

Water-Cooled Servers

Common Designs, Components, and Processes

White Paper Developed by

ASHRAE Technical Committee 9.9,
Mission Critical Facilities, Data Centers, Technology Spaces
and Electronic Equipment



Atlanta

© 2019 ASHRAE

1791 Tullie Circle, NE · Atlanta, GA 30329 · www.ashrae.org · All rights reserved.

This white paper was developed by
ASHRAE Technical Committee (TC) 9.9, Mission Critical Facilities, Data Centers,
Technology Spaces and Electronic Equipment.

ASHRAE is a registered trademark in the U.S. Patent and Trademark Office, owned by the American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.

ASHRAE has compiled this publication with care, but ASHRAE has not investigated, and ASHRAE expressly disclaims any duty to investigate, any product, service, process, procedure, design, or the like that may be described herein. The appearance of any technical data or editorial material in this publication does not constitute endorsement, warranty, or guaranty by ASHRAE of any product, service, process, procedure, design, or the like. ASHRAE does not warrant that the information in the publication is free of errors, and ASHRAE does not necessarily agree with any statement or opinion in this publication. The entire risk of the use of any information in this publication is assumed by the user.

No part of this publication may be reproduced without permission in writing from ASHRAE, except by a reviewer who may quote brief passages or reproduce illustrations in a review with appropriate credit, nor may any part of this publication be reproduced, stored in a retrieval system, or transmitted in any way or by any means—electronic, photocopying, recording, or other—without permission in writing from ASHRAE. Requests for permission should be submitted at www.ashrae.org/permissions.

Contents

Objective of the White Paper.	1
Introduction to Water-Cooled Servers	2
Design Considerations for CDU Implementations.	10
Design Considerations for Non-CDU Implementations	13
Water Quality.	19
Wetted Materials	23
Filtration.	25
Pressure Testing Requirements for Water-Cooled Servers	31
Fluid Couplings Used in Water-Cooled Systems	34
Flexible Hose Used in Water-Cooled Systems	39
Clarifications of the FCS and TCS Loops Described in <i>Liquid Cooling Guidelines for Datacom Equipment Centers</i> , Second Edition	44
Future Work and Additional Resources.	45
References	46

Acknowledgments

The ASHRAE TC 9.9 committee would like to thank the following people for their work and willingness to share their subject matter knowledge to further the understanding of water cooling as applied to servers deployed into data centers.

- Matt Archibald—Lenovo
- Paul Artman—Lenovo
- Tozer Bandorawalla—Intel
- Jerrod Buterbaugh—Lenovo
- Neil Chauhan—DLB Associates
- Shanjiu Chi—Huawei Technologies Company, LDT
- Bryce Cox—Oregon State
- Dustin Demetriou—IBM
- Mike Ellsworth—IBM
- Vadim Gektin—Huawei Technologies Company, LDT
- Joshua Gess—Oregon State
- David Grant—ONRL
- John Gross—J.M. Gross Engineering
- Sakamoto Hitoshi—Huawei Technologies Company, LDT
- Hugh Hudson—Daikin Applied
- Hui Jia—Huawei Technologies Company, LDT
- Devdatta Kulkarni—Intel
- Elizabeth Langer—CPC (Colder Products Company)
- Tim McCann—Hewlett Packard Enterprise
- Mike Patterson—Intel, Retired
- Roger Schmidt—IBM, Retired
- Austin Shelnutt—Dell
- Sean T. Sivapalan—Intel
- Vali Sorell—Sorell Engineering
- Chris Stack—DLB Associates
- Mark Steinke—Cray Inc.
- Jei Wei—Fujitsu
- Casey Winkel—Intel

These people invested a significant amount of their time in conference calls, writing drafts, drawing figures, and editing and reviewing text. Special thanks to Austin Shelnutt and Mark Steinke for each leading a subgroup focusing on specific sections of the white paper and to Mark Steinke for doing much of the editing in the final stages. Also thanks to Roger Schmidt for his leadership in driving the creation of this white paper.

Objective of the White Paper

With more water-cooled IT products arriving in the marketplace, ASHRAE TC 9.9 felt the need to outline some of the common processes, parts, and materials for focus in use for future water-cooled designs. Some parts in a water-cooled IT system will be specific to the product design, such as cold plates, manifolds, arrangement of piping, pumps, valves, and so on, but others such as quick connects, hoses, hose connections, materials, and water chemistry fall more into the category of common parts that can be used by all current and potential manufacturers of water-cooled IT equipment. This white paper is an attempt to provide and make available those items that could be classified as *common*. The material published in this white paper compliments the materials published in *Liquid Cooling Guidelines for Datacom Equipment Centers*, second edition (ASHRAE 2014).

Beyond the objective stated above, this white paper also corrects misunderstandings in the latest edition of *Liquid Cooling Guidelines for Datacom Equipment Centers*, and provides guidance for avoiding common mistakes based upon the book's content. We also expect that much of the content in this white paper will be incorporated into a future third edition.

Introduction to Water-Cooled Servers

When cooling data centers and the heat-generating IT equipment (ITE) contained within, many feel that liquid cooling is the way of the future. While liquid cooling has been on the verge of disrupting the market for a long time, only modest progress has been made. It was common opinion among industry experts that the first to market with a liquid-cooled solution would lose. This may seem counter-intuitive, but makes sense given that ITE manufacturers have been hesitant to invest heavily in liquid cooling technology because the market for that type of equipment has been sparse. In turn, the market has been reluctant to embrace liquid cooling when air cooling continues to be the predominant cooling medium for servers in the marketplace and where liquid cooling is perceived as a niche market.

ITE manufacturers at both the server and component (i.e., processor) level have extended air cooling capability by designing using improved packaging materials for enhanced heat transfer, higher server airflow, and increased ΔT s across the servers. Coupled with cooling efficiency best practices in the data center, such as the use of hot aisle/cold aisle containment, more useful cooling can be delivered with the same airflow.

These incremental advances have breathed new life into a technology previously thought to be approaching its practical limitations. Not long ago, owners and designers believed cabinets approaching 20 to 30 kW load were at the ceiling of air-cooling capability for maintaining component temperatures within industry specifications. More recently, some air-cooled server products have achieved cabinet heat loads upwards of 40 to 50 kW thanks to advances in air moving technology within the server and data center. The ability to move additional airflow and enabling higher rack heat loads is not without cost, however; Figure 1 illustrates how increasing air cooling density not only increases the amount of power used to move that air, but also reduces cooling efficiency. The figure is representative of one server vendor's cooling power efficiency (the amount of fan power divided by the total system power) across a range of different power density servers. Figure 1 also includes one of this same vendor's liquid-cooled servers to illustrate how liquid cooling can break the trend of ramping cooling power costs associated with higher powers of ITE within a rack.

It is not obvious when it will not be possible to solely rely on air cooling for a server. At some point, the air-cooled devices will reach their limits—either through physics limitations or exponential increases in cooling power con-

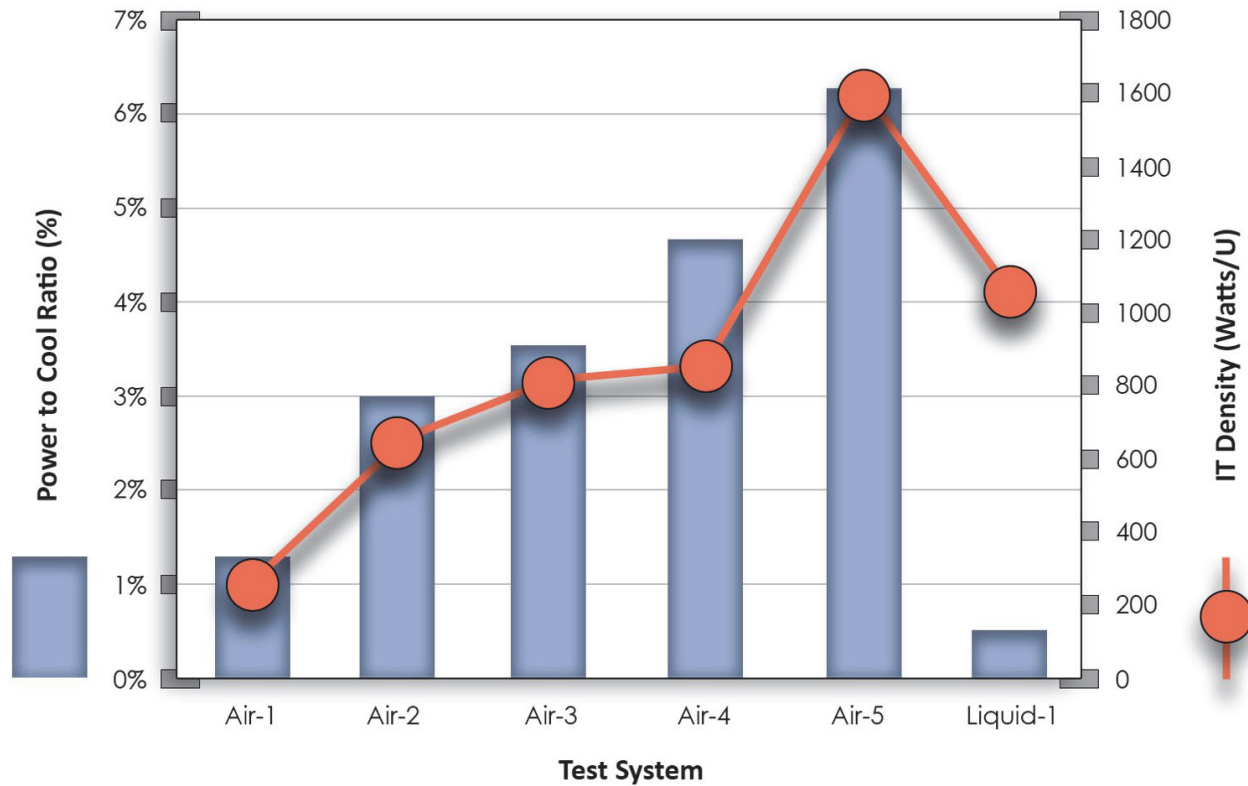


Figure 1 Power-to-cool ratio for various servers.

sumed—and the ITE manufacturers will need to adjust their designs to accommodate liquid cooling. These modifications will require the elimination of air-cooled heat sinks replaced by cold plates mounted to processors and the possible elimination of some and possibly all server fans. They may also require a form factor change, perhaps changing some configurations to allow boards to be mounted and packaged differently with respect to the motherboard and other components. Other modifications that cannot be foreseen now may also be required.

The increasing heat density of modern electronics is stretching the ability of air to adequately and efficiently cool the electronic components within servers as well as the datacom facilities that house these servers. To meet this challenge, water is now being used at the rack or board level. Water has more than 3500 times the heat capacity of air. The ability of water to carry much larger amounts of heat per volume or mass offers tremendous advantages.

Water-cooled servers have been around for more than a half a century. In the past, high-performance mainframes were often water-cooled and supplied as such by the ITE manufacturer. These systems were often highly customized and costly. Increased standardization of liquid-cooled designs will help expand their use by minimizing piping/connection concerns, increasing volumes, and reducing overall cost.

Water-Cooled Product Trends

ASHRAE TC 9.9 recently published *IT Equipment Power Trends*, third edition (ASHRAE 2018), which documents power trends of air-cooled equipment. Those trends are for volume servers and are presented for seven application workloads. Trends for two-unit (2U), two-socket servers from 2015 to 2025 are shown contrasted with power trends reported in the previous, second edition of *IT Equipment Power Trends* in Figure 2. The seven categories for workloads are scientific, visualization and audio, analytics, storage, business processing, cloud and communications. The third edition trends display bands where the band thickness is the variation of measured rack powers for racks of 2U, two-socket servers based on measured rack powers for each workload type. As shown in this figure and as expected, the category with the highest maximum power trend is the scientific computing equipment. This trend is approximately twice as high as the next nearest category (visualization and audio). The geometry for the rack powers displayed are for a 42 U, 19-in. (48.26 cm) standard rack with a depth of 43 in. (109.22 cm).

