Title: "How-to" Guide for Transitioning from Air to Liquid-cooled High Performance Computing Systems

Authors: Dave Martinez, Sandia National Laboratory,
David Sickinger, National Renewable Energy Laboratory

Contact: Natalie Bates, Energy Efficient HPC Working Group
nbates@lbl.gov

Abstract. This document is the result of a collaborative effort to provide a guideline for transitioning from air to liquid-cooled High Performance Computing (HPC) solutions. While the main focus of this effort is directed at HPC cooling, the same ideas are applicable and relevant for any liquid-cooled computing system.

This guide is intended for people who are responsible for operating data centers that contain high-performance computing (HPC) equipment and have not yet transitioned from air to liquid cooling. It gives advice for taking an incremental "small-steps how-to approach" for transitioning without a lot of redo and cost. It assumes that you already have a data center with air-cooled equipment, and that some modifications will be needed to support liquid cooling.

We consider this version to be 0.9. It has been extensively reviewed by members of the Energy Efficient High Performance Computing Working Group (EE HPC WG) Liquid Cooling How-to Team. We are seeking further review and input from the broader EE HPC WG membership. We are posting it on the Liquid Cooling How-to Team website and will seek feedback while at SC23. We expect Version 1.0 will be posted in early 2024. We welcome case study and other reference material. All comments, suggested edits, corrections, and other material should be addressed to nbates@lbl.gov.

Table of Contents

1 Introduction

1.1 Who is the expected Reader? Why is it important to them? What is the Scope?

This guide is intended for people who are responsible for operating data centers that contain high-performance computing (HPC) equipment and have not yet transitioned from air to liquid cooling. It gives advice for taking an incremental "small-steps how-to approach" for transitioning without a lot of redo and cost. It assumes that you already have a data center with air-cooled equipment, and that some modifications will be needed to support liquid cooling.

There are two key reasons to move to liquid cooling:

- Reduce operational cost and improve energy efficiency:

Liquid cooling is key to reducing energy consumption for HPC. This is because the heat capacity of liquids is orders of magnitude greater than that of air. Once heat has been transferred to a liquid, it can be removed from the data center efficiently.

- Stay on the chip technology road map:

Liquid cooling is not only a matter of efficiency, but a requirement to stay on the chip technology road map. "Chip manufacturers are responding to demands for higher performance with higher-frequency parts, raising thermal design power while lowering package case temperature requirements. There will be an increasing number of device inventory or stock keeping units (SKUs) that require liquid cooling" [*Emergence and Expansion of Liquid Cooling in Mainstream Data Centers*, white paper developed by ASHRAE Technical Committee 9.9, 2021].

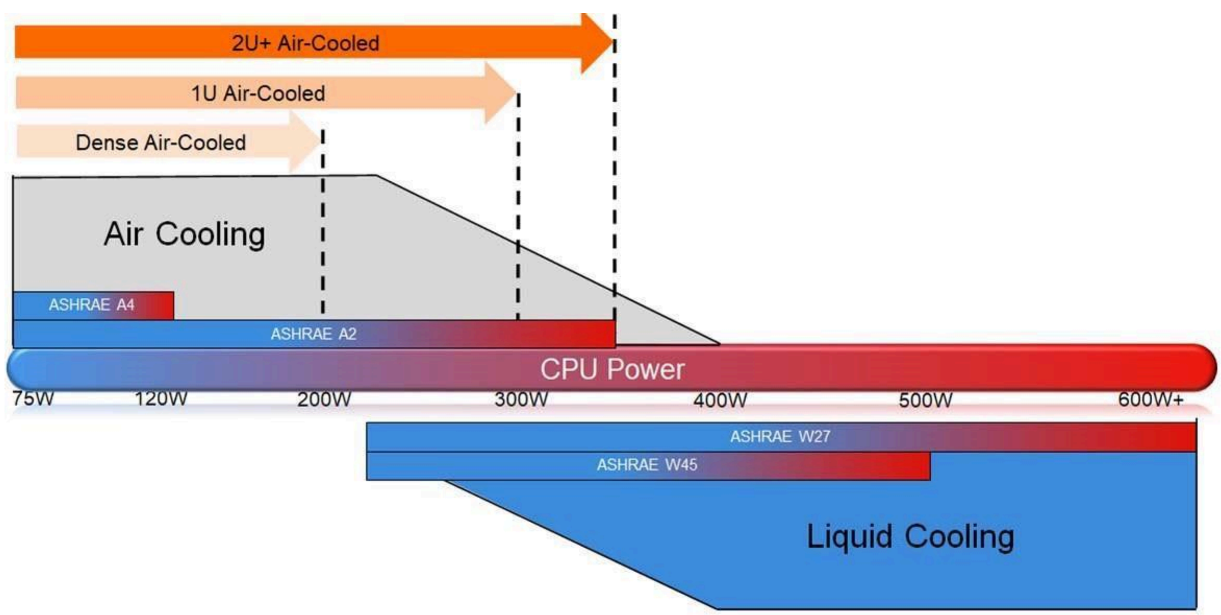

Figure 1. Air cooling versus liquid cooling, transition, and temperatures

Figure source: ASHRAE white paper

The primary authors of this guide are Dave Martinez from Sandia National Laboratory and David Sickinger from the National Renewable Energy Laboratory. Collectively, they have 30+ years of hands-on experience with liquid cooling, including multiple liquid-cooling technologies. They are not only experts, but advocates for liquid cooling in HPC.

This guide was written with extensive review and input from many other people who are active members of the Energy Efficient High Performance Computing Working Group (EE HPC WG). The EE HPC WG has 900+ members from around the globe, with ~50% of the membership from major supercomputing centers, 30% from the vendor community, and 20% from academia. The EE HPC WG encourages sharing experiences and collaboration.

This guide is intended to bring a much-needed "small-steps how-to approach" to transitioning from air to liquid cooling and to leverage, without replicating, excellent work already published by the EE HPC WG, ASHRAE TC 9.9, the Open Compute Project, and The Green Grid.

1.2 What types of liquid cooling exist?

For many, liquid cooling may seem like a confusing and overwhelming landscape. The next section of the Introduction provides a useful typology of liquid-cooling technologies to help resolve confusion.

There are a variety of different liquid-cooling technologies that can be used within data centers. Liquid-cooling fluids on the information technology (IT) loop side can range from treated water to glycol-based solutions to dielectric fluids (nonconductive). Several liquid-cooled solutions are considered hybrid solutions, in that the majority (but not all) of the heat load from IT equipment is captured and removed by the liquid—with the remaining heat load removed by traditional air cooling. However, some liquid-cooled solutions capture practically all the heat load without the use of any fans, which is the most energy-efficient. Most liquid-cooling approaches involve a cooling distribution unit (CDU), which interfaces with the facility cooling loop and provides cooling liquid at the appropriate temperature, pressure, and chemistry for the IT equipment.

