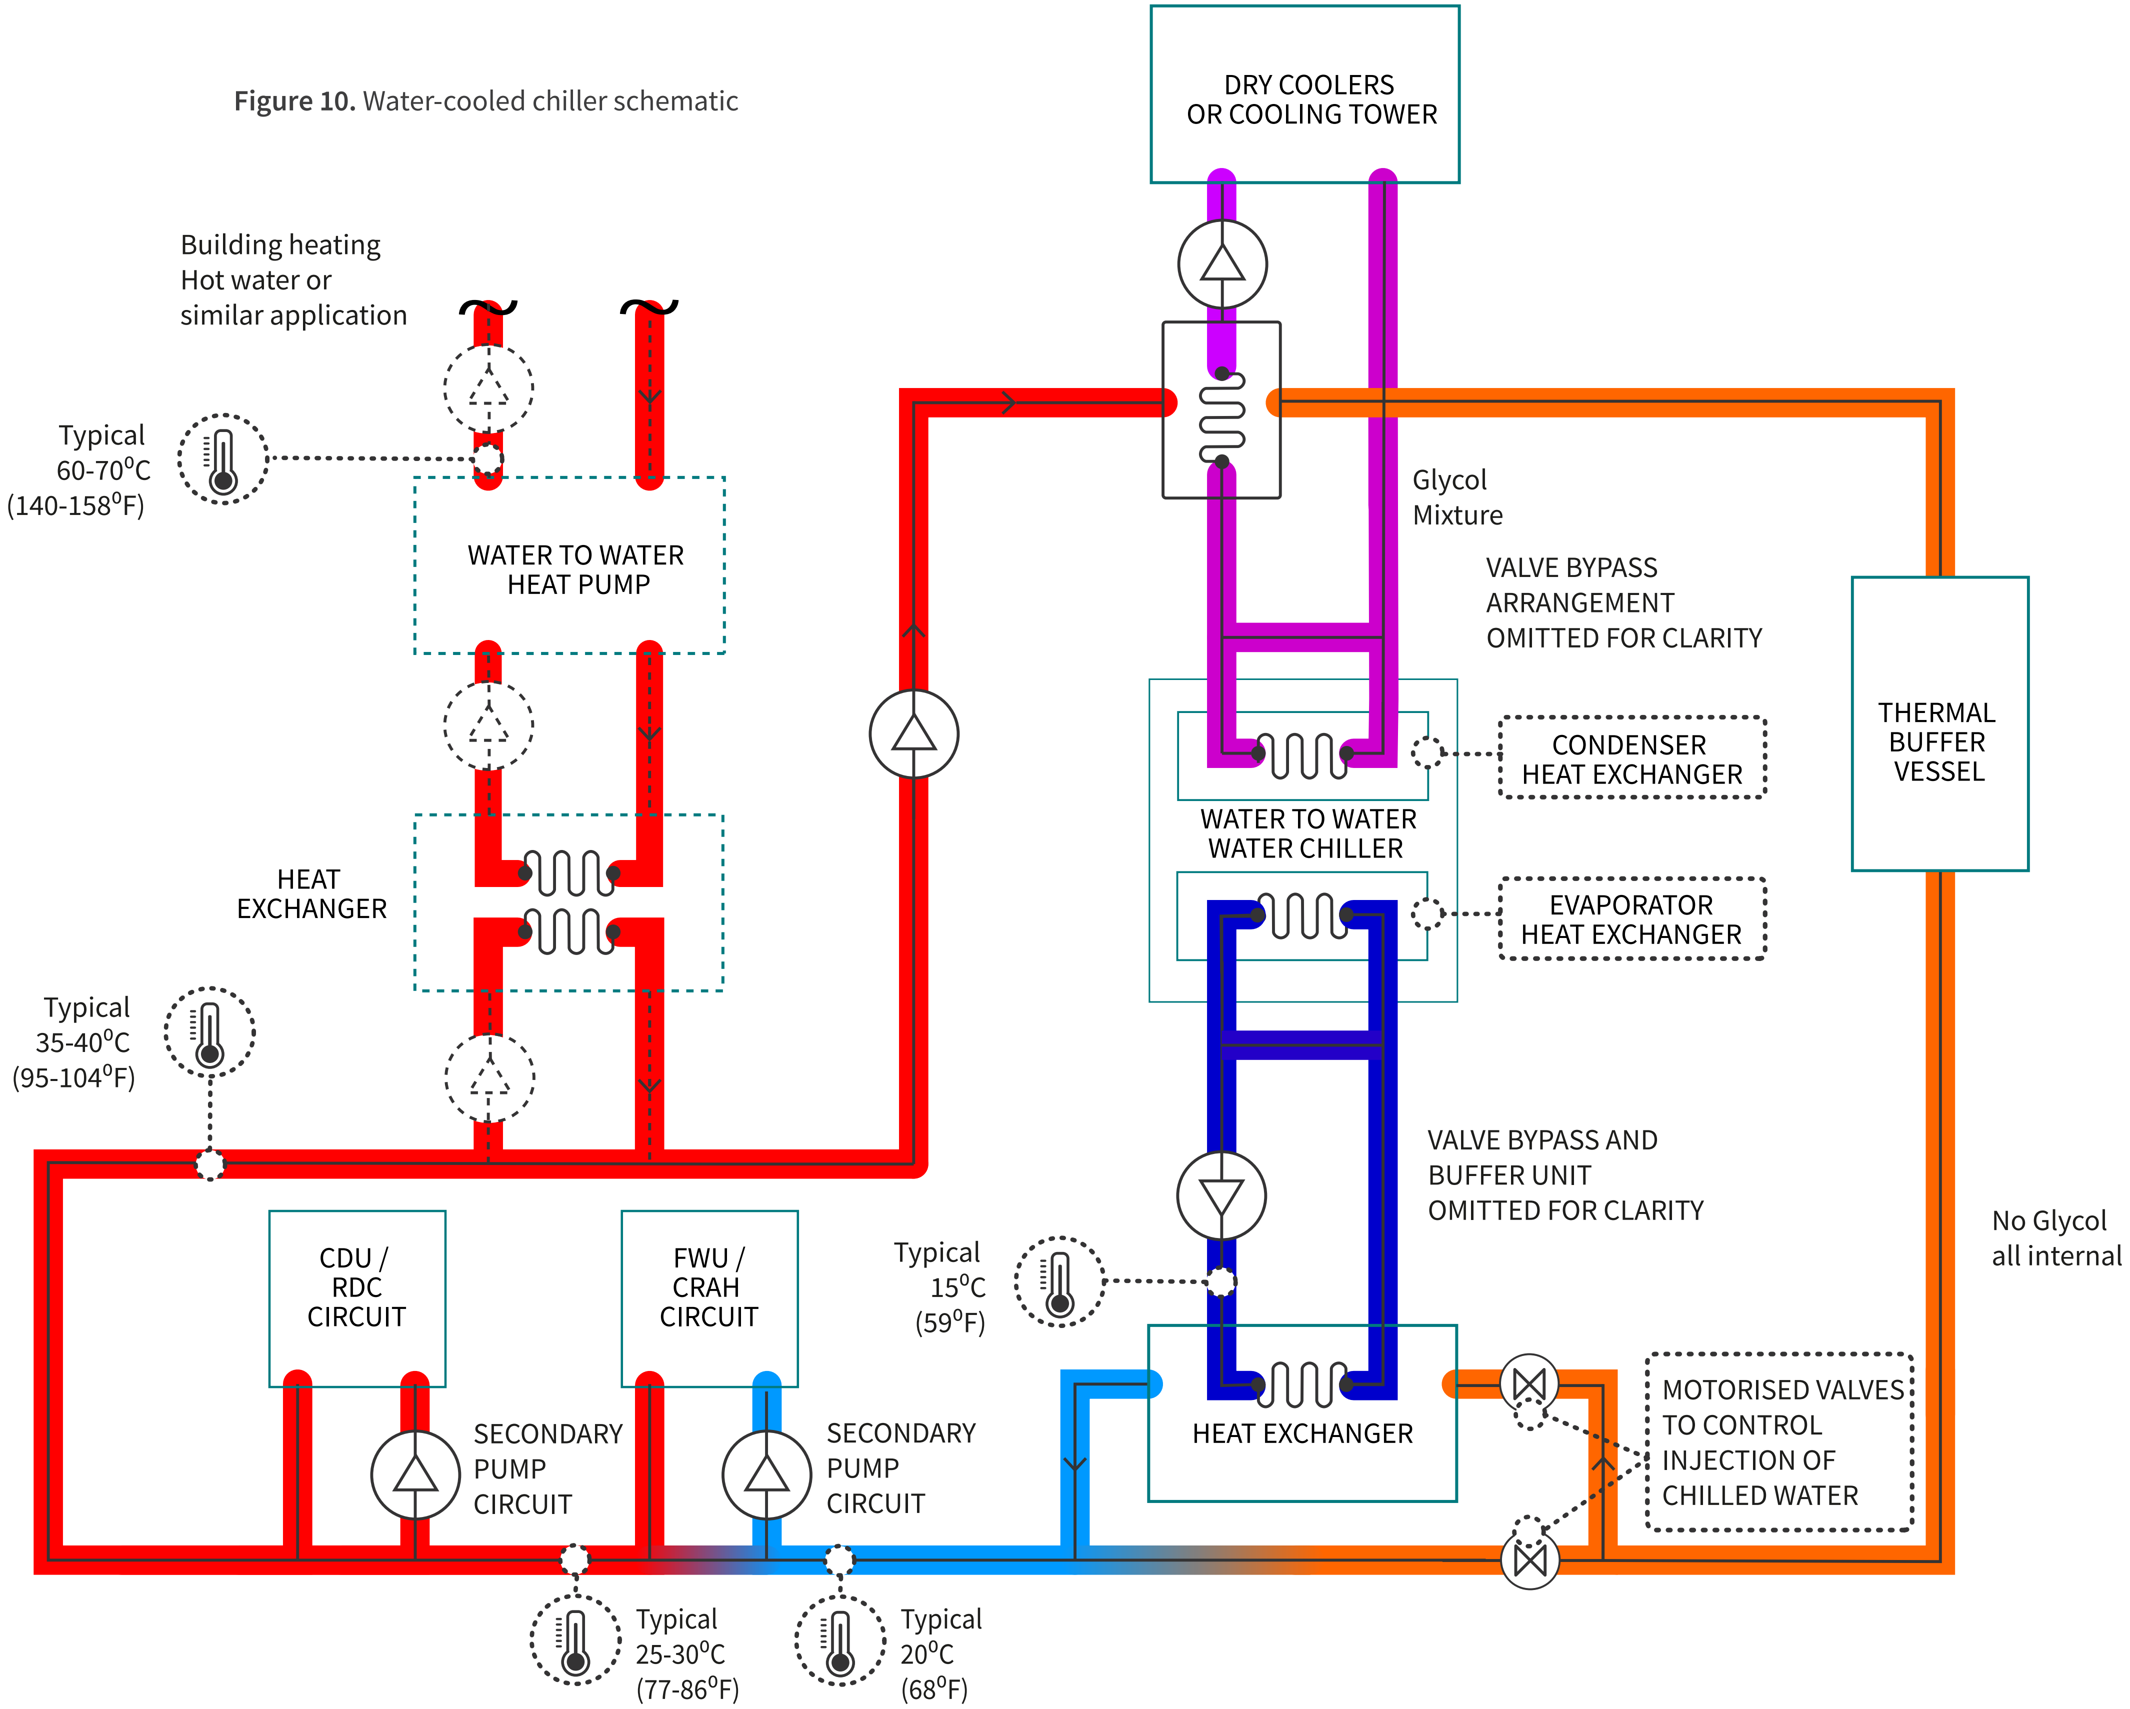

Engineering the Cooling

Energy for Cooling Equipment

As described above, higher water temperatures would lead to more efficient heat rejection, lower running costs and reduced carbon emissions. The water-cooled server requirements currently limit this scope but there are other benefits:

1. Water-cooled servers remove much of the fan cooling, equating to 10-15% of the overall computing power.
2. The omission of computer room air handler (CRAH) or fan wall units (FWU) can save 5-10% overall power.

Overall, using water-cooled servers will provide efficiency benefits and, as these are being used on high-power deployments, the savings will be significant.

There is still little integration between the IT systems and the cooling and there are opportunities for a proactive approach that would provide further energy savings. This would only apply to owner operators who have control over the IT systems. An evolving metric, TUE (Total-power Usage Effectiveness), could provide a more holistic approach to efficient data usage. [25]

Water Quality

Liquid cooling to AI servers requires high-quality fluid to provide reliable operation as well as effective heat rejection. The current server arrangements all use cooling distribution units or similar heat exchangers that hydraulically separate the facility water system (FWS) from the server /technology cooling system (TCS). The latter have their own downstream filtration system.

The FWS must align with the following, which is standard quality for a sealed cooling system but with enhanced filtering to protect the heat exchangers.