Water-cooled servers are not nearly as prevalent as air-cooled servers. A method must be established that can best display some of the trends of water-cooled ITE. Because there is no known database for this data, original equipment manufacturer (OEM) server companies were contacted for this data. The data offered is the maximum achievable power for the available water-cooled designs. Although this is not measured data, it is shown here because that is what was made available. Future updates of this data may encourage OEM companies to

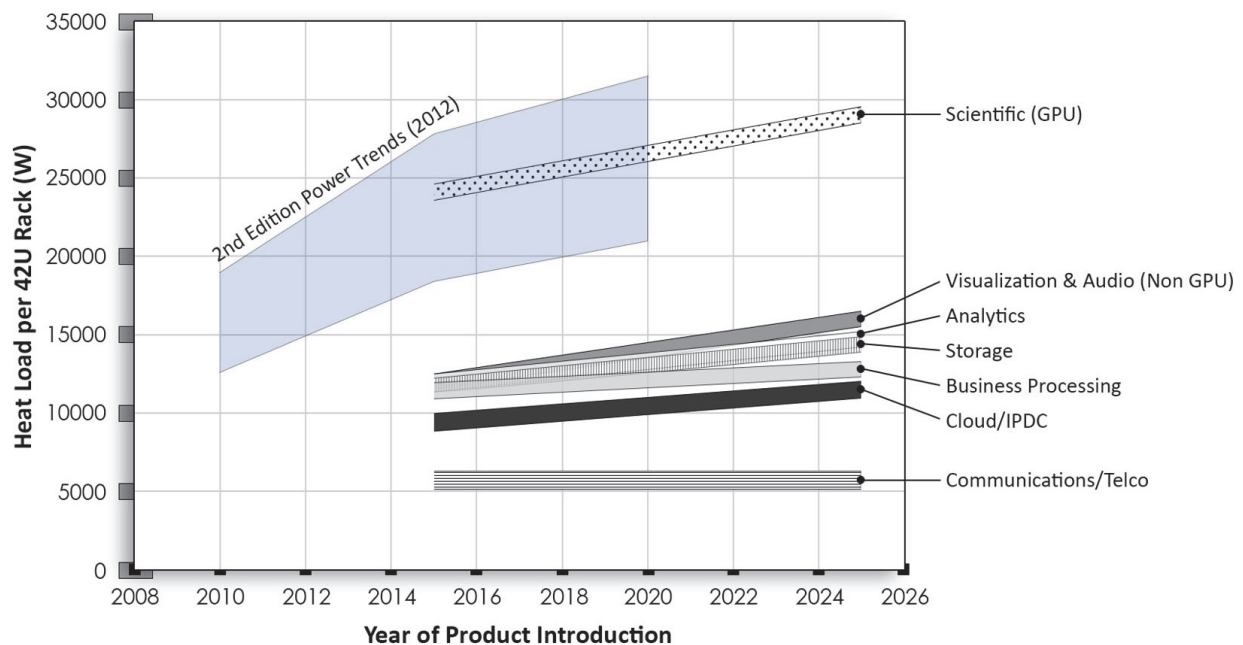


Figure 2 Air-cooled 2U two-socket server power trends (ASHRAE 2018).

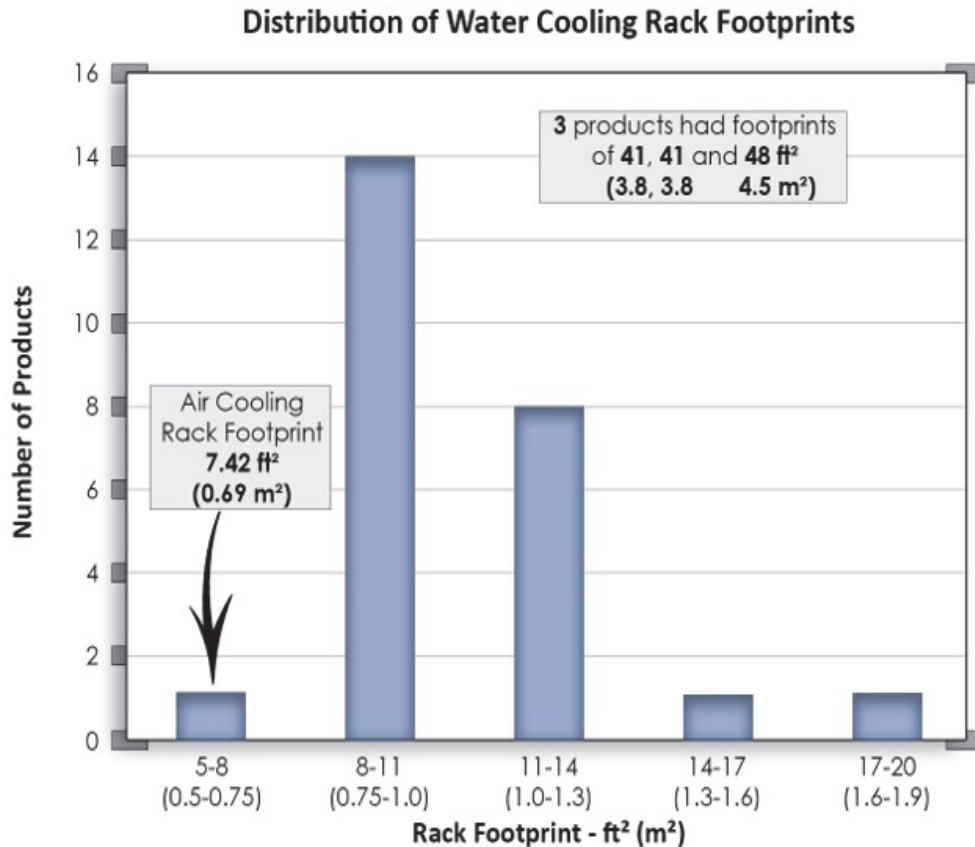


Figure 3 Distribution of water-cooled server rack footprint areas (ft²).

provide more data and representative measured data to better compare to the air-cooled power trend data. In addition to displaying liquid-cooled server power trend data, it will also be informative to show the number of water-cooled products announced annually for several years and whether the number of products is increasing.

The other obvious difference between the air-cooled and water-cooled racks is that there is no standard rack for water-cooled equipment. In fact, there is a very wide variation of footprint for water-cooled server racks. Whereas the air-cooled rack power trends used the standard 19-in. rack footprint of 7.42 ft² (0.72 m²) for all air-cooled power trends, the water-cooled racks varied in footprint from 6.78 to 48.20 ft² (0.63 to 4.48 m²). The data for the footprint of water-cooled racks is displayed along with the air-cooled rack footprint used for the air-cooled trends in Figure 3. This large variation in rack footprint drove the creation of the water-cooled trends based on power/rack footprint rather than just power as in the air-cooled server rack trends shown in Figure 2. Although the height of the water-cooled racks should be fairly constant, this too is varied. However, most designs were racks with heights of 42 U. No accommodation in the graphical display was made for variations in rack height.

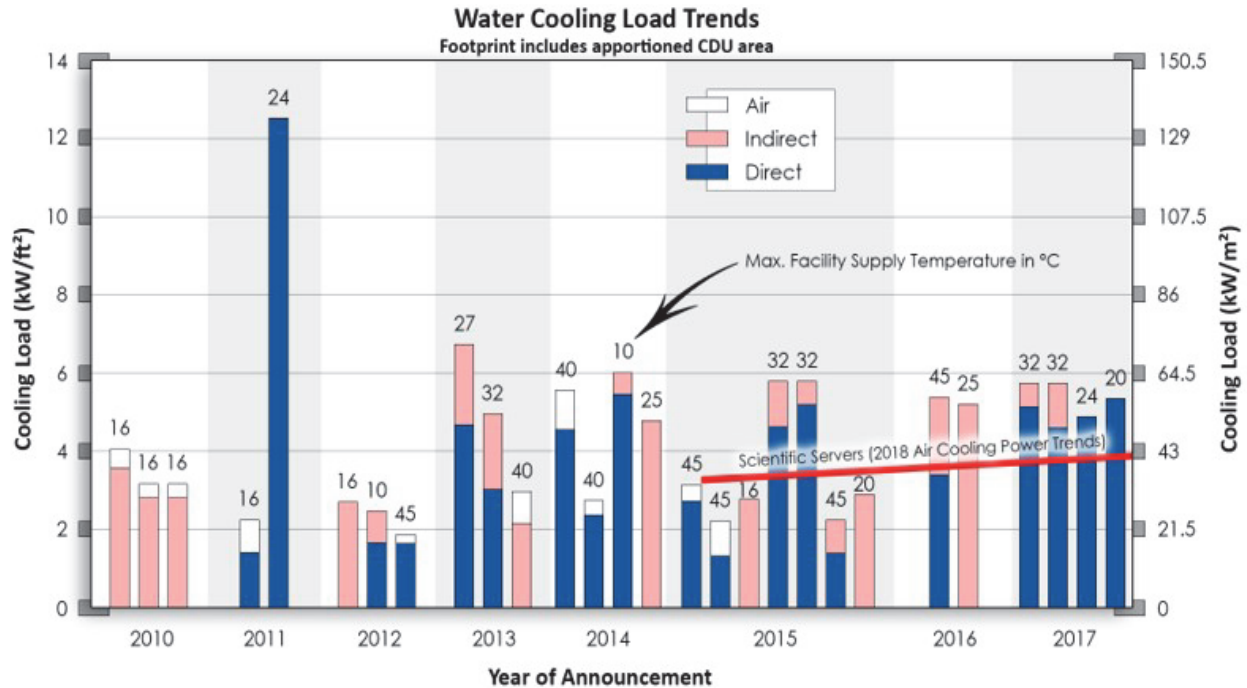


Figure 4 Water cooling load trends from 2010 to 2017.

Data obtained through personal communication with HPE, SGI, Fujitsu, IBM, Dell, Huawei, and Cray.

As shown in Figure 3, the average area of the 28 water-cooled server products is 14.63 ft² (1.36 m²) with a wide spread of 6.78 to 48.20 ft² (0.63 m² to 4.48 m²). This figure clearly shows the contrast between the rack area used for the air-cooled servers and the water-cooled server racks. Only one of the water-cooled products falls within the same category of footprint similar to the rack areas used to display the air-cooled power trends in Figure 2.

The data from the 28 water-cooled products announced from 2010 to 2017 is shown in Figure 4. As explained previously, the vertical axis is the server power divided by the rack footprint. Variations in height are not accommodated in this figure, but it is noted that there was not a large variation of rack heights. One other caveat in the denominator for the area is that if the water-cooled server rack requires a coolant distribution unit (CDU) this area is apportioned to the rack area. For example, if a CDU supports 4 server racks, 1/4 of the area of the CDU is added to the rack area to make up the denominator for power/rack footprint. The average percent CDU area compared to the supported rack space is 16% with the range spanning 13% to 25%. All the water-cooled server rack data is collected from personal communications with seven server OEMs—HPE, SGI, Fujitsu, IBM, Dell, Huawei, and Cray.

The number of bars shown for each year in Figure 4 do not show any overriding trend, the average being a little over three water-cooled products per year averaged over the eight years. Only one year (2015) showed any substantial growth from the average (seven products), but this was not sustained as the following year

had only two products and the next year four. At the very least, the data shows no growth trend in the number of products designed by IT OEMs for water cooling over time. Each bar displays the total maximum heat load dissipated by the water-cooled server rack. The amount of power dissipated is categorized into three areas—air-cooled, direct water-cooled, and indirect water-cooled. These are defined as follows:

- **Air cooled**—power that is transferred directly to the room air and cooled via traditional data center cooling
- **Indirect water-cooled**—power that is transferred indirectly to water through an air-to-water heat exchanger located within or on the rack (could be a rear door heat exchanger)
- **Direct water-cooled**—power that is transferred directly to an attached heat transfer component such as a cold plate

The total heat load is only through indirect water cooling in four of the 28 racks (that is, the use of a rear door heat exchanger). In three of the 28 racks the total heat load is achieved by only direct water cooling, with the remainder of the products using a combination of direct and indirect cooling and in some cases some heat being exhausted from the rack to the ambient air. But for those that exhaust some heat to the ambient it is clear from the graph that this was minimized, leading one to conclude that the designs are focused on driving as much heat as possible to water, either direct or indirect.