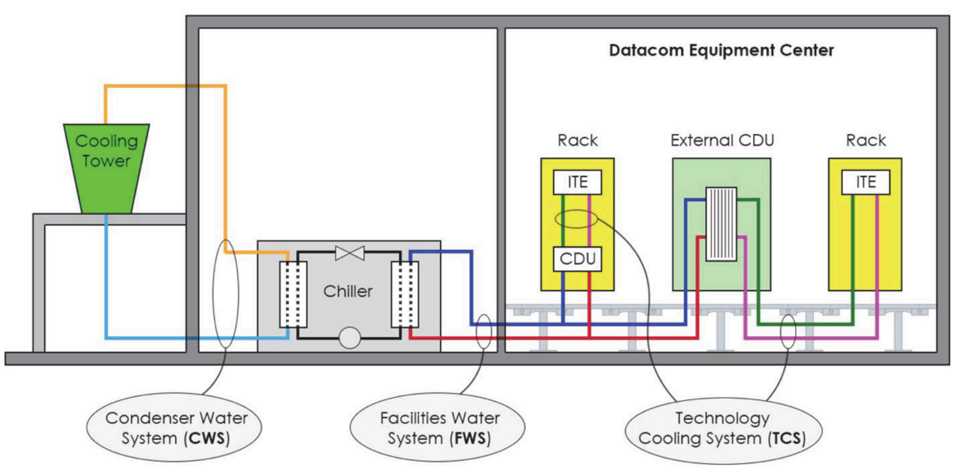

Figure 2. CDU liquid-cooling system within a data center

Figure source: ASHRAE white paper

Liquid-Cooling Technology Types

The authors categorize liquid-cooling technologies into three technology classes:

- Localized air-to-liquid heat exchanger: Involves air-to-liquid heat exchanger (coil) placed at the back of a server-rack that captures server heat, resulting in energy efficiency gains by transferring heat to liquid closer to the server-rack source (versus computer room air-handler or air-handler units).
    - Variations include rear-door heat exchanger (also referred to as cooling doors), enclosed cabinet, in-row coolers, and overhead cooling, all utilizing some type of air-to-liquid heat exchanger close-coupled with the IT source.
    - Options for cooling doors include active system where additional fans draw air through coils, or passive system where server fans move air through coils.
- Cold plates: Standard fin-based heat sinks on chips are replaced with cold plates that have liquid flowing through channels to remove heat (solution generally involves a CDU).
    - Variety of technology solutions exist, and this approach can also be applied to memory and other heat-producing components.
- Immersion: Electronics submerged in a dielectric fluid (nonconductive); approaches include:
    - Single-phase immersion cooling, in which heat load is transferred to dielectric oil/fluid that is pumped around electronics by a CDU.
    - Two-phase immersion, in which an engineered dielectric fluid with a boiling point below IT components' maximum operating temperatures removes heat load by undergoing a liquid-to-gas phase change, with that vapor transferring heat to a vapor-to-liquid heat exchanger, which then condenses back to a liquid in a passive cycle.

The cold plate and immersion classes are considered direct-liquid-cooling technologies. HPC data centers have been early adopters of direct liquid cooling due to rack power densities (the authors have experienced densities of 60 kW per compute rack in 2013, and recently surpassing 200+ kW per compute rack).

2 Operational Efficiency

Operational efficiency and cost are often overlooked in the eyes of computer system administrators, who are generally more focused on computer users, applications, and performance. It is often difficult to see over organizational silos. IT administrators and facility administrators often do not work together very closely, both being very busy and occupied with their domains. An effective organization will consider the total cost of ownership for effective decision-making.

"IT administrators see the cost of a server, but not the cost of the facility to house it or the energy required to keep it running, which over the life of the server will be **three** to five times the cost of the server." --Kenneth Brill and Gail Dutton, AFCOM Data Center Management, Sept/Oct 2008

The transition from air to liquid cooling can be seen as a part of an overall program for reducing operational costs, improving energy efficiency, and increasing resiliency. The following is a list of items to consider in addition to liquid cooling when embarking upon this type of overall improvement program.

Ongoing Consolidation of Computer Environment

Be proactive with consolidating your systems. If you are moving toward liquid cooling, make sure that you are being effective in your air side first. You will want to get rid of systems that are just idling and not really doing anything. Also, have a strategic plan for placement assistance on your data center floor. Evaluate how you are with hot and cold aisle containment. How would you do it if you are going to do it with liquid cooling—for example, with cooling doors. Sketch it out ahead of time so you will not be waiting until the last second trying to work on designs for future needs.

Decommissioning Satellite Data Centers

Especially if you have a big campus, you may have multiple data centers. You may want to try to consolidate throughout your campus or even beyond your local campus. This may lead to decommissioning some data centers that could result in greater density in your remaining data center(s)—and toward liquid cooling. In any case, ensure that you have a consolidation process and plan in place that includes room and space for future systems.

Fiber and Other Obstructions

If you have a raised floor, most of your fiber and power might be under the raised floor. One good practice is to start looking at overhead cable trays for the fiber and some of your power (perhaps using bus systems). This would allow you to make room for the liquid cooling. A good practice is to put your liquids and piping under the floor. Although it is safe to put piping on top of racks, some people have a fear of water near the racks. It also enables you to start cleaning up systems under the floor that have probably existed for many years.

Strategic Floor Use

Look at your data center floor strategically. Perhaps you have non-rack equipment like lockers and file cabinets on your data center floor. Take them off the floor. Make your data center a "data center proper" and house nothing but computer equipment on the floor. Think about how you would use that space. Perhaps you will be placing a CDU at the end of the row? You may need a small part of your data center floor for transitional space. When you are moving to liquid cooling, you may want to maintain some of your air-cooling environment. You could rack and stack some of your systems over to the side while you make way for your liquid cooling. This would allow for minimal disruption to the continuity of business.

Standardized Computer Room Air Conditioning

Take a holistic view. Look at the supporting systems outside of the data center as well as the computer systems within the data center, for which you will need to engage both facilities and IT employees. Understand what the facilities employees have planned. Is there an opportunity to do some standardization of their pieces of equipment? What do the IT employees need in the future for power, cooling, etc. When you are in this transitional phase from air-side to liquid cooling, there's a lot of interdependencies and a need to work with facilities and IT.