Liquid Cooling Guidelines for Datacom Equipment Centers
Second edition (ASHRAE 2014)

- A pH between 7 and 9.
- Consistent benzotriazole (BTA) concentration or other corrosion inhibitor.
- Regular use of a deionization filter is needed, after which BTA replenishment to the recommended concentration will be necessary.
- The water should be filtered to remove solid particles with sizes 10 microns and larger or at a minimum of the smallest passage in the system.

The table below shows the differences between the facility water system (FWS) and the technology cooling system (TCS). This paper is based on the FWS and TCS being downstream of the CDU.

Table 2. Water quality guidelines for the FWS and TCS

Parameter	**FWS** (Table 5.3 ASHRAE 2014)	**TCS** (Table 6.2 ASHRAE 2014)
pH	7 to 9	8.0 to 9.5
Corrosion inhibitor(s)	Required	Required
Biocide	—	Required
Sulfide	<10 ppm	<1 ppm
Sulfate	<100ppm	<10ppm
Chloride	<50 ppm	<5 ppm
Bacteria	<1,000 CFUs/mL	<100 CFUs/mL
Total hardness (as $CaCo_3$)	<200 ppm	<20 ppm
Conductivity	—	0.2 to 20 microhmo/cm
Total suspended solids	—	<3 ppm
Residue after evaporation	<500 ppm	<50 ppm
Turbidity	<20 NTU (Nephelometric)	<20 NTU (Nephelometric)

Filtration levels are also key to liquid cooling to servers with in-line duplex strainers required on the final feed to CDUs. The strainer mesh will be sized to suit the specific application and will be around 500 microns (mesh size 35).

The “open commute project” provides updated recommendations from water quality experts. [33] Water quality guidelines are shown on the following page.

Engineering the Cooling

Table 3. “Open commute project” water quality guidelines

Parameter	Performance	Notes
Total Suspended Solids (TSS)	< 5 ppm	Normally should not be present and ideally should be close to 0. There should be no visible solids.
Total Dissolved Solids (TDS)	< 1000 ppm after treatment, then monitor rate of rise	Related to conductivity and may increase with water treatment chemistry. Establish baseline after treatment and track changes over time.
Conductivity µS/cm or µmhos/cm	< 1,500 µS/cm @25°C after treatment, then monitor rate of rise.	Related to TDS. Addition of chemicals can affect conductivity. Baseline prior to and after treatment. If conductivity rises during operation, analyze to determine the cause, and take corrective action. Temperature dependent.
Corrosion byproducts/Ionsmonitor rate of rise.	Cu ions: < 0.2 ppm Fe ions: < 0.1 ppm	Baseline metal ions using Inductively Coupled Plasma (ICP) Lab test for soluble ions including iron (none expected), and copper should be less than 0.1 ppm at fill. If any metal ions are high in the fill water, consider remediation in the treatment plan. Investigate the cause of any metal ion increases over time (e.g., Cu ions > 0.2 ppm).
pH	Fill: 8.5 -10.5 Operational: 7.5 -10.5	Optional buffer (e.g., with borate) to maintain higher pH. Operating above a pH of 9.5 may reduce biocide additions.
Total hardness Total Ca / Mg as $CaCO_3$	< 30 ppm	High hardness values indicate the use of poor makeup water quality
Turbidity	<5 Nephelometric Turbidity Unit (NTU)	Lab test. The measure of particles in a fluid that affect the clarity of water. Fill should be clean and the quality should not deteriorate. There should be no visible discoloration or opacity.
Microbiological control – bacteria	Fill: <1 Colony Forming Unit (CFU)/ml, Operational: <100 CFU/mlUnit (NTU)	Fill with sterilized water with no detectable bacteria present. Use test methods for detection down to 1 CFU/ml or better. Other technology can be used for tracking microbial activity on site as long as the correlation can be established with conventional plate count enumeration.
Corrosion Inhibitor: Azoles (e.g., tolytriazole-TTA or other azole products)	>100 ppm or per treatment plan	For corrosion protection of copper and yellow metals. Azoles levels increased due to high surface to volume ratio. Other corrosion inhibitor, e.g., molybdate is not expected. Use only if required. Do not use nitrites (potential nutrients for microbes).
Dispersant polymer	5-20 ppm (typical range, if required by treatment plan)	<50 ppm if 304 or 321 stainless steel (SS) present
Chloride	<50 ppm if 304 or 321 stainless steel (SS) present	Confirm Chloride level compatibility with all SS in the system. Chloride must be low if 304 or 321 SS is used. Chlorides can concentrate in the system. Chloride may also require more azole and may interact with molybdate.