The air-cooled power trend for scientific equipment from Figure 2 is shown in Figure 4 to contrast the power capabilities of water cooling with that of air. Because the air-cooled power trends within *IT Equipment Power Trends* (ASHRAE 2018) were displayed as power only and not power/footprint, the air-cooled power trend graph in Figure 2 needs to be converted to power/area as shown in Figure 4. The footprint of the air-cooled racks shown was 24×43 in. (60.96×109.22 cm).

The comparison of the water-cooled products to the scientific air-cooled products from 2015 to 2017 shows that some of the water-cooled products were equal to or slightly lower power than the air-cooled products.

This begs the question as to why these are designed as water-cooled products. It could be that some of the components, like processors, are of high enough power that they require water cooling. But for those that do exceed the air-cooled scientific servers, the power is in general 50% higher than for the water-cooled products. Another benefit of adopting liquid cooling is improved data center energy efficiency.

Lastly, the tops of each bar in Figure 4 display the maximum facility water temperature that can be employed to cool the server rack. The first seven products shown have maximum facility water temperatures of 50°F to 75.2°F (10°C to 24°C). Of the remaining 21 products, eight have maximum facility water temperatures of greater than or equal to 104°F (40°C), indicating a clear focus on the use of economizers and the possible use of the heated exhaust water from the servers to heat local buildings in the winter.

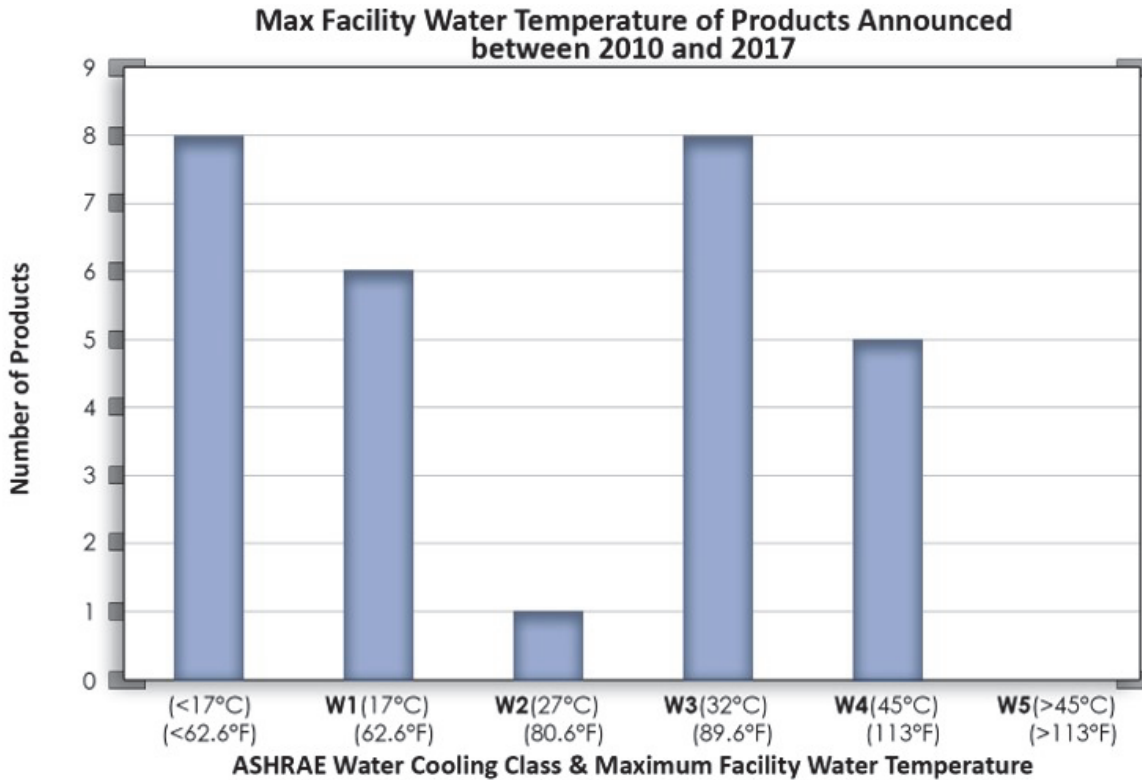


Figure 5 Maximum facility water temperature of products shipped.

A summary of the maximum facility water temperature allowed compared to the ASHRAE water cooling classes for all the water-cooled products announced over the last eight years is shown in Figure 5.

If a product can support a facility water temperature up the maximum allowed in that class then it is placed in that category. For example, if the maximum facility water temperature for an announced product is 75.2°F (24°C) then it is placed in the W1 class.

Water-Cooled Server Data Center Implementations

Several implementations of water cooling could be deployed, such as the water coolant removing a large percentage of the waste heat via a rear door heat exchanger. Another is with the water coolant passing through cold plates attached to processor modules and other high powered electronic components within the rack. The most common method to implement these water cooling designs is to deploy a coolant distribution unit (CDU) external to the datacom rack.

Another method to deliver water cooling to the ITE, although not as common is the non-CDU liquid cooling system. In this case, the facility water is provided directly to the ITE rack for cooling the components. These two-different implementations and some design considerations for each will be described in more detail in following sections.

Other Design Implementations

In addition to discussing the type of CDU implementations, the following sections will discuss several other design considerations needed for water-cooled servers. Some of these elements are common to fluidic design practices in general, such as the importance of approved water chemistry and understanding wetted materials. Some of the design considerations are unique, such as very small passages of cold plates and service requirements for hardware failures. Each section will address one aspect of water-cooled server design and describe some of the areas of concern and focus for the ITE designer.

Design Considerations for CDU Implementations

The most common form of liquid cooling implementation within the data center utilizes a CDU to separate the facility water system (FWS) from the technology cooling system (TCS). By using liquid-to-liquid heat exchangers, CDUs can reject heat from ITE without exposing the sensitive IT cooling components to the typically less regulated FWS water. CDUs may be located within an IT rack, providing TCS liquid distribution to equipment within a single rack, or installed as a floor-standing external unit that distributes TCS liquid to a plurality of racks. Figure 6 shows an example of a CDU-based liquid cooling installation and specifically references each of the three liquid loop types discussed in this paper.

The CDU plays an important role in data center liquid cooling by providing the following functions:

- Transferring heat from TCS to FWS
- Circulating TCS coolant

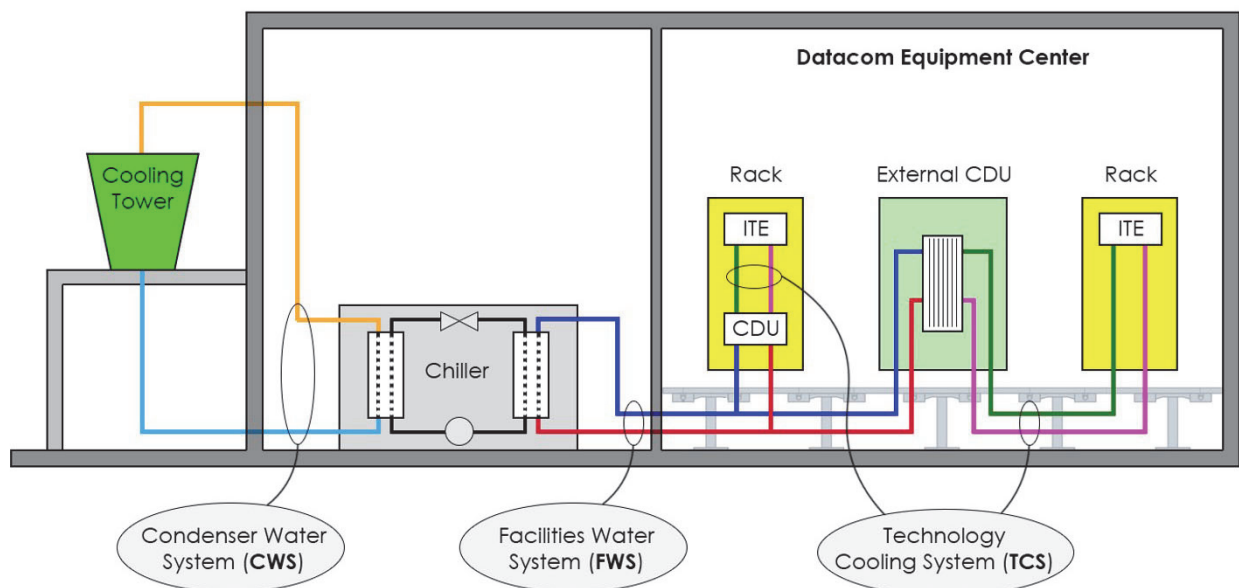


Figure 6 CDU liquid cooling system within a data center.

- Allowing for a TCS coolant other than water (e.g., refrigerant or engineered fluid)
- Preventing condensation within the ITE by regulating TCS water above room dew point
- Establishing and maintaining a coolant quality and chemistry different from that associated with the FWS and more suitable to the TCS
- Supplying flexible coolant temperature to ITE

Water Quality for a CDU Implementation

By using a liquid-to-liquid heat exchanger as the transfer medium between TCS and FWS, the CDU is able to provide a critical benefit to ITE: separation of water quality control domains. Yet, while the CDU is able to shield ITE cooling components from the FWS water, the CDU itself must be designed to operate harmoniously with both FWS and TCS water loops. The key differences between TCS and FWS water quality are articulated in later section titled Water Quality. It is important to understand the differences in both water types so that proper application of their respective guidelines is present in the design and deployment of a CDU.

The primary side of the CDU's internal liquid-to-liquid heat exchanger is the section of the unit that interfaces directly to the facility water supply. Flow control valves, piping, heat exchanger plates, filters, and all other wetted equipment attached to this section must be chemically compatible with the constituents of the FWS loop.

The secondary side of the CDU's internal liquid-to-liquid heat exchanger is the section of the unit that interfaces directly to the TCS. Flow control valves, piping, heat exchanger plates, filters, and all other wetted equipment attached to this section must be chemically compatible with the constituents of the TCS loop.

Another important design consideration for a CDU-based implementation is the change in the TCS water quality state over the lifetime of the deployment. TCS fluid chemistry is often proprietary in nature as selected/designed by the IT manufacturer. As such, chemical compatibility with the components in the TCS loop tends to be quite good and the fluid shows good stability over time within a completely closed environment. However, issues can arise with water quality within the closed TCS loop. Foreign contaminants can be introduced over time as new ITE elements are introduced into the system. A new server node that is poorly purged/cleaned from the factory may expose an existing TCS loop to new contaminants upon installation.

Larger, external CDUs may provide a better return on investment for scaling deployments of liquid servers within a data center by enabling future liquid expansion, but the risk of foreign contamination and pipe scale increases with the number of devices serviced and because, over time, branches of liquid conduits can become stagnant. Because the TCS loop is isolated from the FWS, there is often some ambiguity in ownership of its administration between facility teams and the IT group. For this reason, it is necessary to have an established TCS water quality control owner; ensuring that onboarding new equipment, servicing exist-

ing equipment, and general contaminant permeation over time does not degrade the cooling capabilities of the ITE or the CDU itself.