Raised Chiller Plant Water Temperatures

The authors refer to this as "climatizing." It is slowly, incrementally, raising the temperature of your data center—making sure you do not adversely affect the IT systems, or the IT system administrators (sys admins). There are a lot of sys admins that think there is a problem if they do not feel cold air in front of their computer and hot air coming out the back. If the process is done over a few months, then you will get to the point where it will not be a noticeable change. It will make your systems run very evenly and more efficiently.

Energy Efficiencies: Use as Critical Criteria in Purchase Decisions

There is going to be a cost difference between liquid-cooled systems and air-cooled systems. The upfront capital expenditure for liquid cooling is a little bit more than air cooling, but be prepared to show the operating expenditures and a net savings. It is pump energy versus fan energy, and this includes fans within the IT system. Liquid cooling may allow for your computer systems to run with higher clock speeds than the clock speeds that they run with in an air-cooled environment. With liquid cooling, you can actually "over-clock" at times. Higher clock speeds and overclocking may allow for the computers to get more done and save energy in the long run. You could site case studies that demonstrate this.

Installation of Hot and Cold Aisle Containment Systems To Direct Air Flow

If you are not utilizing hot and cold aisle containment, it would be a good idea to start there. If you are going to implement active cooling doors over multiple IT rows, another option is a laminar flow arrangement. This is where cold air enters front of racks in the first row, and cooling doors make the exiting air room-neutral (same or slightly lower temperature) that can serve as the supply air to the next row of racks without putting perforated tiles on the floor. While this is not direct (on-the-chip) liquid cooling, there are energy efficiency gains by transferring heat to liquid closer to the server-rack source.

Updated Environmental Monitoring System

Use your existing environmental monitoring system to look at how efficient you are running with your air-side cooling. Deploy additional monitoring as needed that records flows, as this will enable insight into how the systems all react to each other within the room. When you implement liquid cooling, your data center room might run a little warmer, but much more energy efficiently.

Take a Holistic View of the Present Mechanical Systems (Feasibility Study)

In considering the transition to liquid cooling, take a holistic view of the present mechanical systems. Look at your chilled water plant, your supporting infrastructure from the chilled water plant into the building, and how it is distributed within your data center. Then look at your computer room air conditioning units and how their fans are running; are they 80% or 90%? How are the pumps running—speeds like 90%–100%? Look at your big supporting fan systems and your mechanical cooling system, your chillers and your towers. If you were to model it, then you could see how everything reacts if you were to proceed with incorporating liquid cooling. Then you can get a good cost analysis of the benefits from this. This is an ongoing process; there are always improvements that you will want to make.

<u>Updated Environmental Present Conditions</u>

Look at what your outdoor temperatures are—maybe you can leverage the outdoor air at times.

3 Know Where You Are Through Metrics

The power usage effectiveness (PUE) family of data center metrics offers a great way to gain insight into your data center operation. If not already, you will want to gather meter data and track these metrics—especially PUE and ultimately transitioning to [total-power usage effectiveness (TUE)](). This is part of an ongoing commissioning process for continuous improvement.

Green Grid has defined several key metrics to assess the cutting-edge performance of a data center. These metrics include (formulas in Figure 3):

- PUE power usage effectiveness
- ITUE IT-power usage effectiveness
- TUE total-power usage effectiveness
- ERE energy reuse effectiveness
- WUE water usage effectiveness
- CUE carbon usage effectiveness.

Data Center Metrics: The PUE Family

Energy:

$$PUE = \frac{\text{"Facility energy"} + \text{"IT energy"}}{\text{"IT energy"}}$$

$$ITUE = \frac{\text{"IT energy"}}{\text{Total energy into the "}\textbf{Compute Components}\text{"}}$$

$$TUE = PUE \ x \ ITUE$$

$$ERE = \frac{\text{"Facility energy"} + \text{"IT nergy"} - \text{"Reuse energy"}}{\text{"IT energy"}}$$

Sustainability:

$$WUE = \frac{\text{"Annual site water usage"}}{\text{"IT energy"}} \quad \left(units \ \frac{L}{kWh}\right)$$

$$CUE = \frac{\text{"Total CO2 emissions"}}{\text{"IT energy"}} \quad \left(units \ \frac{kgCO2eq}{kWh}\right)$$

Figure 3. Data Center Metrics: The PUE Family

Figure source: NREL

3.1 PUE and the Drive to 1.0

PUE is a simple metric, although with limitations, that captures the data center energy use overhead for IT. PUE is defined as total energy divided by IT energy, where total energy includes IT as well as cooling, power distribution, and other overhead associated with the data center; for example, lighting. PUE is never less than 1.0, by definition. PUE was a very successful initiative driven by The Green Grid that has helped to galvanize data center energy efficiency improvements. The Green Grid is an organization that targets business enterprise, but PUE has been adopted by the HPC community as well. Driving toward a PUE of 1 has been an effective stake in the ground for collective action. If you are measuring PUE already, it is worth reviewing how it was implemented.

3.2 Move towards ITUE/TUE – capture the benefits of fan-less solutions

Part of "working on the 1.0" includes refining the PUE measurement methodology to account for cooling and power distribution losses inside the compute system. One of the primary gaps of PUE is that it ignores parasitic power delivery and cooling system energy use that occurs inside the IT "box." The PUE metric puts this in the denominator as well as in the PUE calculation. For example, shifting energy from the more efficient data center room fans to the lower efficiency IT internal fans may have a negative effect on the overall energy consumption but would improve the PUE. Getting the accounting right is particularly important for HPC, as many of the system vendors include these "infrastructure" functions within the system, with liquid cooling as a prime example.

The EE HPC WG has defined two metrics to give insight into this gap in energy efficiency tracking. The first is ITUE, which is analogous to PUE but for the IT equipment. ITUE is simply total power into the IT divided by the total power to the compute components. In this way, it is a PUE for the cluster. Cooling fan energy, power supply inefficiencies, and energy loss in voltage regulators are all in the numerator, while CPU and memory energy use are in the denominator. As with PUE, a value of ITUE = 1.0 is ideal and means all energy entering the box is used directly in computing.

3.3 Energy Reuse Effectiveness and Heat Reuse

With heat reuse, my PUE is 0.8. Perhaps you have heard something like this? "I am re-using waste heat from my data center on another part of my site and my PUE is 0.8!" Unfortunately, it cannot be so. The definition of PUE does not allow for a number less than 1.0. Energy reuse is not strictly adding to the energy efficiency of either the compute system or the data center; however, it can reduce the energy requirements for the surrounding environment.

The Green Grid, in collaboration with EE HPC WG, has created the ERE metric, which is defined as the total energy minus reused energy, divided by IT energy. With no heat reuse, ERE is simply PUE. But, unlike PUE, less than 1 is possible and desirable. Heat reuse is starting to be adopted and could soon be a well-tested best practice.