Sustainable Energy Solutions

Sustainability is already a key driver in data center design and operation, and this will become even more important as the power demand for AI servers continues to grow. Tetra Tech will publish a follow-up paper on “Sustainable Energy Solutions” to cover both low- and no-carbon technologies.

Focus of this Green Paper

This paper will:
- Analyze the GHG emissions associated with liquid-cooled AI data centers.
- Explore effective carbon offset strategies.
- Provide recommendations for achieving carbon neutrality.

Key takeaways:
- Understand the magnitude of power requirements for new liquid-cooled AI data centers.
- Heat reclaim opportunities from liquid-cooled AI data centers.
- Diverse offset options exist but require careful evaluation.
- Achieving carbon neutrality is possible but demands a strategic approach.

One specific feature of adopting liquid cooling will be better opportunities for heat reuse. The increased water temperatures will reduce the uplift requirements from heat pumps, which should be included in all new heat rejection systems. A big challenge for this technology is finding a local user, which will become a factor for future developments.

> **There are approximately 8,000 data centers globally, consuming over 2% of the global energy. With the increased cooling demand from liquid-cooled AI data centers, reclaiming heat and utilizing green and renewable energy will be a requirement. Tetra Tech High Performance Buildings Group is publishing a follow-up ‘Green Paper’ on Sustainable Solutions for Liquid-Cooled AI Data Centers.**

Summary

Taking the Heat Out of AI

With the dramatic increase in the deployment and use of AI, an uptake in liquid-cooled servers is expected in 2024. As AI server power continues to rise, liquid cooling will be required across many data centers.

The cold plate vs immersion debate will gather pace as liquid cooling deployment increases and the server manufacturers will have a big say in the outcome.

For cooling engineers, both provide a similar challenge of removing substantial amounts of heat, but the final server cooling technology has a limited impact on the FWS. Both systems currently need water with an optimum (on server) temperature of around 30°C (86°F), requiring FWS flow temperature at circa 25°C (77°F). The key will be to design flexibility into the FWS along with the necessary resilience and thermal ride-through capabilities.

The bigger engineering challenge, and FWS design decisions, will be for hybrid cooling arrangements that need to support both air and liquid-cooled servers. To allow for future increases in AI server loads, facilities must include much more flexibility in the FWS design, along with more space for future heat rejection and the corresponding flexibility in the electrical infrastructure. This is where cooling system development now needs to focus.

Our job as engineers continues to be to reject this heat as effectively and efficiently as possible but also look much closer at opportunities to re-use heat wherever possible.