Filtration for a CDU Implementation

While filtration is addressed in its own section later in the paper, it is important to note the specific role of filtration within a CDU environment. Within the CDU, the liquid-to-liquid heat exchanger is typically a flat plate construction with nominal plate spacing ranging from 0.079 to 0.315 in. (2 to 8 mm). The primary side of the CDU's liquid-to-liquid heat exchanger operates using FWS water, which can have a variable degree of filtration applied to it upstream of the CDU depending on the customer. Additionally, the FWS lines may sit stagnant between facility mechanical commissioning and the deployment of a CDU that attaches to the FWS water. Endpoint filtration at the CDU (on the supply side) is necessary in order to prevent particulate build up within the plate heat exchanger. The exact plate spacing for a given CDU heat exchanger should be used to design filtration in accordance with the recommendations of the Filtration section of this paper. Finally, filtration is also recommended on the secondary (TCS) side of the CDU heat exchanger as well to mitigate foreign particulate infiltration caused by servicing or commissioning equipment, along with general scale formation.

Temperature Control

Regulating the supply temperature of the TCS loop is one of the primary functions of the CDU. The supply water temperature provided by the CDU may be controlled by using any combination of bypass loops, proportional control valves, and/or variable speed pumps as a response to variation in FWS water temperatures and the heat load of the ITE. Typically, the TCS loop is further defined by the highest allowable temperature to support liquid cooling of the ITE, referred to as the *upper bound*. In addition, the lower bound temperature is defined for the TCS loop. A good definition for each limiting temperature is as follows.

- **TCS upper bound**—The maximum TCS water supply temperature that can be sustained at a given flow rate to provide sufficient cooling capability to the attached ITE. Temperatures beyond this limit will require additional flow rate and/or may cause ITE to exceed design temperature.
- **TCS lower bound**—The minimum TCS water supply temperature that may safely be provided to the ITE. This limit may be dynamic in nature and depend upon the dew point within the cooling space. The CDU is responsible for monitoring the ambient dew point and elevating the secondary water loop (TCS) supply temperature to at least 3.6°F (2°C) above the room dew point to prevent condensation. This provides a buffer and accounts for some sensor uncertainty.

Design Considerations for Non-CDU Implementations

Today, most implementations of liquid cooling for datacom equipment employ a discrete FWS demarcation point by way of a CDU. There are some scenarios, however, where coupling ITE directly to FWS can be advantageous. As previously discussed in this paper, the CDU provides many valuable benefits that need to be considered when planning a non-CDU implementation:

- Condensation prevention
- Water quality isolation
- Flexible coolant selection
- Flexible coolant temperature supplied to ITE
- Operating pressure reduced for ITE equipment

Figure 7 shows an example of data center liquid cooling system that delivers facility water directly to the ITE equipment. In this diagram, ITE refers to *information technology equipment* (often server hardware), RFU refers to a *rack filtration unit*, and FFU refers to a *facility filtration unit*.

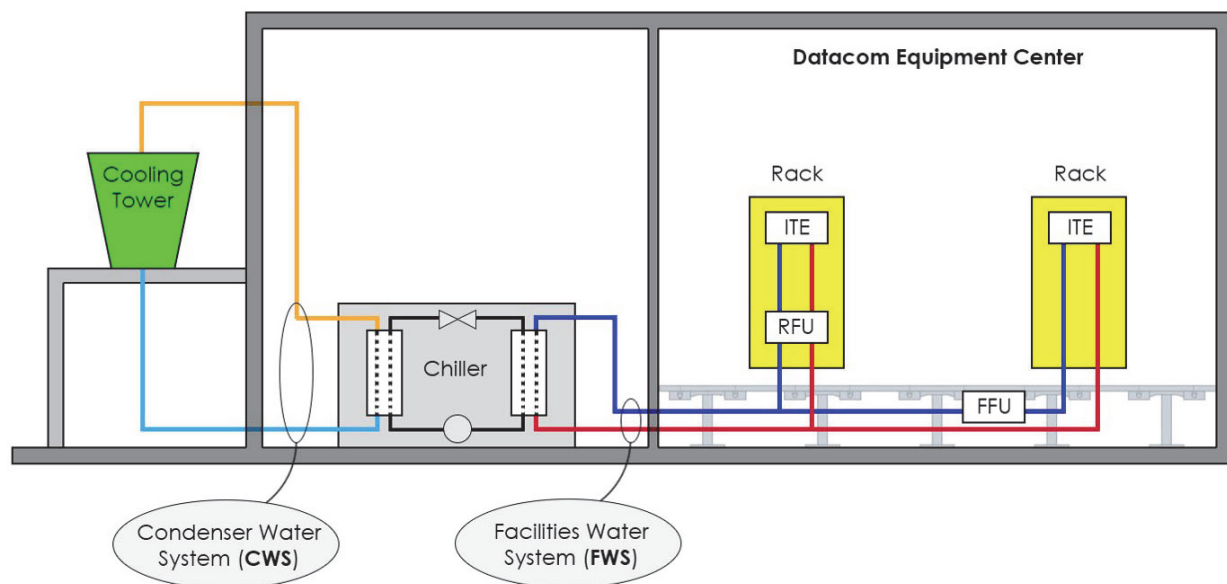


Figure 7 Non-CDU liquid cooling system within a data center.

Some of the key advantages of a non-CDU implementation include the following:

- Reduced data center floor or IT rack space consumption
- Closer coupling of ITE heat sources with final heat dissipation medium
- Shared pumping redundancy
- Improved liquid cooling efficiency

While the benefits of a non-CDU implementation are fairly attractive, the aforementioned CDU functions are still necessary. In a non-CDU implementation, however, these CDU functions must be accounted for elsewhere in the overall system (either in the ITE rack itself or in a tightly controlled FWS).

Water Quality for Non-CDU Implementations

With a CDU in place to reject ITE heat load into a FWS (primary side) water loop, the water quality between the ITE and the CDU is much easier to control and maintain. This enables manufacturers of ITE liquid cooling devices to employ a wide selection of constituent materials that may otherwise be sensitive to elements found in the FWS. In most cases, the result is that a designer of liquid cooling devices intended to operate with FWS directly brought into the ITE must conform to the same material compatibility limitations as other components typically found in the FWS such as computer room air conditioning/handler units, pipes, strainers, fittings, valves, pumps, heat exchangers, and CDUs themselves.

Corrosion and scaling of materials used in non-CDU liquid-cooled ITE can be accelerated through the use of wetted materials that are chemically incompatible with various waterborne constituents. For this reason, wetted materials should be selected based upon those that are known to be compatible in the same water loop under a given definition of water quality standards. ASHRAE provides a valuable discourse on the best practices and principles for wetted materials in equipment that uses FWS water in Chapter 5 of *Liquid Cooling Guidelines for Datacom Equipment Centers*, Second Edition (ASHRAE 2014). In addition to detailed explanation of various failure modes and material choices, Chapter 5 presents a tabulated reference for water quality standards that apply to equipment operating on the FWS.

In non-CDU implementations, the facility design team must work closely with the ITE manufacturer to assure the materials and water quality selected are appropriate for both the facility and the ITE.

Filtration Specifics for Non-CDU Implementations

As with CDU implementations, non-CDU water-cooled ITE requires end-point filtration on the supply side of the FWS. In Figure 7, the RFU and FFU are presented as a rack-level filtration unit and a facility-level filtration unit aimed at providing end-point filtration for devices serviced by the facility water supply. These devices should be positioned at the junction of FWS and ITE where the RFU would reside on the ITE-side of the junction and the FFU would reside on

the facility-side of the junction. It is not required that both RFU and FFU be present, but it is important that one or the other is implemented.

Filtration Sizing

Selection of an RFU or FFU filter particulate size is a balance between cooling equipment reliability, liquid flow impedance, facility filtration, and service frequency. As will be discussed in Section 7, the maximum allowable particulate size for a given application is dictated by the smallest flow path in the system. Particulate that is of equal or greater dimensional magnitude as part of the flow path will easily plug or impede that portion of the flow path. It is understood that interior flow paths may shrink over time because of various forms of scale and surface corrosion. Furthermore, multiple solid particles may coalesce in the fluid stream, creating partial or complete blockages of specific flow paths.

Today, the most commonly used liquid heat transfer apparatus in electronics cooling is the microchannel cold plate, shown in Figure 8. These devices maximize use of the caloric density (specific heat) of liquid coolants by passing them through channels that are highly effective as an attribute of the very small passages. Through advanced manufacturing methods, typical fin spacing on these liquid heat sinks is often less than a few tenths of a millimeter.

Because ITE liquid cooling devices typically have very narrow dimensions across the heat transfer surfaces, it is imperative for there to be appropriate filtration planning between data center operators and the ITE liquid cooling solution provider. Decreasing the size of fin spacing in a cold plate design typically results in superior heat transfer performance but, often, at the cost of static pressure drop across the inlet and exit ports of the cold plate. Figure 9 shows a graphical illustration of this parametric behavior.

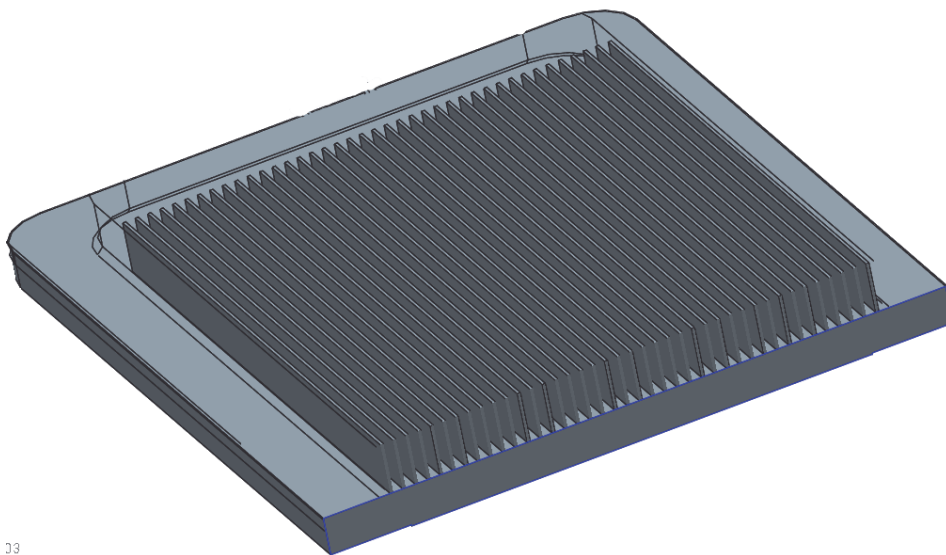


Figure 8 Cut-plane view of a microchannel cold plate.

Because a non-CDU implementation of ITE liquid cooling must operate reliably using FWS, cold plate fin spacing will necessarily be larger than the cold plates that operate in a tightly controlled TCS loop as part of a CDU implementation. Larger fin spacing typically results in an increase in cold plate thermal resistance (as shown in Figure 9), but the tradeoff can be recovered in the elimination of approach temperature across a CDU heat exchanger.

Striking a balance between diminished thermal performance of larger fin spacing and the level of filtration that is economically suitable for most data centers, nominal fin spacing for a cold plate intended to be used in a non-CDU implementation is suggested to be between 0.079 to 0.315 in. (0.7 and 1.0 mm). In line with this recommendation and the filtration sizing guidance in the Water Quality section, it is suggested that RFU or FFU filtration be sized at 100 μm or smaller to prevent solid particulate from building up within the cold plate structure and impeding the flow of vital cooling fluid.