4 What is Your Vision?

You now have a basic understanding of:

- Why you need to go to liquid cooling and what types of liquid cooling are available
- What other things you might do to improve your operational efficiency
- What basic metrics you should implement in order to track your progress.

Prior to diving into the hands-on for liquid cooling, take a strategic look and set a vision. At a minimum, this includes building a cross-functional team. The team can work on right-sizing, total cost of ownership, and risk mitigation. It will also be the same team that will be necessary for implementing liquid cooling.

4.1 Building a Cross-Functional Team

In the foreseeable future of HPC, more speed will equate to more density, power draw, weight, etc. This is very unlikely to change, at least in the next 5 years. In the past, facility systems were pretty much an afterthought or only considered midway through the purchase of HPC systems; the facility team was brought in very late in the process, and facility systems were designed reactively. The computer and facilities are so closely coupled now that even designs for new computer systems consider the facilities impacts in the conceptual design stages. Furthermore, older data centers are having to beef up the infrastructure and modify existing structures that were never intended for this type of capacity.

It helps to be proactive, look toward the future, and build as much flexibility as you can, not only for the system that will be arriving, but for systems that are several years out.

Put together a strategic objective intended to ensure you maintain executive commitment to plan, fund, schedule, govern, identify, and prioritize the goals needed for changing current operational models. This strategic objective requires funding and commitment from the leadership team. It is imperative that there is a comprehensive understanding of the return on investment and total cost of operations for the three key elements of right-sizing corporate data centers:

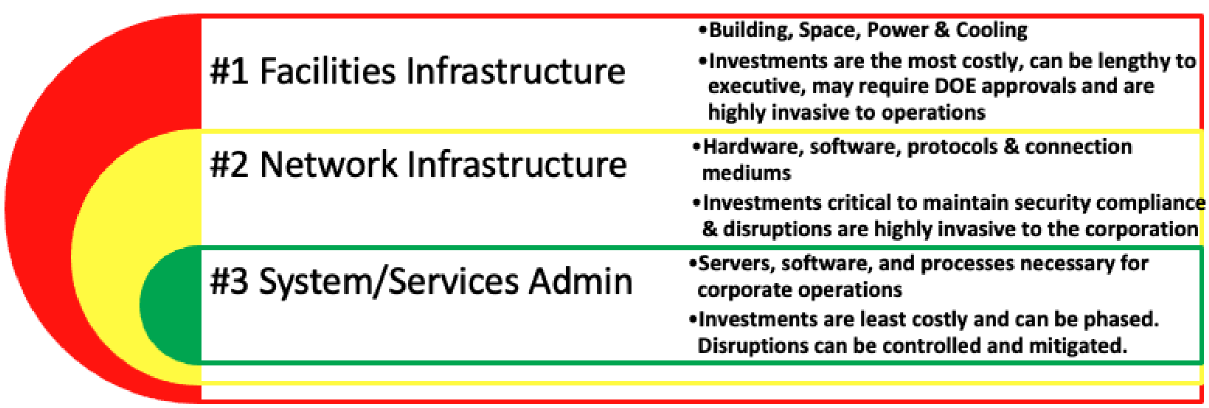

Figure 4. Key elements of right-sizing corporate data centers.

Figure source: Sandia

The interdependence of these three elements is often talked about, but the work in each of these areas is still done in silos. It is imperative that all three of these key elements are analyzed and evaluated for any IT project, solution, procurement, and deployment. If they are not analyzed, then the decision/action is based on assumptions and could be costly, and, in a worst-case scenario, render the project, solution, or deployment impossible to achieve.

This inability to ensure the three key elements are the basis for every IT decision is why right-sizing the data center seems unattainable. Continuation of a siloed operational model will continue to prevent controlling and managing the constraints and variables that will provide the ability for the data centers to be fluidly responsive to change.

4.2 Rightsizing your system

The next step is to understand what right-sizing the data center means and the questions and challenges that preclude taking real action. There is no crystal ball and no definitive markers that provide conclusive, predictive data enabling development of a concrete plan for right-sizing the data center for the next 5 to 10 years. However, there are initiatives and actions that can be undertaken to reduce and minimize the variables that result in costly, and potentially avoidable, data center modifications.

The following are examples of what can happen if you under- or over-design systems:

- Stranded capacity is a physical constraint. What this essentially means is that the data center has to have sufficient power, space, and cooling resources to ensure that IT systems can operate at full utilization. These resources are interrelated and can result in tremendous costs to a data center if not properly balanced.
- Trapped capacity is a managerial constraint. Once the calculations have been completed to ensure the system can be supported by the power, space, and cooling available, trapped capacity occurs when the workload does not use the intended power, space, and cooling that has been allocated to it. As an example: if an HPC has a technical power requirement of 10 MW and is only using 6 MW, the delta (4 MW) is trapped capacity. This power could be used for another system. So, sub-optimization of the actual IT capability is the cause of trapped capacity, and although the equipment is “on,” it is not fully optimized and produces higher operational and maintenance costs.

5 Start Implementing Liquid Cooling

Finally, it is time to dive into the hands-on for liquid cooling. The following is a suggested approach to getting started with first direct-liquid system (cold-plate or immersion) at your site. Taking these baby steps will provide staff experience working with liquid-cooling systems and interacting with vendors to configure liquid-cooled systems. The intended audience is one that has a site operations/facilities department that can assist with facility preparation.

5.1 Test bed system

The assumption for a first testbed system is a single rack utilizing an in-rack CDU that requires a 1” facility water (hydronic) system connection. While it is possible to retrofit old servers with cold plates, the authors suggest buying a rack from a vendor with servers already designed and configured for liquid cooling with an in-rack CDU included (rack appearing on left side in Figure 2, and sample of actual hardware shown in Figure 4). This approach also allows you to require the selected vendor to be responsible for the secondary fluid system (ASHRAE terminology is technology cooling system used between CDU and IT equipment), from verifying fluid material compatibility, to treatment of secondary fluid, configuring CDU setpoints, and the ongoing maintenance during a multiple-year warranty. For cold-plate based solutions, in-rack CDUs usually require 4 U of rack space and can be placed in the bottom or top of the rack (where possible, the authors prefer to run facility hydronic systems below racks to minimize risk, but also have top-fed racks).