References

1. NVIDIA, *NVIDIA H100 Tensor Core GPU Datasheet*, https://resources.nvidia.com/en-us-tensor-core/nvidia-tensor-core-gpu-datasheet, 7/2023
2. Eren Çam, Zoe Hungerford, Niklas Schoch, Francys Pinto, Miranda, Carlos David Yáñez de León, *International Energy Agency Electricity 2024, Analysis and forecast to 2026*, 1/2024
3. Intel, *Intel® Core™ Ultra 5 processor 125H, 18M Cache, up to 4.50 GHz Specifications Sheet*, https://www.intel.com/content/www/us/en/products/sku/236848/intel-core-ultra-5-processor-125h-18m-cache-up-to-4-50-ghz.html, No date
4. Jessica Gullbrand (Intel), Nigel Gore & Jason Matteson (Iceotope), Elizabeth Langer (CPC – Colder Products Company), *Open Compute Project Liquid Cooling Cold Plate Requirements Document*, 10/9/2019
5. Ruiyu Sun, Yongwei Li, Shifeng Wang, Guofeng Chen, & Qingming Fu (JD.com), Jun Zhang, Nishi Ahuja, Yuyang Xia, Hongxing Zhou, Jinwen Yang, & Qing Qiao (Intel), *An Advanced Liquid Cooling Rack Design for Data Center*, 6/4/2020
6. Frank van Gool, *CFD Study Chilled Water Buffer Vessel*, 6/21/2021 (Confidential)
7. Cam Turner (CoolITSystems), Christopher Chapman (Boyd), Benjamin Sutton (Calyos), Elizabeth Langer (Colder Products Corp), Jessica Gullbrand (Intel), Jordan Johnson (Intel), Sean Sivipalan (Meta), Jeff Grau (Parker Hannifin), Joshua Harmon (Staubli), Drew Tuholski (Vertiv), Mick Jones (Vertiv), Lilach Butchmits (ZutaCore), Shahar Belkin (ZutaCore) and Philip Yu (Veolia WTS), *Cold Plate Cooling Loop Requirements, Rev 2*, No date
8. Samantha Yates, David Zhou, and Edward Kung (Intel), Huang Jiaming, Zhao Shuai, and Li Rengang (Inspur), Rachit Malik, Jordan Johnson, Liguang Du, Na Chen, and Berhanu Wondimu (Intel), Wu An, Li jinbo, Bai Xinlu, Ren Yuying, and Liu Guangzhi (Inspur), Cam Turner and Jerry Kao (CoolIT), Hao Jingyang, Liu Xuejiao, and Max Chen (Lenovo), He Yongzhan (Baidu), Edward MC Wu (Wistron), Songwei Lee (AVC), Benny Huang and Akash Sengupta (Cooler Master), Nick Tan (Delta), and George Lee (BOYD), *White Paper: Cold Plate Development And Qualification*, No date
9. Drafter Unknown, *Cooling Options for HPC Diagrams*, No date
10. Cray Shasta, Site Preparation Guide, Preliminary, 12/2009
11. Bill Kosik (DNV), Brian Rener (SmithGroup), Ameya Soparkar (Affiliated Engineers Inc.), and Robert Sty (HDR, Inc), *Data Centers, Fall Edition*, 2023
12. Don Mitchell (Victaulic); John Menoche (Vertiv); John Gross (JMGross Engineers); Vali Sorell (Microsoft), John Musilli (CPS/Integra), John Bean (GRC); Jorge Padilla(Google); Jeremy Rice (Google); Nishi Ahuja (Intel); Mark Lommers; Michael Gonzalez (CEJN); Cosimo Pecchioli (Alfa Laval); Le Yu; Patrick Giangrosso; Aaron Duda; Brian Evans; Rich Donaldson, Thomas Squillo; Jack Kolar; Sean Sivapalan (Intel); Jason Rafkind; Nishi Ahuja, (Intel); Michael Beatty (Nalco); Brandon Peterson (CoolIT), Jaclyn Schmidt (CoolIT), Masud Karim, John Peterson, Jason Matteson (Iceotope), John Groenewold, Marcus Moliteus (Aligned Data Centers), Matthew Winter(Global Switch), Joe Capes (LiquidStack); Gerard Thibault (KAO), Dale Sartor, David Quirk, Herb Radlinger, Mark Dansie, Bret Lehman (PCX Corp), Madhusudan Iyengar (Google), Caleb Lusk (Rittal), Hamid Keyhani, Rolf Brink (Asperitas), John Fernandes (Facebook); Sam Allen (Burns & McDonnell); Sean Sivapalan (Intel); Rob Bunger (Schneider); Isabel Rao (CoolIT), Raul Alvarez (Submer); Rob Sty; Alex McManis (GRC); Greg Towsley (Ebara); Eugene Maritz, and Mohammad Salehi, *Data Center Liquid Distribution Guidance & Reference Designs*, No date
13. Cosimo Pecchioli (Alfa Laval), Jaime Comella (Cloud&Heat), David Sickinger (NREL), and Otto VanGeet (NREL), *Data Centers Heat Reuse 101*, No date
14. Jabari George, John Fernandes (Meta Platforms Inc.), Juan Carlos Cacho Alonso (Rittal), Kenneth Kiernan (Amphenol Corporation), Michael Thompson (nVent Schroff GmbH), and Philippe Boisvert (Boyd), *Evaluating the Limits of Door Heat Exchanger Solutions Optimized for Open Rack V3*, No date