Pressure Testing for Non-CDU Implementations

Pressure testing of liquid-handling components used within the data center is extremely important to ensure cooling liquid leaks do not create environmental

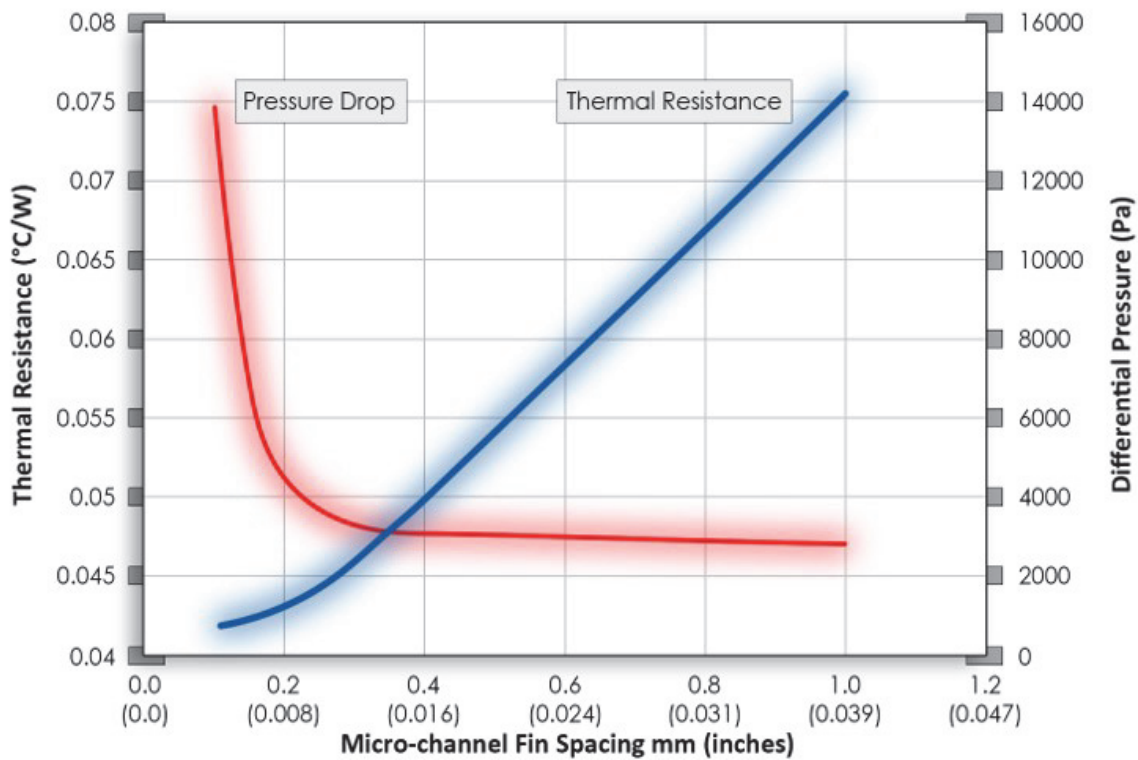


Figure 9 General illustration of the effect of fin spacing on microchannel cold plate performance.

and/or human safety hazards or capital loss due to equipment damage. Equipment designers must consider the type of liquid loop that each device will be connected to. Every device should be designed and tested to operate safely in one of the three water loop types—CWS, FWS, and TCS.

Each of these three loop types will have different characteristics for water chemistry, filtration, and operating pressure requirements. Equipment designed and tested to operate on a TCS loop may not be suitable for use with an FWS loop because of its construction and choice of materials.

Liquid cooling loops designed around non-CDU implementations must ensure all components within the assembly are designed with the static and dynamic pressure limitations of the FWS loop. The pressure considerations of a non-CDU ITE cooling loop will drive significant design differences from the ITE developed to exist on a TCS loop because the pressure conditions are significantly higher and increasingly dynamic in nature on the FWS loop.

Temperature Control for Non-CDU Implementation

One of the critical functions a CDU provides to liquid cooled ITE is the ability to influence the temperature of the liquid coolant delivered to each end device. Condensation on the surface of a cold plate residing within a server can be catastrophic. Many FWS provide chilled water that may be at a lower temperature than the air dew point in the data center. Where a CDU has the ability to elevate secondary coolant temperature (either through flow rate reduction, bypass loops, or auxiliary heat input) above the dew point, a non-CDU implementation of liquid cooling must carefully plan for this scenario.

A data center that is planning to implement non-CDU liquid cooling should have the ability to regulate FWS temperatures such that they are always maintained at least 3.6°F (2°C) above the data center room ambient dew point. This provides a buffer and accounts for some sensor uncertainty. In addition to the data center providing FWS water temperature that is always above the dew point of the data center, there are other recommended condensation mitigation techniques that can and should be employed within the liquid cooling solution of the ITE:

- Collection of telemetry:
 - Supply water temperature
 - Server inlet air dew point
- Methods of control:
 - Proportional control valves to slow the flow of incoming water
 - Binary control valves to halt the flow of incoming water

Furthermore, ITE liquid cooling equipment should be connected to the facility's building management system to send out alarms during excursions when the supply water temperature dips within 3.6°F (2°C) of the incoming air dew point. Other temperature control techniques employed in non-CDU liquid cooling include, but are not limited to, the following:

- Well-insulated/sealed pipe and cold plate walls to shield moist air from the liquid component surfaces
- Upstream liquid-to-air heat exchangers (i.e., rear door heat exchangers) to elevate incoming water temperatures
- Bypass valves to reduce the flow rate of incoming coolant

Water Quality

The main issue in the required water quality for ITE water cooling systems is a misapplication of the water quality recommendations in the second edition of *Liquid Cooling Guidelines for Datacom Equipment Centers* (ASHRAE 2014). It must be understood that there are two different loops described in the book. Chapter 5 deals with the FWS and Chapter 6 deals with the TCS. At the water quality level, these two water streams are usually physically separated by the CDU. They are often additionally separated by design and operational ownership. The FWS system is generally owned by the building owner and the facilities staff. The TCS loop is more likely to have some combination of a CDU vendor, IT system provider, or architecture and engineering firm involved in its specific use and design.

The FWS system is often a campus- or site-wide building system, while the TCS loop is data-center-specific and most often associated with a specific set of IT hardware.

The requirements for the TCS loop are more stringent than for the FWS loop. The FWS loop typically serves a robust heat exchanger with relatively large flow passages. Very often this heat exchanger is a flat plate type in a CDU. The water quality required in the heat exchanger can be met with FWS water. The TCS loop serves the ITE and provides flow for cold plates removing the heat from the electronic components. These cold plates and internal plumbing structures are far more sensitive to the potential water quality issues described in the second edition, such as corrosion, scaling, fouling, and microbial challenges. Because of this, the TCS water quality guidelines (Chapter 6, ASHRAE 2014) are much stricter than that outlined for the FWS (Chapter 5, ASHRAE 2014). These water quality guidelines are outlined in Table 1. Unfortunately, both have been misapplied. All too often this has been a simple error of using the wrong table for the system under consideration.

When a TCS type requirement is incorrectly imposed upon a facility owner and their existing building level cooling loop, there will be few if any buildings able to meet this requirement. The building owner will invariably and rightfully push back on this requirement as it is unnecessary and would come with a very high cost. Similarly, when an FWS water quality specification is applied to the IT side of the cooling loops it puts the ITE at significant risk due the much lower water quality in that loop.

The most important first step to a successful liquid cooling implementation is to apply the FWS guideline (Chapter 5) to the facility side and the TCS guideline (Chapter 6) to the ITE side.

Table 1 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	<100 ppm	<10 ppm
Chloride	<50 ppm	<5 ppm
Bacteria	<1000 CFUs/mL	<100 CFUs/mL
Total hardness (as CaCO ₃)	<200 ppm	<20 ppm
Conductivity	—	0.2 to 20 micromho/cm
Total suspended solids	—	<3 ppm
Residue after evaporation	<500 ppm	<50 ppm
Turbidity	<20 NTU (Nephelometric)	<20 NTU (Nephelometric)

FWS-Water-Quality-Specific Issues

The FWS water quality guidelines are only guidelines. There are many FWS systems in use today that operate outside these guidelines and work without problems. An important part of water quality is water treatment chemistry (i.e., pH, corrosion inhibitors, and so on). As every site has different feed or make-up water quality available, the specifics of the treatment plan must be site specific. The challenges of each site FWS loop in terms of corrosion, fouling, scaling, and microbial issues associated with the specific water chemistry and composition of the incoming water needs to be comprehended. The facility should have water treatment expertise on staff, or the facility should partner with a water treatment service provider to provide expertise on water chemistry and treatment. Ongoing water monitoring and water management is typically a part of the treatment regimen. Involving the water treatment specialist in discussions about the materials of construction, temperature, pressure, and operating requirements of the FWS and/or TCS will go a long way toward successful implementation.

The current FWS water quality guideline does not have any guidance for particulate removal. It is recommended that some level of filtration or screening be provided in the FWS loop. A general guideline for FWS systems is a filter or strainer of 35 to 40 mesh (or 500 to 400 μm). The value of 35 mesh is considered adequate for the plate and frame heat exchangers most commonly found in a

CDU. Other fluid or heat transfer components with narrower flow channels may require additional review. This should be further discussed with the water treatment vendor and the CDU provider, the key design feature being that the mesh size must be fine enough to preclude anything that could foul the downstream hardware (e.g., heat exchangers and so on) from being exposed to such particulate.

TCS-Loop-Water-Quality-Specific Issues

The TCS loop water quality requirements (Table 1) require a higher level of water quality than the FWS loop can generally provide. This guideline and the reasonably tight ionic limitations (i.e., when measured as conductivity) have caused challenges in the field. A TCS system may meet the conductivity requirement initially and then exceed the requirement over time. This could be caused by ionic contaminants entering the system (including CO₂ permeation), or ions leaching from materials of construction which requires remediation, or it could be due to the addition of a corrosion inhibitor, biocide, or pH buffer; all of which should be acceptable. Successful system maintenance requires understanding the cause and taking the appropriate action. The value for conductivity is most important for the quality of make-up water being added to a TCS loop. In addition, conductivity should be charted over time and the root cause of any changes or deviations should be understood.

Filtration is also not covered for the TCS loop in the liquid cooling book second edition (ASHRAE 2014) and should be addressed in the forthcoming third edition. As the TCS loop feeds water to the IT component level liquid heat sinks (i.e., cold plates) that generally have finer dimensions than the FWS, filtration here is even more important. Cold plate cooling performance improves with smaller and smaller channel dimensions. Filtration is needed to preclude fouling of these tight channels. It is recommended the TCS loop filtration have an absolute filter rating of 7 to 10 times smaller than the finest channel dimension in the IT cooling equipment. An absolute filter rating, described later in this paper, is required versus a nominal filter rating. The definitions of *nominal* and the performance of such filters are too loose to risk fouling the IT cooling hardware.

The design of this filtration is also critical for operation of the IT hardware. These filters will need to be maintained, typically through replacement. The anticipated frequency must be low enough to not have an impact, or preferentially an $N+1$ - or $2N$ -type redundant design should be employed. Note also that the design flow and pressure drop of the N configuration should be taken at the “dirty” or loaded state, where the pressure drop across the filters can be significantly higher than clean pressure drop.

Monitoring and Maintenance

The TCS and the FWS loop need to have a monitoring and maintenance plan established. The FWS system will generally be maintained by the building owner and is largely outside the scope of this paper. The TCS loop plan must be under-

stood and can be the responsibility of the IT owner, IT provider, CDU provider, or a third party.

Frequency of monitoring can be affected by many factors, so a single, one-size-fits-all recommendation is not practical. Instead, it is recommended that all parties involved agree to the same monitoring frequency. One methodology to establish monitoring frequency is to monitor more frequently in the initial operation of the system and then to reduce that frequency through the proper application of statistical process control methodologies.

Perhaps as important is that a plan needs to be developed for actions and responses based upon values that are outside of the recommended water quality ranges. For example, if the bacteria count is too high then a possible next step is to flush the system, add more biocide, or repeat the sample (as bacteria testing can often produce false high positives). If high bacteria count is confirmed, then take action. The same process is needed for pH, conductivity, and ionic contaminants. This plan should be developed and agreed to by all stakeholders.