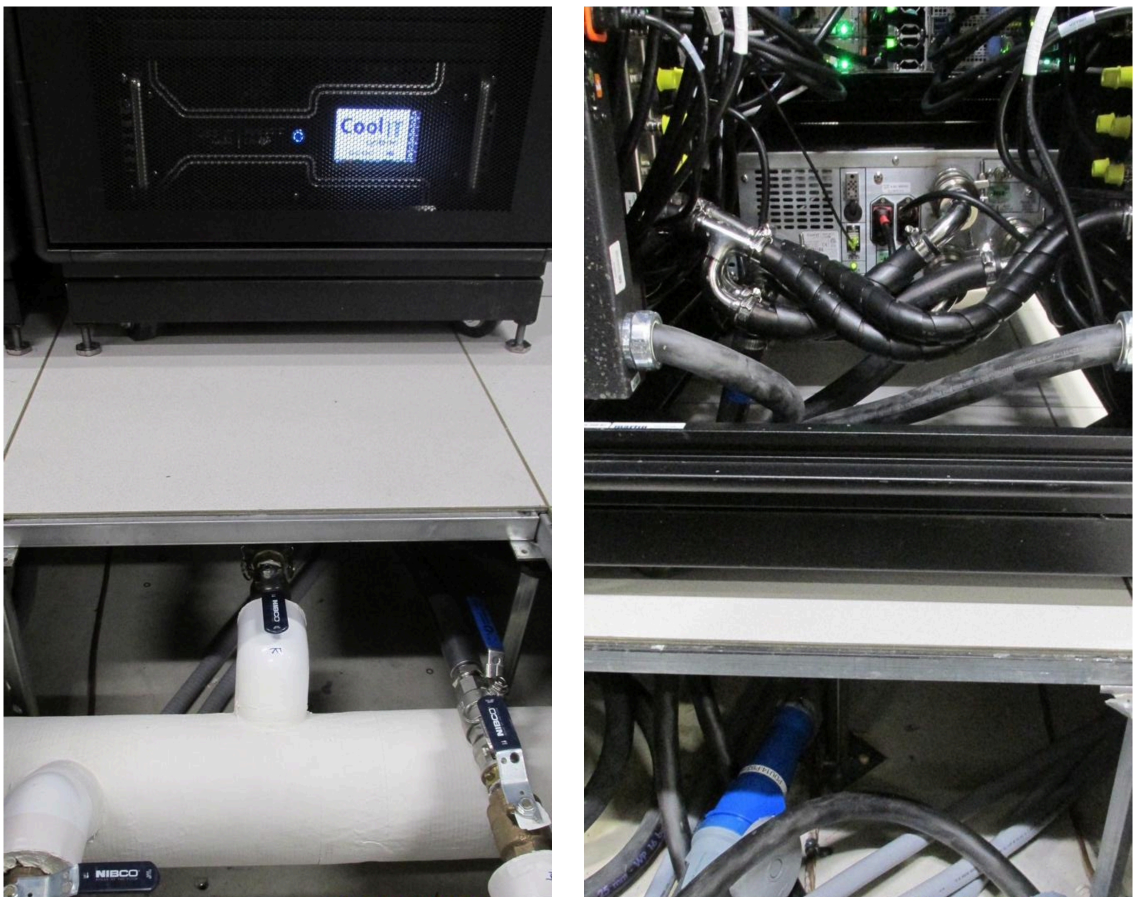

Figure 5. Liquid-cooled rack with 4 U CDU

Source: Sandia

While you can approach your testbed system as a stand-alone demo and just run a single 1” hydronic line, you might want to think about future-proofing the hydronic design to support multiple racks (goes back to your vision), as the hydronic design/installation will likely be the biggest capital expenditure. Facility supply temperatures for liquid-cooling systems (as of early 2023) are typically in the W27 (27°C, 80.6°F) to W32 (32°C, 89.6°F) range. Close coordination with vendors on supported facility supply temperatures is encouraged before purchasing a testbed system.

When talking with vendors, it is worth double-checking how temperatures are being referenced. The authors think of facility water supply temperature as what is provided (entering) the CDU, so that return temperature from the CDU back to the facility is higher. For the secondary fluid system, the supply temperature is what is delivered to the IT equipment and is the critical temperature to manage (the secondary return temperature is back to the CDU and will be higher).

There are a few approaches to facility preparation:

- For a quick setup, if you already have facility flanges/taps nearby, deploying 1” diameter reinforced flexible rubber hoses connected to existing facility flanges/taps is an option. The authors tend to run rigid pipe up close to liquid-cooled racks, then utilize flexible hoses to make the final connections; however, some have run flexible hoses 40 feet in a few cases.
- Work with the site operations/facilities department to identify closest source of facility water that meets ASHRAE W class of equipment purchasing. Facility supply can be cooler, as most CDUs include a valve on the facility side to regulate flow through CDU heat exchanger. The authors have had success in running CPVC schedule 80 pipe (supports higher temperatures and pressures than regular PVC pipe) within a data center as the last-mile hydronics—the connections between permanent facility pipes (such as black steel) and CDUs for IT racks. The pipe diameter could be slightly larger than 1” to accommodate a few IT racks (although price of piping goes up as diameter increases).
- Consider a manifold distribution system upfront (see Advanced Topics section for images). This will have a higher initial cost but will offer more flexibility over the long run (so lower total cost of operations). Work with the site operations/facilities department to have a larger diameter pipe branch run to the data center row where you eventually plan to deploy several racks. The manifold, which can be purchased or custom made, can split into multiple 1” lines supporting many racks deployed at a later date that have some combination of in-rack CDUs and/or door heat exchangers. Note that the largest stand-alone CDUs that serve multiple high-density racks through a shared secondary loop system are starting to require 4” facility connections.

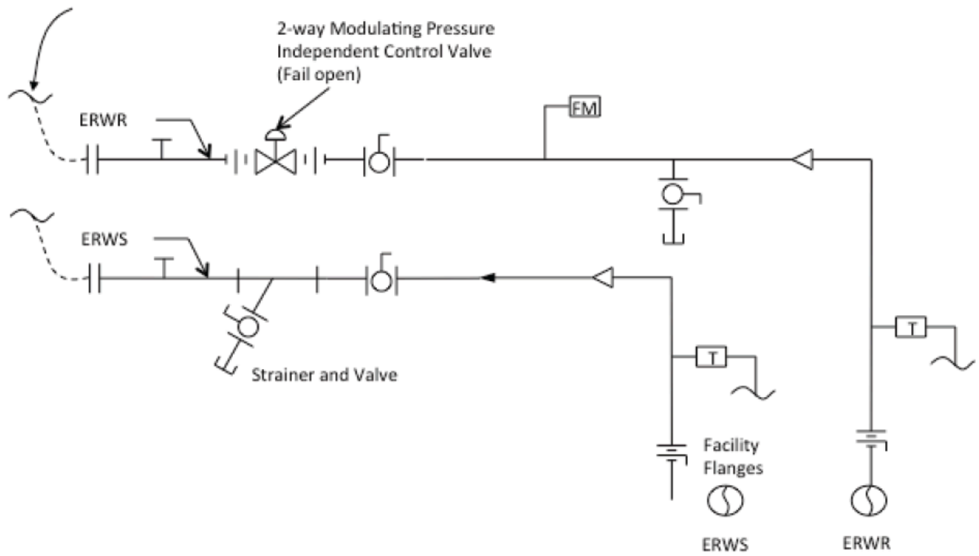

Figure 6. Sample line drawing of facility hydronics to support single 4 U CDU

Figure source: NREL

It is important to understand the liquid-cooled solution selected for testbed system—what percentage of heat load from IT equipment will be captured and removed by the liquid—and what the remaining heat load will be that relies on traditional air cooling. This should be checked during the commissioning phase.