15. Don Mitchell (Victaulic), John Menoche (Vertiv), John Gross (JMGross Engineers), Vali Sorell (Microsoft), John Musilli (CPS/Integra), Michael Gonzalez (CEJN); Tim Marquis (Parker), John Bean (GRC); Jorge Padilla(Google); Jeremy Rice, (Google); Nishi Ahuja, (Intel); Mark Lommers; Cosimo Pecchiol (Alfa Laval); Brian Evans (WSP/KWMC); Rich Donaldson (WSP/KWMC); Thomas Squillo; Jack Kolar; Bret Lehman (PCX Corp), Madhusudan Iyengar (Google); Caleb Lusk (Rittal), Hamid Keyhani (Meta), Rolf Brink (Asperitas), John Fernandes (Meta), Sean Sivapalan (Intel), Rob Bunger (Schneider); Isabel Rao (CoolIT), Raul Alvarez (Submer); Alex McManus (GRC); Rich Whitmore (Motivair); Sam Allen (Burns & McDonnell); Greg Towsley (Ebara); Philipe Boisvert (Boyd); Rob Bunger (Schneider); Ozan Tutunoglu (Nortek); Timothy Shedd (Motivair); and Gary Tinkler (Usystems), *Guidelines For Connection Of Liquid Cooled Information Technology Equipment (ITE) To Data Center Facility Systems*, 4/25/2022
16. Jeff Grau (Parker), Michael Gonzalez (CEJN), Tim Marquis (Parker), Joe Engel (SafeWay), Nick Goenner (SafeWay), Kenneth Kjellberg (CEJN), Amanda Bryant (Eaton), Glenn Charest (Facebook), John Fernandes (Facebook), Ben Kim (Facebook), Michael Gonzalez (CEJN), and Andrew Wasielewski (CEJN), *Hose and Manual Couplings – Best Practices, Rev. 01*, No date
17. Rolf Brink (Promersion), Jessica Gullbrand (Intel), John Bean (Schneider Electric), Nigel Gore (Iceotope), Rick Payne (Flex), Jimil Shah (TMGcore Inc), Rick Margerison (TMGcore Inc), Kevin Wirtz (Cargill), Kristin Anderson (Cargill), John Bean (GRC), Andy Young (Asperitas), Ashley Hessin (Vertiv), Nigel Gore (Vertiv), Michael Jones (Vertiv), Eduardo de Azevedo (Shell), Volker Null (Shell), Punith Shivaprasad (Shell), Eleanor Jones (M&I Materials), Sayan Sengupta (M&I Materials), Raul Alvarez (Submer), David Montes (Submer), Peter Cooper (Submer), Michael Sakamoto (UL), and Kai Zhou (UL), *Immersion Requirements Rev 2.0, First Amendment*, 8/18/2023.
18. Michael Berktold, Jessica Gullbrand (Intel), Fred Rebarber, Nigel Gore (Vertiv), Matthew Archibald, Michael Thompson (nVent), and Cam Turner (CoolIT), *Leak Detection and Intervention, Rev. 1*, No date
19. Punith Shivaprasad (Shell), John Bean (GRC), Jimil M. Shah (TMGcore), Eduardo de Azevedo (Shell), Rolf Brink (Asperitas), Sayan Sengupta (M&I Materials), Kevin Wirtz (Cargill), Peter Cooper (Submer), Rick Margerison (Chain Enterprises, TMGcore), Stephen Pignato (3M), Phil Diffley (LiquidStack), Gustavo Pottker (Chemours), Mustafa Kadhim (Iceotope), Volker Null (Shell), David Thomas (Neste), and Kai Zhou (UL Solutions), *Material Compatibility in Immersion Cooling*, 11/28/2022
20. Cheng Chen (Meta Platforms Inc.), Dennis Trieu (Microsoft Corporation), Tejas Shah (Intel Corporation), Allen Guo (CoolerMaster), Jaylen Cheng (Wiwynn Corp), Christopher Chapman (Boyd Corporation), Sukhvinder Kang (Boyd Corporation), Eran Dagan (Habana Labs), Assaf Dinstag (Habana Labs), and Jane Yao (Enflame Technology Co., Ltd), *OCP OAI System Liquid Cooling*, No date
21. Cheng Chen, Yin Hang, Noman Mithani, Chris Malone, Yueming Li, Wenying Zhang, John Fernandes, Kalpak Dhake, Jaret Wyatt, Jarrod Clow, and Darron Young, *Practices and Insights into Liquid Cooling on Meta's AI Training Platforms*, No date
22. Unknown (Tetra Tech), *No Time (for the Cooling) to Die, Testing 7MW of Water-Cooled IT Load Presentation*, No date
23. Tim Shedd (Dell Technologies) and Emily Clark, Ph.D, *Performance Comparison of Five Data Center Server Thermal Management Technologies Presentation*, No date
24. Bharath Ramakrishnan (Microsoft), Husam Alissa (Microsoft), Dennis Trieu (Microsoft), Robert Lankston (Microsoft), Mark Shaw (Microsoft), Zaid Kahn (Microsoft), and Christian Belady (Microsoft), *Warm Liquid Cooled Cloud Facilities: What Do We Achieve?,* No date