Another monitoring decision is that of on-line monitoring versus periodic samples. On-line monitoring will provide the best visibility to system health and stability, but also at the highest cost. The balance between cost and risk around lower water quality visibility must be weighed by the system owner and the choice made based on that analysis. Most quality parameters can be measured on-line. However, bacteria cannot, and as such a periodic water sample must be taken for monitoring. It seems prudent to measure several of the other analytics at the same time. The majority of the ionic contaminants could be measured on-line but this is at great expense and TC 9.9 is unaware of any installations where this is warranted. An intermediate, but still at some cost, approach is to monitor conductivity, pH, and turbidity. These can each be measured on-line with basic instruments. But even this level of monitoring is likely only warranted for the largest systems.

There has been some discussion of biofilms in the industry, as more liquid cooling systems are used. These are naturally occurring phenomena in water systems and in and of themselves are not a problem, if water chemistry is stable and ionic and organic loading is minimized. Biofilms are also kept in check by proper fluid design. Minimizing low velocity areas and avoiding dead-legs in the fluid stream are key.

Another interesting challenge is that of material compatibility. When reviewing materials in the TCS loop, we begin by considering what materials are common or acceptable in other cooling loops. The possible issue here is that typical building cooling loops may be in the 41°F to 59°F (5°C to 15°C) temperature range. Materials that work well at these temperatures may not work well at all at the warmer temperatures in our TCS cooling loop. With loop designs for W2 through W4, we can see fluid temperatures of 122°F to 140°F (50°C to 60°C) (see *Liquid Cooling Guidelines for Datacom Equipment Centers*, second edition). Material compatibility should be examined at these temperatures rather than the initially considered cooler loop temperatures.

Wetted Materials

Proper wetted material selection is just as important as proper water chemistry to ensure the health of the water-cooled server. In this context, the term “wetted material” shall reference any material that comes in direct contact with the water-based cooling fluid. The ASHRAE TC 9.9 liquid cooling book (ASHRAE 2014) does offer some limited recommendations on which wetted materials to avoid and which to use. As more water-cooled servers have made it to the marketplace, the list of wetted materials is increasing.

In this white paper, the recommended wetted materials list will be amended based upon what has recently been released by ITE manufactures. It is important to note that this is not an endorsement of such materials but, merely the latest in a growing list of materials that are being used by ITE manufactures for the water loop. The designer must understand the risk associated with the material selections and abide by materials compatibility, local ordinances, and best practices when choosing a wetted material.

One unexpected aspect of this topic is the sheer number of materials that do contact the working fluid. A great amount of due diligence is required to fully identify wetted materials in the entire water loop. As understanding of material compatibility with the working fluid evolves, the designer should continuously monitor the wetted material lists from end to end for the entire system. The designer should also require component suppliers to provide a wetted materials list to ensure overall compatibility.

To serve as general guidance, the wetted materials that were commonly found in commercial and available hardware will be listed. Table 2 shows the commonly found wetted materials in the current water-cooled servers. Materials in bold are originally listed in the liquid cooling book (ASHRAE 2014).

In addition, careful consideration must be given to the use of solder and brazing. Soldering is not recommended as solder joint reliability is poor because of the relatively high porosity in solder joints. Brazing is the recommended method for joining water-carrying copper hardware. Neither brazing or soldering should be used for joining steels or stainless steels (ASHRAE 2014).

Table 2 Common Wetted Materials Found in Water-Cooled Servers

Material	FWS	TCS
Acrylonitrile butadiene rubber (NBR)	X	
Aluminum and alloys	X ^a	X ^a
Brass with <15% zinc	X	X
Brass, chrome plated	X	X
Brass, nickel plated	X	X
Carbon steels ^b	X	
Copper^c	X	X
Copper alloys: <15% zinc and lead free^c	X	X
Polyoxymethylene (POM)	X	
Ethylene propylene diene monomer (EPDM)	X	X
Fluoroelastomer (FKM)	X	
Fluorinated ethylene polypropylene (FEP)		X
Polyamide (PA)	X	
Polychloroprene (CR)	X	
Polyethylene		X
Polyoxymethylene (POM)	X	
Polyphenylene sulfide (PPS)		X
Polytetrafluoroethylene (PTFE)	X	
Polypropylene (PP)	X	
Polysulfone or polyphenylsulfone (PSU, PPSU)	X	X
Silicone	X	
Stainless steel, solution treated and passivated ^d	X	X
Thread sealant ^e	X	X
Polytetrafluorethylene tape ^e	X	X

a. Though normally avoided in liquid cooling loops because of its potential as a galvanic corrosion catalyst of other metals, aluminum has been effectively used in FWS and TCS cooling loops provided there is proper surface treatment of the metal (via coatings or other proprietary methods) such that the aluminum is not directly exposed to the same liquid as other metals in the cooling system. Aluminum may also be considered in a TCS environment that leverages a non-conductive (dielectric) liquid.

b. For use with CDU implementations only. From the liquid cooling book (ASHRAE 2014) and regarding the FWS side, carbon steels may be used provided that steel-specific corrosion inhibitor(s) are added to the system and proper inhibitor concentration is maintained.

c. In reference to the liquid cooling book, copper is only listed in the TCS wetted material list. There are several versions of copper alloys listed in the FWS section of the liquid cooling book (ASHRAE 2014).

d. *Metal passivation* refers to the process of either creating new or restoring previously lost corrosion resistance properties on the surface of various metals. Many metals or alloys that exhibit strong corrosion resistance properties can lose those features at the surface of the metal during heat working. An example of this is 300 series stainless steels, which derive their strong corrosion resistance from chromium oxide formation across the surface when the constituent element chromium is exposed to atmospheric air containing oxygen. During weldment, the extreme localized temperatures liberate the chromium from the surface, exposing iron to the atmospheric air, leading to iron oxide corrosion formation that then may propagate throughout the metal structure. Various solution treatment options (such as nitric acid) can quickly restore the chromium concentration in the surface of the metal by dissolving unbound iron. With this in mind, some metals on this list are good choices for use in FWS and TCS but need passivation consideration if high temperature work is completed upon the stock.

e. Both liquid thread sealant and tape have been effectively used in TCS and FWS environments. It is important, however, to follow best practices under the consultation and/or supervision of certified professionals with both sealing methods such that robust seals free of loop contamination are maintained.

Filtration

Most modern data centers automatically dispense chemicals to continuously treat water used in the facility water supply. While the chemical consistency of water can be regulated reasonably well, there are often challenges in ensuring solid particulate levels are well filtered at all points throughout the piping system. The reason it is challenging to control particulate levels in all liquid branches is largely attributed to pipe scaling that occurs downstream of centralized filtration devices. Often, liquid branches that are plumbed for future deployment of ITE can remain stagnant for long periods of time as the data center space is gradually filled out. This can result in sediment formation from scale and corrosion on the interior pipe surfaces. For this reason, localized filtration on the supply side of water-cooled datacom equipment (for both CDU and non-CDU implementations) is required to prevent water-borne sediment from fouling heat transfer surfaces and flow paths within the cooling devices.

Maintaining fluid quality is critical for liquid cooling applications reliability and performance. A critical ingredient to maintaining fluid quality is filtering the fluid. There are several considerations to keep in mind when choosing a filter. The particulate filter size, the filter effectiveness, and the resulting pressure drop through the filter.

There are two competing goals in selecting a filtration strategy. First, the finer the filter the less chance of fouling or blocking the downstream narrow flow channels. On the other hand, these finer filters are invariably more expensive, have a larger pressure drop, and/or require more filtration elements to get suitable flow.

There are two main terms used to define a filter's capability: the *beta ratio* and the resulting *filter capture efficiency* for any particular particulate size. Beta ratio is commonly used in the oil, gas, and hydraulic industry, but is a relatively new concept for water and water/glycol used in liquid cooling and is introduced for liquid cooling here.

The beta ratio is defined as the ratio of the number of particles of a given particulate size between the upstream of the filter and the downstream of the filter as shown in Equation 1. Beta ratio is meaningless unless the particle size corresponding to the rating is included. This rating is based on the testing method from ISO 16889:2008 (ISO 2008), which is a multipass method for evaluating filter performance.

$$\text{Beta}(x) = \frac{\text{Quantity of particles of size } (x) \text{ upstream}}{\text{Quantity of particles of size } (x) \text{ downstream}} \quad (1)$$

The filter efficiency is defined as shown in Equation 2.

$$\text{Filter efficiency} = \left(\frac{\text{Beta}^{-1}}{\text{Beta}} \right) \cdot 100 \quad (2)$$

For example, if a 50 µm filter has a beta ratio of 75, then there is only one particle greater than 50 µm downstream of the filter for every 75 50 µm particles entering the filter. Following the equations above, a filter with a beta ratio of 75 is also 98.7% efficient at capturing 50 µm particles. Higher beta numbers represent better filtration for a given particle size.

Filters are rated in two ways: absolute rating and nominal rating. The *absolute rating* refers to the largest diameter spherical particle that can physically pass through the filter under laboratory conditions. The absolute rating does not necessarily reflect the filters ability in real-world applications. The *nominal rating* refers to the number of particles that can pass through the filter of a particular size. In the example above, the filter has a nominal rating of 75 for a 50 µm particle and is 98.7% efficient at removing particles of 50 µm in size.

It is common to refer to a filter with a nominal rating greater than 75 as effectively an absolute filter because of its ability to filter a high percentage of particulates for a given size. When selecting a filter, it is recommended to choose a filter with a beta rating of 75 or greater for the smallest particulate to be filtered.

A given filter will have a different beta ratio for each particulate size. Consider a hypothetical filter, for which Table 3 shows the beta ratio and resulting filter efficiency. Say, for example, the intent is a beta ratio of 75 for 50 µm particles. The filter may also have a beta ratio of 200 for 100 µm and 20 for 5 µm. Therefore, depending on the filtration needs of different particle sizes, an appropriate filter can be chosen.

Figure 10 shows three levels of particle filter quality for a hypothetical filter. Any given filter will have different filter efficiencies for each particle size passing through the filter. Depending on the particle size(s) of interest, the beta ratio and filter efficiency can be understood and can be used to choose the desired filter. Review the values toward the top of Figure 10 to aid in the description of the beta ratio and filter efficiency. The first value is there are 200 particles of 100 µm in size entering the filter. On the exit side of the filter, there is only one particle of 100 µm exiting; therefore, the filter captured 199 of the 100 µm particles meaning the filter is 99.5% efficient at capturing 100 µm particles. The beta ratio is 200 for

Table 3 Hypothetical Filter with Multiple Particle Beta Ratios and Resulting Filter Efficiency

Particle Size, µm	Particles In	Particles Out	Beta Ratio	Filter Efficiency, %
50	200	1	200	99.5
10	100	1	100	99.0
5	75	1	75	98.7

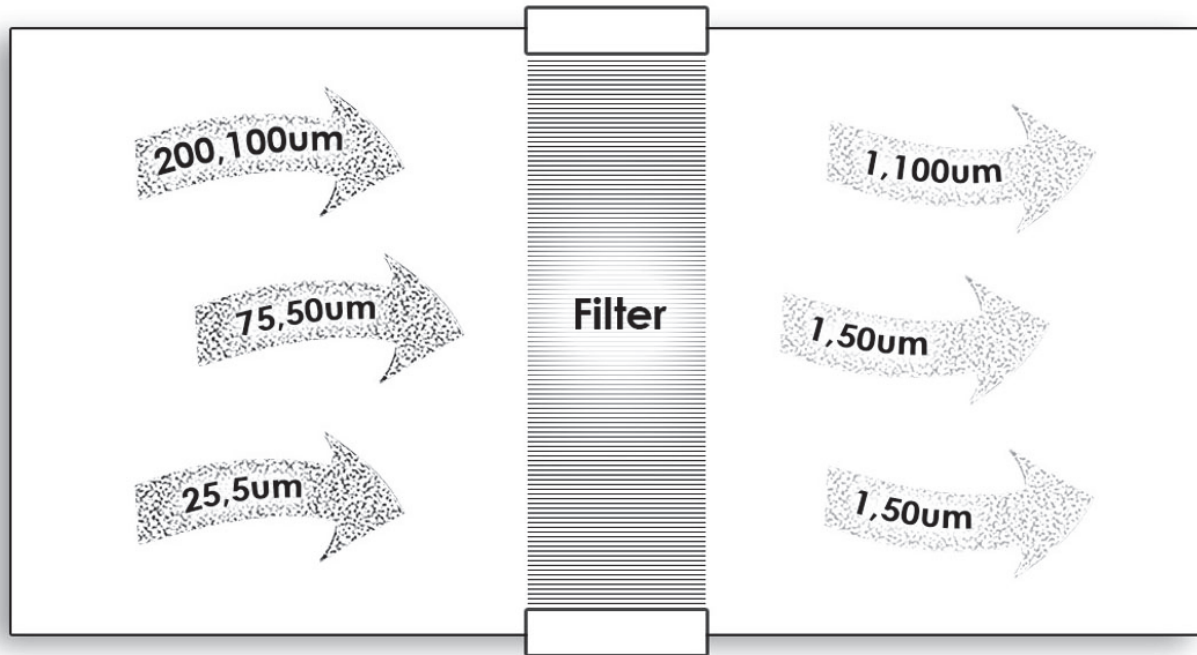


Figure 10 Hypothetical filter with multiple particle beta ratios.

the 100 μm particles. The next two lines in the figure are values for different sized particles, 50 μm and 5 μm , passing through the same filter and have different efficiency values and beta ratios.