When working with facility hydronic systems, the following are good practices to follow:

- Flush new piping sections with an adequate amount of water to remove debris and contaminants introduced during installation
- Hydrostatic pressure testing, used to check for leaks before connections to CDU are made (fill line with water for pressure testing; do not use high-pressure air in system)
- Install high-point valves to vent air from system during fills (or assist with draining sections that have been isolated)
- Verify connection fitting types (imperial vs metric) and coordinate with all parties involved.

There are plenty of data center operators (and vendors) with liquid-cooling experience. It would be worthwhile to reach out and talk with a data center operator with liquid-cooling experience to help train staff on safety, signage, policies, and other considerations before purchase, delivery, and installation of your testbed system.

<u>Commissioning Considerations</u>

Suggest including facility-related measurements in acceptance test plan to understand/verify percentage of IT equipment heat load captured direct to liquid:

- Measure total power usage of testbed rack, measuring at rack power distribution units and/or at breakers feeding rack equipment. Also capture power usage of CDU.

- Calculate heat captured direct to liquid by measuring delta-T and flow rate through CDU and comparing to total rack power. Run using three different loads and allowing enough time (~10–15 minutes) to reach a semi-steady state before calculating: idle, typical test harness, under max load. Benchmark applications such as LINPACK/HPL or SPEC's Benchmarks and Tools can be used for load.
- Recommend allowing a few days for these tests to be conducted.
- More advanced test runs:
  - Assuming initial measurements are taken with conservative temperature, you can now elevate supply temperatures to get a sense of the impact to facility infrastructure as you expand and potential efficiency gains from elevated temperatures.
  - Compare against a similarly configured air-cooled system, run same benchmarks. Something to consider when purchasing testbed rack is including mix of servers that are either air-cooled or direct-liquid cooled.

The EE HPC WG has a document specifically focused on commissioning liquid cooled systems. Read that document, especially Section 1 (1.1.8/1.1.10/1.2.1) and Section 2.1 and checklist in Section 2.3. Required data typically includes pressure, temperature, flow, Btu, kW, kWh, valve position, differential pressure, and other points.

After commissioning, work with friendly users during additional testing phase (~2 months) to allow data center and system administrator staff an opportunity to get familiar with testbed system before it goes into full-production mode with full user community.

5.2 Advanced Topics/Considerations

- Experiment with door-heat exchanger for non-compute racks
- Depending on facility conditions, consider micro-pod (requires just adding power)
- Review against vision, adjust based on testbed experience:
  - Testbed helped establish team and understanding
  - Feasibility study: work with engineers and alignment with vision.

<u>System-Level Considerations:</u>

CDU and hydronics

- Stand-alone CDUs are starting to reach MW scale, with 4" diameter supply/return feeds. These are usually server multiple racks (scalable units) with secondary loops designed to support multiple racks.
- Review electrical grounding of CDUs and liquid-cooled racks
- Monitoring of secondary fluid (both in-rack and stand-alone)—review with vendor how to get this data
- Options to review with vendor: supplied manifold above racks (scalable units) or build-your own (rack manifold/manifolds to rack manifolds).

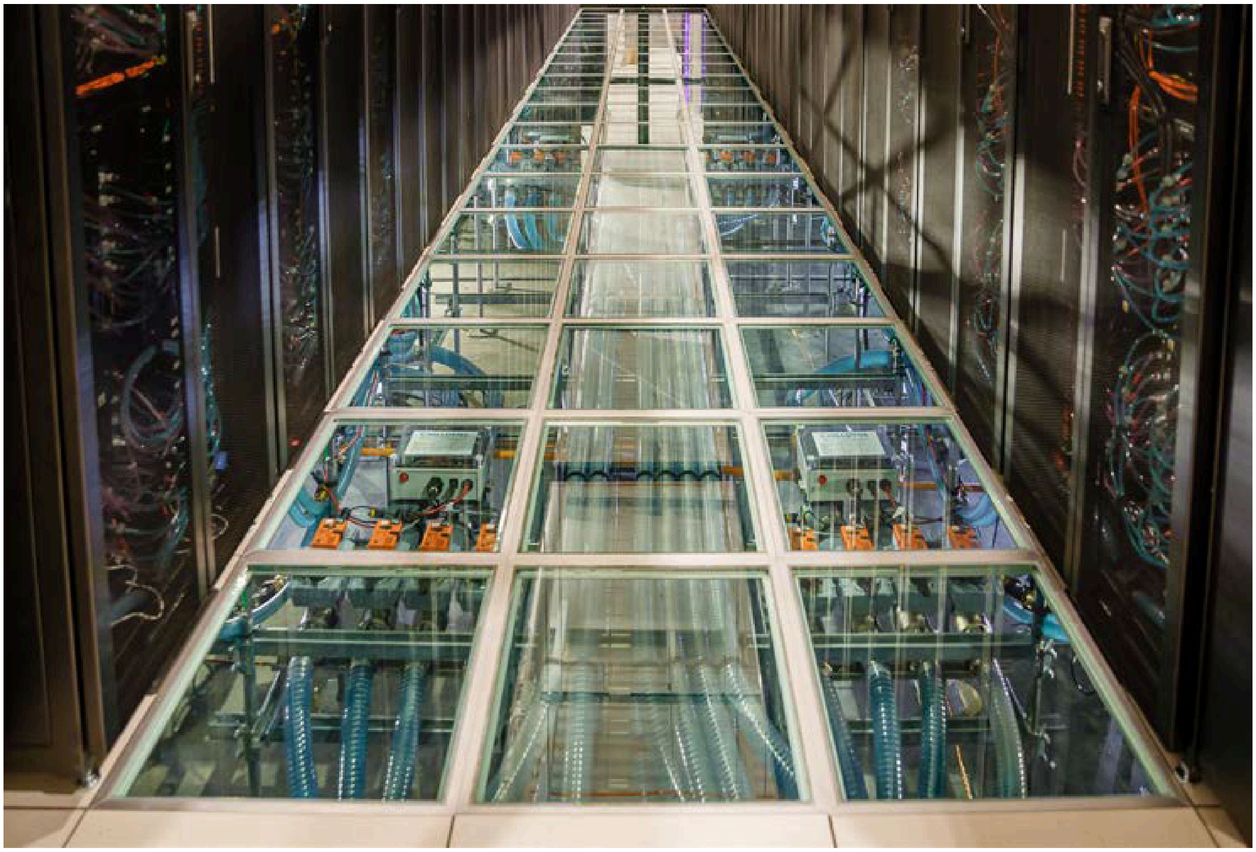

Figure 7. Sandia's liquid-cooled Attaway HPC system

Water flows through flexible tubes and pipes visible through glass panels above two rows of Attaway's compute racks. At the end of the racks, CDUs control water flow and rates. Image Credit: Sandia.