References

25. Michael K Patterson, PhD, PE (Intel Architecture Group Intel Corporation), William Tschudi and Henry Coles (Lawrence Berkeley National Laboratories), Lawrence B Seibold (Silicon Graphics Corporation), David J Martinez (Sandia National Laboratories), Vali Sorrell (Syska Hennessy Group), Natalie Bates (Energy Efficient HPC Working Group): *TUE, a new energy efficiency metric; moving PUE inside the box*, No date
26. https://www.theregister.com/2024/03/27/nvidia_blackwell_efficiency/
27. https://www.nvidia.com/en-us/data-center/gb200-nvl72/
28. https://www.forbes.com/sites/moorinsights/2024/03/26/nvidia-gtc-2024-wrapup-blackwell-mediatek-omniverse-and-vision-pro/?sh=752996a94698
29. https://energyinnovation.org/2020/03/17/how-much-energy-do-data-centers-really-use/
30. Keegan Yaroch & John Cuthbert (Dow Inc.), David Miller (Third Coast Chemicals), Sean T. Sivapalan (Intel Corporation), Sean Barlett (Meta), *Guidelines for Using Propylene Glycol-Based Heat Transfer Fluids in Single-Phase Cold Plate-Based Liquid Cooled Racks*, No date
31. https://www.servethehome.com/2-phase-immersion-cooling-halted-over-multi-billion-dollar-health-hazard-lawsuits/
32. Paul Artman (Advanced Micro Devices), Tozer Bandorawalla, PhD (Intel), David Grant, PE, CEM, DCEP (Oak Ridge National Laboratory), John Gross, PE, ATD (J.M. Gross Engineering), Jason Matteson (Director of Product Strategy), David Moss (Dell), Chris E. Peterson (Dell ISG), Suresh Pichai, PE (Equinix Global Engineering), Mani Prakash, PhD (Intel), Matt Shumway (Seagate), Mark Steinke, PhD (Advanced Micro Devices), *White Paper Developed by ASHRAE Technical Committee 9.9, Mission Critical Facilities, Data Centers, Technology Spaces, and Electronic Equipment: Emergence and Expansion of Liquid Cooling in Mainstream Data Centers*, No date
33. Dale Sartor, Retired (Lawrence Berkeley National Laboratory), F. Philip Yu and Craig Myers (NALCO Water), Sean Barlett (Meta), and Sean T. Sivapalan (Intel), *Liquid Cooled Rack Using Water-Based Transfer Fluid*, No date
34. https://www.nextplatform.com/2022/05/10/intel-pits-new-gaudi2-ai-training-engine-against-nvidia-gpus/

TETRA TECH
High Performance Buildings Group

Cosentini | GLUMAC | HOARE LEA | NDY

Tetra Tech High Performance Buildings Group Research Committee

Michele Barbieri: Mechanical Engineer
Steve Bonser: Electrical Engineer
Katlyn Coolbaugh: Mechanical Engineer
David Elliott: Mechanical Engineer
Ian Fan: Mechanical Engineer
Ian Francis: Mechanical Engineer
Mitchell J. Graf, PE, LEED AP: Senior Mechanical Engineer
Derek Main: Electrical Engineer
Monica Mikulak, PE, LEED AP BD+C: Mechanical Engineer

Review Committee

William P. Bahnfleth, PhD, PE: Professor of Architectural Engineering at Penn State University
Franco Caroli: Account Manager South EMEA, Submer
David Martinez: Sandia National Laboratory
Steve Harrington, Ph.D: Founder & CEO, Chilldyne, Inc.
Mark S. Myers: Datacenter Platform Architect
Steve Straus, PE: Mechanical Engineer, Tetra Tech HPBG
Peter Gross, PE: PMG Associates Consulting and Advisory
John Sasser: Chief Technical Officer, Sabey Data Centers
Jason Pullen: Kao Data
Sam Khalilieh, PE, LEED AP: National Director, Advanced Manufacturing & MCF, Tetra Tech HPBG
Allison Boen: President, Alcatex, Inc.
Darren McSorley: NDY

www.tetratech.com
www.tetratech.com/solutions/high-performance-buildings/