To put debris and contaminate size into perspective, Table 4 gives examples of typical sizes of common particulates. Table 4 is taken from The Engineering Toolbox (www.engineeringtoolbox.com).

A fluid loop has two main ingredients that drive filtration requirements: the fluid connectors and the cold plate microchannel geometry. Cold plate microchannel passage dimensions range between 500 to 50 μm . It is informative to compare this typical value with the contaminants listed previously or even to other contaminants of interest. An industrial survey of filtration should be made with all involved parties: the CDU, quick disconnect, ITE manufacturer, system integrator, and end user should agree to the right ratio of minimum channel to filtration level. Experience suggests filter sizes ranging from 1/2 to 1/10 of the smallest size passages in the cooling loop. Most of the filtration industry refers to these ratios from the perspective of the filter and its size, not the passage size in the system, as the filter designers do not know what the passage size in the system will be while designing the filters. As such, the naming convention is from the filter perspective and is referred to as 2X for the 1/2 and 10X for the 1/10 protection ratio. As an example, a water-cooling loop that has a cold plate with 50 μm gap sizes could use a filter sized between 25 μm (2X) to 5 μm (10X). Therefore, a fluid passage size of 50 μm and a filter with 25 μm passages filters to a 2X ratio. The choice of the exact filter size will be guided by the amount of filtration pressure drop that can be overcome in that system.

Table 4 Examples of Particle Sizes

Particle Description	Approximate Minimum Particle Size, μm	Approximate Maximum Particle Size, μm
Sand	100	N/A
Human hair	40	300
Pollen	10	>500
Mold	3	30
Bacteria	1	50
Atmospheric dust	<1	40
Metallic particles	≤ 1	>20

Practical considerations can drive the ratio and the designer needs to investigate filtration recommendations from all component manufacturers (CDU, manifolds, fluid disconnects, cold plates, and so on). If the 10X value is used, the hardware will obviously be better protected, but at the expense of pressure drop across a finer filter (driving pump capital and operating costs). This could be countered with an increase in the filter cross section area but that increases filter and filter housings with added cost and larger size.

Filters can be implemented as full through flow, side stream or perhaps a combination of both. With a full through flow filter implementation, all of the fluid passes through the filter at all times. In general, this implementation may utilize a filter closer to 2X to minimize pressure drop while providing adequate filtration.

Alternately, a side-stream filter implementation could have a filter at the 10X or higher for a fraction of the flow. Over a longer period of time, all the water will be filtered at the finer level. Particulate loading or fouling from something internal breaking down is generally gradual and such a side-stream filter could prove very effective at preventing a problem. A step function increase in contaminant may happen during a refill of the loop water, and it is beneficial to protect against that case by using a filter finer than the main loop filter.

Figure 11 shows a schematic of such a side stream filter. The main cooling loop pumps can be used to provide the differential pressure across the side-stream filter taking some small percentage of the flow back to the inlet of the pumps. Values of 5 to 10% of the full loop flow rate are not uncommon for the volume of water recirculated through a side-stream filter. Note that this set-up can be simple but monitoring flow decay as described in a subsequent section is critical. In the set-up shown, no redundancy is needed. It should not be an issue to take the filter off-line for maintenance. Other set-ups with a side-stream pump can also provide side-stream filtration but these are typically more expensive and complicated. Side-stream filters can provide significant filtration improvements and IT protec-

tion improvements without the cost and energy needed for such fine filtration for the full stream.

Fluid connectors are the other driver that require filtration. Leakage of quick-disconnect fittings can occur in particle-laden fluids. Make sure to refer to the fluid connector supplier for guidance. The appropriate filter is the highest beta ratio with the smallest particle size between the requirements for the microchannel and the fluid connectors.

Another consideration for filtration is startup versus normal operation. Starting the system with the cleanest water as possible, with filtration as described above, is the most important factor for the long-term cleanliness of the system—start clean, stay clean. Depending on the system, the contamination and debris in the loop may be much higher at initial startup for a new deployment. One option may be to start with a larger filter requirement (i.e., lower beta ratio with larger particulate size) and slowly increase the filter efficiency until the fluid is at its target quality and the filter is also at its target operating capture efficiency size. This cleanup of the system should be done before installation of the ITE incorporating the tightest channels, since using those channels as part of the cleanup defeats the function of the tight channels.

Filter monitoring and maintenance is also a critical part of a successful operation. As stated previously, monitoring filter loading by measuring pressure drop will allow the owner to properly time and execute filter changes. The ideal sce-

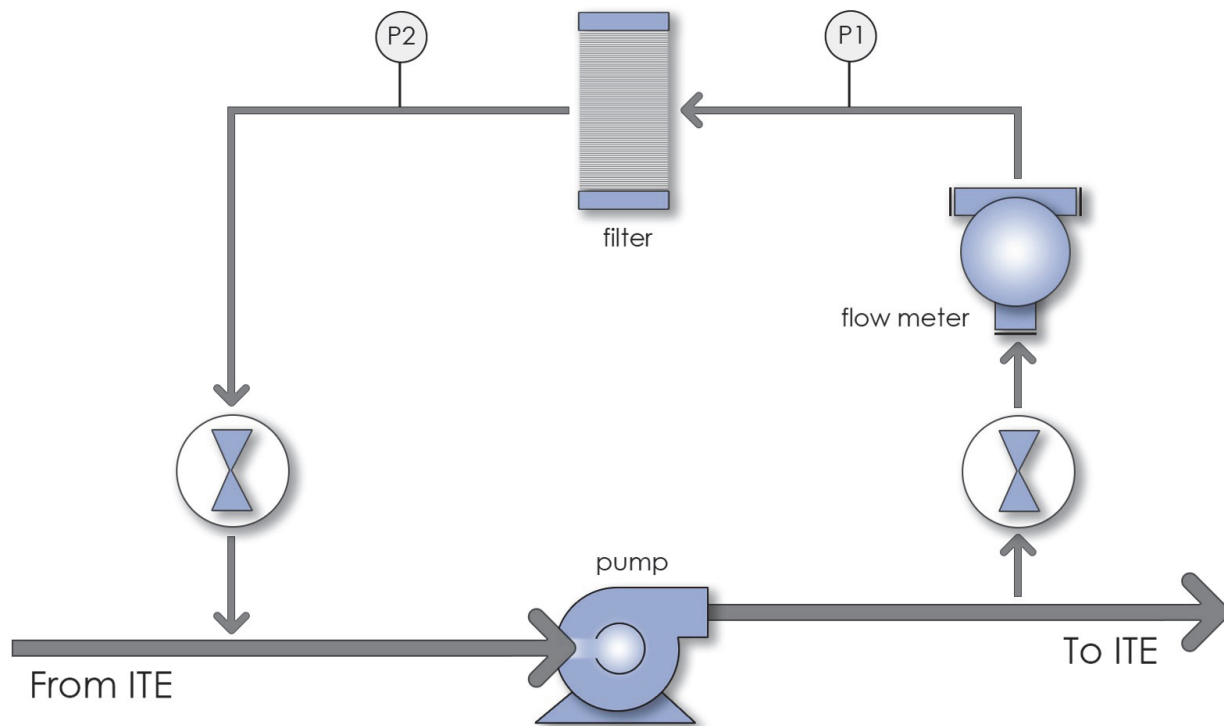


Figure 11 Example of a side-stream filter in a cooling loop.

nario is to do this with on-line instrumentation but can also be done successfully by periodic monitoring using a clipboard and a spreadsheet. The concept of flow decay needs to be implemented for tracking purposes. This is done by dividing the flow through the filter by the pressure differential across the filter (Equation 3).

$$\frac{\text{Flow}_t}{\Delta P_t} \rightarrow \frac{\text{Flow}_1}{P_{in,1} - P_{out,1}} \rightarrow \frac{\text{Flow}_2}{P_{in,2} - P_{out,2}} \quad (3)$$

This value should be trended over time. Checking the flow or pressure alone is insufficient. Both flow and pressure are needed as the ratio decreases by reduction of fluid flow and/or increased pressure drop, the trend can be used to extrapolate the timing for filter replacement. During initial system start-up, this check should be done daily. Once a pattern is established it is expected that the required frequency would move to weekly, then to monthly, or more, depending on the stability of the system. Appropriate use of statistical process control methodologies should be the guide to setting up this system, but as stated, it can be as simple as regularly writing down pressure drop and flow rate and entering them in a spreadsheet, tracked over time.

Pressure Testing Requirements for Water-Cooled Servers

Safety is paramount in anything ASHRAE TC 9.9 does, and it is with this consideration that a section on pressure testing liquid-cooled ITE is included. The applicability of standards is best left to the customer and their local code compliance authorities or representatives. With this in mind, TC 9.9 feels its role is to provide guidance and links to the range of standards and not to develop a pressure testing or IT safety standard—those already exist.

The primary safety standards for ITE are IEC 62368-1, *Audio/Video, Information and Communication Technology Equipment—Part 1: Safety Requirements* (IEC 2014) and IEC 60950-1, *Information Technology Equipment—Safety—Part 1: General Requirements* (IEC 2005).

It should be noted that 60950-1 is an earlier standard that is prescriptive in nature; the IT industry is transitioning to 62368-1, which is hazard-based. It is beyond the scope of this white paper to discuss, but a helpful summary is available from Underwriters Laboratories (UL 2012).

In addition to the IEC codes for ITE, the other codes that data center practitioners need to be aware of are the ASME B31.n series, specifically, B31.3 (ASME 2016) for interconnecting piping. The ASME has defined piping safety since 1922 and engineers in the facilities world are familiar with this series. These codes typically apply to the site-based interconnecting piping between the system providers of hardware and/or the facility. Note also that any toxic, flammable, or corrosive fluid used in liquid cooling is specifically excluded from these discussions and should be treated separately.

Reading the two pressure test standards in IEC 62368-1 and ASME B31.3 results in different guidance for pressure testing. At the risk of oversimplifying (the reader should examine both standards and fully understand the implications and methodologies of each), the IEC standard requires a leak test at 5x the normal design pressure, while the ASME standard requires a test at 1.5x the design pressure.

ASHRAE TC 9.9 began discussions with IEC TC 108 (the committee responsible for IEC 62368-1) and was informed the 5x value has been changed to 3x under normal operating conditions and 2x under abnormal and single fault conditions in IEC 62368-1, third edition (IEC 2018).