Facility hydronics: large-scale direct liquid-cooled deployment considerations:

- Install valves, circuit setter, strainer—consider how to isolate to repair as well
- Control valves
- Temperature gauges and pressure gauges (visual plus communication to gather data in database)
- Manifold distribution systems.
- Leak detection and floor drain.

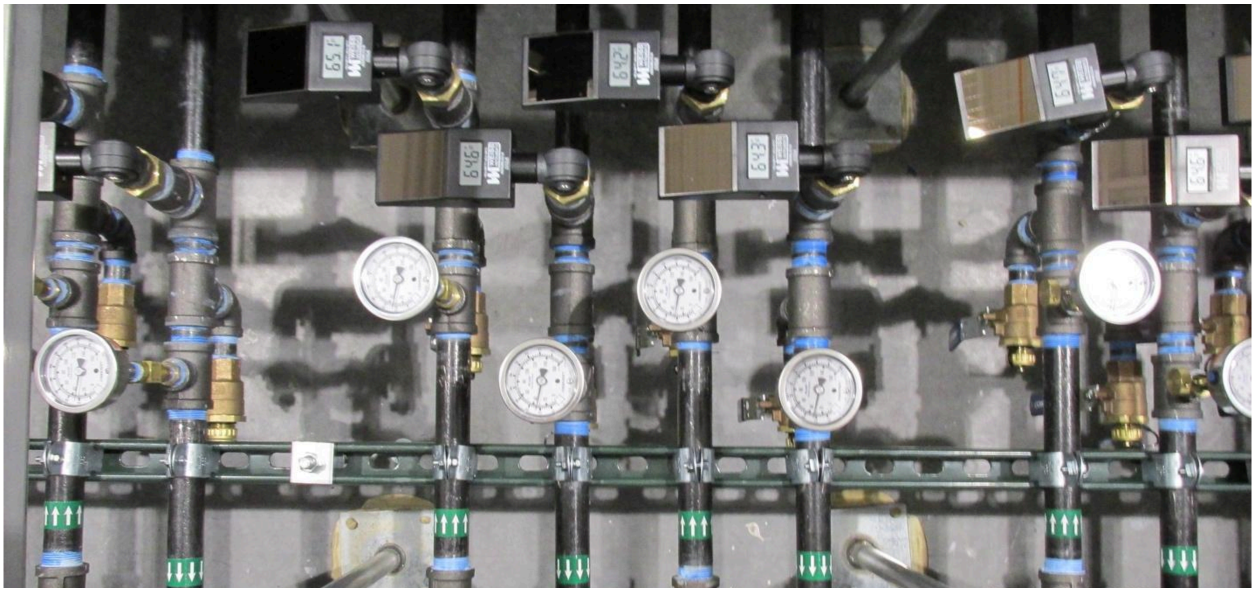

Figure 8. Sample manifold distribution system

Image Credit: Sandia

Other system considerations:

- Weight of liquid-cooled racks increasing as density increases, as well as immersion tanks.
- Safety considerations: safety data sheets on secondary fluids, signs if temperatures over 120°F (scalding considerations), establishing processes to deal with leaks.
- Resiliency approaches: N+1 pumps in CDU for scalable compute unit, operation with loss of control signal, cooling hardware that is supplied with dual-feed power.

Facility Level Considerations

Other facility considerations:

- Option to consider: containerized/modular data center approach—if severe limitations with existing data center and not planning on building a new data center building
- Facility monitoring of hydronics (elevated temperatures)
- Ongoing commissioning: possible to have multiple and different direct liquid-cooled hardware operating together, consider ongoing commissioning and review of facility/system data whenever large IT systems added or removed

The EE HPC WG has published a paper on operational data analytics that provides a model for understanding types of analytics and maps them to HPC site requirements.

6 Making the Financial Case for Transitioning to Liquid Cooling

As stated in the Introduction, there are two very good reasons to transition to liquid cooling. First, to reduce operational cost and improve energy efficiency. Second, to stay on the technology road map. When making the financial case for liquid cooling, the focus is on reducing operational cost and improving energy efficiency. Staying on the technology road map may be critical for your site's mission and meeting user requirements, but for most readers, it is likely to be an intangible benefit.

6.1 Testbed

What might a cost benefit analysis look like for a liquid-cooling testbed? The costs are site-specific and many of the benefits are intangible, so the testbed analysis will be "soft."

The benefits of the testbed:

- Gain the experience of working with liquid cooling (e.g., skills, general experience, expectations, culture)
- Test the facility side of the liquid-cooling system (e.g., water delivery, water quality)
- Show the cost-effectiveness of liquid cooling (e.g., operational expenditure testing for air vs. liquid)
- Demonstrate the performance gain of liquid cooling (e.g., time to solution, energy to solution).

The costs of the testbed:

- Although it is likely to be the highest-cost line item, whether you include the compute equipment cost in the testbed cost-benefit analysis is not clear. If the compute equipment is required and has already been budgeted for the site mission and user requirements, you might choose to exclude it.
- There are likely to be changes required to your data center to support this testbed. For example, you may have to bring a minimum of 1" facility water (hydronic) system to the new computing equipment. Other changes may include facility water hookup, taps, and instrumentation for testing (e.g., flow meters, gauges, thermometers).
- You probably want a contract to provide services for verifying fluid material compatibility, treatment of secondary fluid, configuring CDU setpoints, and handling any ongoing maintenance.
- Finally, there are costs to collecting data to demonstrate the operational expenditure and performance gains from liquid-cooled systems. This may require contracting for resources to collect data while running benchmarks and/or applications. This may require power meters as well as experienced personnel.
- If you want to get a sense of price difference of racks based on air vs. liquid cooling, you can request a quote for a similarly configured air-cooled rack (if the power density can still be supported by traditional air cooling). The additional cost of the compute equipment that is required for liquid cooling would likely be included.

6.2 Larger Compute System

Beyond an immediate and focused cost-benefit analysis for the liquid-cooled testbed, you could also do a cost-benefit analysis for a larger computer system, perhaps even for your next acquisition.

The benefits of a liquid-cooled larger computer system:

- Liquid cooling should yield reduced electricity costs to power the cooling system. Your testbed should have provided you with data that supports this assertion.
- Liquid cooling should result in lower water usage.
- The overall footprint of a liquid-cooled system is smaller than an air-cooled system.
- You may be able to reduce the total amount of heating, ventilation, and air conditioning installed and/or capital expense. You may also be able to reduce the need for higher-voltage power distribution that supports heating, ventilation, and air conditioning. This depends upon your site and the size of the larger computer system you are going to install.
- Liquid-cooling lowers sustained ambient noise, which may be a "soft" benefit.
- If you are building a greenfield site, you may be able to use a level concrete slab instead of raised flooring.
- There may be an increased opportunity for heat reuse. Some examples include heating office space or heating water for bathroom sinks.

You will need to understand and document the key assumptions of the facility cooling infrastructures being compared. Establish a baseline for your current facility. Each facility is somewhat unique, but you need to determine what will change by transitioning from air to liquid cooling. The next step is to describe the aspects of your current facility that have to change either by adding, removing, or modifying the existing cooling infrastructure. Adding pumps and pipes, removing chillers, and enhancing telemetry are all examples of required modifications.

After key assumptions are documented, you can start to lay out all of the potential costs the two scenarios of air cooling and liquid cooling. There are two major kinds of costs; these are capital and operational costs. The capital costs will most likely include mechanical systems and may also include electrical systems and building systems. Operational costs will include energy and may also include labor and other miscellaneous materials.

Here are some examples of costs to consider:

- Initial installation costs for planning, design, construction, and commissioning
- Cost of training personnel on liquid cooling
- Sustaining costs of service, maintenance, repair, ongoing commissioning
- Costs of decommissioning.

7 Liquid Cooling Resources

7.1 EE HPC WG

The EE HPC WG has several references that pertain to HPC center mechanical cooling systems, with a specific focus on liquid cooling. These references and more can be found on the Lawrence Berkeley National Laboratory Center of Expertise for Energy Efficiency in Data Centers website.

Hot for Warm Water Cooling provides insight into how you might set a target for your site's ultimate energy-efficient liquid-cooling solution. Sites across multiple geographies face different environmental conditions and cannot all adopt the same liquid-cooling solutions. This paper analyzes the wet and dry bulb temperatures for geographies distributed all over the United States and then describes a typology of liquid-cooling solutions that are appropriate for those different geographies.

Systematic Approach for Universal Commissioning Plan for Liquid-Cooled Systems provides a lot of good advice on the fundamentals that every site should consider for liquid-cooled design, commissioning, and acceptance testing. Read this document thoroughly before proceeding from your testbed to a more substantial deployment of liquid cooling in your site.

Owner's Project Requirements Guideline for Data Center Mechanical Controls provides a guideline for facility and operations managers to make sure that their mechanical cooling control systems are well designed upfront in the process of installing a new supercomputer. It will serve as a guiding checklist for the reader to use to inform their own generation of project-guiding documents.

Liquid Cooling Controls Team Whitepaper is a guideline for data inputs to consider when designing a liquid-cooling control system. This paper includes a general background section on the transition from air to liquid cooling.

TUE, a new energy-efficiency metric applied at ORNL's Jaguar is a paper that proposes two new metrics: ITUE (IT-power usage effectiveness), similar to PUE but "inside" the IT and TUE (total-power usage effectiveness), which combines the two for a total efficiency picture. It includes a demonstration of the method, and a case study of measurements at ORNL's Jaguar system.

A Conceptual Framework for HPC Operational Data Analytics provides a broad framework for understanding trends in Operational Data Analytics (ODA) for High-Performance Computing (HPC) facilities.

The Open Specification for Liquid Cooled Rack Transfer Fluid follows the Construction Specifications Institute (CSI) format for the transfer fluid that is used in the secondary or closed loop between the heat exchanger (e.g., CDU) and the cold plates or other heat exchangers within the rack/servers. This document provides a specification template for furnishing and installing a water-based transfer fluid and associated equipment.

7.2 ASHRAE Data Center Resources

The ASHRAE Technical Committee 9.9 is a leader in data center environmental knowledge and industry best practices. There are 14 books in the ASHRAE Datacom Series available for purchase, including *Liquid Cooling Guidelines for Datacom Equipment Centers*, along with publicly available white papers and technical bulletins.

An ASHRAE TC 9.9 white paper, *Emergence and Expansion of Liquid Cooling in Mainstream Data Centers*, indicates liquid cooling will become more mainstream due to increasing socket power (CPU/GPU/TPU) and is the source for Figure 1 that shows when a transition from air to liquid cooling based on socket power is likely. Note that ASHRAE has also recently renamed the ASHRAE water classifications to clearly define the server supply temperature range. The upper temperature limits are now included in the W number (in degrees Celsius). The classes are newly named as follows: W17 (previously W1), W27 (W2), W32 (W3), W40 (new), W45 (W4), and W+ (W5). This change to ASHRAE W classes has been included in the fifth edition of *Thermal Guidelines for Data Processing Environments* (released in 2021) and will be included in the upcoming third edition of the *Liquid Cooling Guidelines for Datacom Equipment Centers*.

ASHRAE Links

- *Water-Cooled Servers Common Designs, Components, and Processes* white paper
- *Emergence and Expansion of Liquid Cooling in Mainstream Data Centers* white paper
- ASHRAE Datacom Series, list of books available for purchase.

7.3 Open Compute Project

The Open Compute Project has a number of relevant projects dedicated to liquid cooling:

- Advanced Cooling Facilities
- Cold Plate
- Immersion
- Door Heat Exchanger
- Heat Reuse
- Hardware Management for Liquid-Cooling Workstream.
- Guidelines for Using Water-based Transfer Fluids in Single Phase Cold Plate Based Liquid Cooled Racks

7.4 The Green Grid

Available from the public library area that can be accessed with a free guest account.

- [Liquid Cooling Handbook]
- [PUE]
- [Carbon Usage Effectiveness]
- [Space Usage Effectiveness]
- [Water Usage Effectiveness.]

7.5 Case Studies

- Sandia HPC data center case study
    - *Sandia National Laboratories' Holistic Data Center Design Integrates Energy- and Water- Efficiency, Flexibility, and Resilience*
- MORE LINKS from EE HPC WG community! Please email suggestions to Natalie nbates@lbl.gov