One concern on the part of TC 9.9 is that even the 3x requirement may cause design issues and preclude liquid cooling in some instances. When the liquid cooling loop is driven by a CDU or a facility cooling system the required test pressure

could be well over 100 psig (6.9 bar) and could cause problems for some of the heat transfer components. Of particular interest are the cold plates attached to the silicon components. These are generally thin, rectangular components where the lid could deform at high pressures. Designing the structural hold-down mechanisms to allow a 3x or 5x pressure test could render the thermal solution too expensive or too large for use in today's data-center-based ITE.

There is a belief that the original IEC 5x test point came from a consideration of a fully contained liquid cooling loop inside a laptop or server, leveraging heating, venting, and air-conditioning standard requirements based on a two-phase cooling system.¹ IEC 62368-1 has a broad applicability. This is reinforced by the use of liquid filled component (LFC) throughout the IEC pressure test section. In the case of a component inside the ITE, it is reasonable to assume the design pressures of an LFC are much lower than a row or room-based cooling loop and, as such, the 5x test value is not a problem. Additionally, the work on the original IEC 62368-1 code was done before liquid cooling had become as common as it is today.

With all this in mind, it is worth considering the ASME code. In the case where a CDU or facility-level system provides the cooling fluid to the ITE and is the home of the pump, all of the site-based piping is tested at 1.5x under the ASME code. It then makes sense that the rest of the cooling loop inside the ITE should also be tested to 1.5x. Recall that in this design there is no pump inside the ITE, it is entirely passive. Additionally, the CDU or facility system regularly includes pressure relief valves. The ability to use 1.5x as a pressure requirement in the design of the IT cooling hardware allows a wider, more economical adoption of liquid cooling.

To this end, TC 9.9 and the IEC safety committee have agreed to collaborate on exploring the right value for pressure testing externally connected ITE over the coming months and to provide guidance that is both practical and safe for the IT hardware.

As of this writing, TC 9.9 recommends following the IEC code for IT hardware assemblies delivered to the data center, and recommends the facility follow ASME B31.3 for the site-installed interconnecting piping. The reader is encouraged to stay aware of the IEC code as changes are expected, such as publishing the 3x requirement and possibly going to 1.5x in the version of the standard expected to be published in 2021.

The discussions and review of this topic made TC aware that this topic is confusing and additional guidance for the end user may be required. The following is intended to make this clearer:

- Pressure testing the ITE at the data center is not recommended, aside from checking for leaks when the final connection is made. The IT manufacturer is responsible for ensuring the equipment is sound when it is shipped. This equipment or assembly may be at the blade or server level

1. *IEEE PSES Newsletter*, Vol. 3 No.4, pages 24–35 by Lal Bahra, PE, published Dec 2007. References include UL 1995, the heating and cooling equipment (e.g., HVAC) standard. The UL1995 5th Edition, published in 2015, reduced the hydrostatic 5x test to 3x, amending the original 5x requirement.

up to the rack or rack/CDU combination. Shutting down or taking off-line operating ITE such that new assemblies could be pressure tested when added to a liquid loop in the field should be avoided.

- The IT manufacturer is responsible for pressure testing the assembly they will ship. This is generally done to the IEC code.
- The end user and/or their construction contractor is responsible for the pressure testing of the site-designed and installed interconnect piping/hosing between the site-based systems and IT-supplied hardware. This is generally done to the ASME B31.3 (or international equivalent) code. This is typically done prior to the installation of the IT kit.
- It is the responsibility of the end-user/owner to ensure the ITE pressure rating exceeds that of the cooling loop it will ultimately be attached to. Detailed discussions between the owner, IT supplier, and CDU supplier (if different than the IT supplier) are strongly recommended.

TC 9.9 expects that additional specific guidance on pressure testing will be included in the upcoming third edition of *Liquid Cooling Guidelines for Datacom Equipment Centers*.

Fluid Couplings Used in Water-Cooled Systems

Successful implementation of data center water cooling systems requires consideration of fluid connection points, which are critical to overall system performance and reliability. These points commonly involve quick-disconnect fluid couplings, allowing for connection and disconnection during operation. The types and requirements of connections can vary based on the portion of the system they serve.

In the water-cooled data center seen previously in Figure 6, large flow fluid connection points can be found at the condenser water system (CWS) and facilities water system (FWS) separating external cooling towers and chillers from datacom equipment. Within the ITE environment, key connections may be made at the interface of the CDU to the TCS as well as through the datacom equipment cooling system (DECS) at the rack level, or to rear door heat exchangers.

The following are possible serviceable fluid coupling points for each system:

- **FWS**—chiller, in rack-CDUs
- **TCS**—CDU secondary side branch, rear door heat exchangers, ITE cabinets/nodes

The primary attributes to consider when specifying fluid couplings, and recommendations for use in water cooling systems are outlined in the subsequent subsections.

Key Characteristics

The primary driver of using fluid couplings in water cooling systems is to be able to make and break fluid connections without affecting the TCS loop operation, thus creating a serviceable link between equipment. In electrical-based systems, this process might be referred to as *hot-swappable elements*. The goal of this section is to identify key elements of the fluid coupling that relate to system operation, performance, and reliability.

When specifying fluid couplings for water cooling systems, it is important to fully consider the materials of construction, both wetted and non-wetted. Particularly for coupling components exposed to the fluid flow path, materials should be chosen that are compatible with the fluid that therefore reduce the risk of corrosion, contamination, physical failure, and so on. The material selection can be key to other system attributes as well such as temperature, pressure, condensation, weight, and regulatory requirements such as Restriction of Hazardous Substances (RoHS) (EU 2011), UL94 (UL 2013), and others.

Fluid Coupling Physical Attributes

Many styles of fluid couplings are available for use in water-cooling systems, but the general anatomy is relatively common. The coupling consists of two halves, often gendered male/female (plug/socket, insert/body, etc). For sensitive environments, such as TCS and DECS, coupling shutoff with minimal fluid spillage is recommended. The fluid couplings are typically separated into two main categories—wet break and dry break.

A *wet break* connector refers to a coupler that is not closed when the connection is broken. The fluid is allowed to escape without blockage. The attached piping should be drained before separation. Figure 12 shows an example of some wet break connectors.

A *dry break* connector refers to a coupler that is closed or sealed when the connection is broken. The fluid is not allowed to escape from either side of the connection. Both halves of the coupling have a shutoff feature to seal off fluid flow during disconnection. The amount of fluid released can vary. Couplings with this type of shutoff may also be described as having *nonspill* or *flush face* valves. The method of connecting the coupling is often manually through a latching mechanism integrated into the fluid coupling (direct dock) or mounted directly to a piece of equipment with separate retention features (blind mate). Figure 13 shows examples of a dry break quick disconnects.

Coupling Terminations

The coupling set can integrate with the water cooling system through a variety of termination styles. When using flexible hose or tubing, consideration should be given to the style of hose barb connection, and whether additional clamps or ties should be employed to ensure reliability during operation. More details on flexible hose use are available in the following section. When using rigid tubing or mating the connector directly with a port or manifold, a threaded termination may be used. Tapered pipe threads may require additional sealing paste on installation to ensure a leak-free connection point. Straight thread terminations with elastomeric O-ring sealing such as Society of Automotive Engineers (SAE) or British Standard Pipe Parallel (BSPP) can be a reliable termination option for TCS or DECS, where risk severity of leaks is more critical. For fluid connections near ITE, consideration of vibration should be taken at installation and appropriate installation torque should be applied to resist loosening over time.



Figure 12 Wet break quick disconnects.



Figure 13 Dry break quick disconnects. Both plug and sockets of three different models represented.

Fluid Coupling Performance

In order to ensure consistent and reliable operation of the water cooling system, fluid couplings should be specified accordingly for the segment of system in which they operate. For fluid connections at the FWS, where emphasis is on high flow and operating pressure, couplers have large throughput with low impedance. Fluid connections at the TCS and DECS are generally smaller in nominal flow size, but more frequent as fluid is distributed throughout the rack. Across the entire system, drips and leaks from fluid connections can be problematic, but are particularly so in the DECS where critical and sensitive components are in operation.

Coupling Performance Attributes

A proposed summary of fluid coupling performance attributes is provided in Table 5. For more information on pressure requirements and testing, refer to the previous section, Pressure Testing Requirements. The attributes will be defined as follows:

- **Flow rate**—Volumetric fluid flow rate supported for safe, reliable operation. When specifying a connector set, preferred installations promote maximum flow with minimal pressure loss. Typical ranges for water-cooled server applications may be 0.1 to 10 gpm (.38 to 31.85 L/m).
- **Flow coefficient**—Dimensionless factor used to characterize the flow performance of a fluid coupling set; a correlation of volumetric flow rate to pressure loss for a known fluid. Often reported as a C_v (or K_v) value, where a C_v value is representative of the volumetric flow rate expected to result in 1 psi (6.9 kPa) pressure loss. Expected flow coefficients of fluid couplings for water-cooled server applications may be between 0.1 and 8.0+, based on the size of the fluid loop being served.
- **Operating pressure**—Maximum internal fluid pressure recommended for safe, reliable operation. Special consideration should be given for

Table 5 Fluid Coupling Performance Attributes

Attribute	Unit	Description
Rated flow rate	L/min (gpm)	Volumetric fluid flow rate supported for safe, reliable use
Pressure—Operating	bar (psi)	Maximum rated fluid pressure supported for normal operation
Pressure—Burst	bar (psi)	Minimum rated pressure to cause leak, catastrophic failure
Flow coefficient	—	Correlates volumetric flow rate to pressure loss through fluid coupling (i.e., C_v , K_v)
Spillage	cc (mL)	Volume of fluid loss on disconnection of coupling set
Temperature—Operating	°C (°F)	Rated temperature range for normal operation
Temperature—Storage	°C (°F)	Rated temperature range for storage, shipping
Connection force	N (lbf)	Force needed to fully connect coupling set, often will increase along with system pressure
Connection cycles	cycles	Minimum rated number of mechanical connections

extreme temperature applications, or when in combination with high flow rates. An operating pressure from 0 to 120 psi (0 to 827.4 kPa) may be considered for water-cooling applications.

- **Burst pressure**—Minimum internal fluid pressure for catastrophic failure of a fluid coupling. Common failure mode under burst pressure conditions may be elastomeric seal extrusion. Typical pressure limits for fluid couplings may be 2x to 5x operating pressure.
- **Spillage**—Volume of fluid present upon disconnection of fluid coupling pairs. Amount will vary depending on the type of shutoff integral to the connectors. Common fluid coupling embodiments may include poppet or flush-face valve types with flush-face generating the lesser amount of spillage. A flush-face connector may also be referred to as a *nonspill* or *dripleless* coupling, where the mating surfaces may only be minimally wetted on disconnection. Fluid coupling spillage is often a function of system pressure and flow rate on disconnection. Consult fluid coupling suppliers for product-specific information.

Table 6 Typical Fluid Couplings Characteristics

Attribute	Units	System Segment			
		Cabinet	Chassis	Drawer	Blade
Mating style	—	Direct	Direct	Direct	Direct, blind
Termination	—	NPT, ORB	NPT, ORB	NPT, ORB	ORB, barb
Nominal size	cm (in.)	5.1–10.2 (2–4)	5.1–10.2 (2–4)	1.3–5.1 (0.5–2)	≤1.3 (0.5)
Flow rate	L/min	<1000	<100	100	≤5
Pressure—Operating maximum	bar	10	10	10	5
Pressure—Burst minimum	bar	35	35	35	17
Spillage per disconnection	cc	—	—	≤15	≤1
Temperature—Maximum	°C	200	200	150	150
Connection force	N	—	—	—	≤50
Connection cycles	No. cycles	100	100	100	1000

Modes of Failure

The following are possible situations in which a fluid coupling may fail:

- **Pressure spike (beyond rating)**—catastrophic physical coupling failure, leak
- **Disconnection under high flow and pressure**—seal extrusion, incomplete shutoff
- **Water hammer**—seal extrusion, leak
- **Damage/wear of sealing components**—leak, unable to connect
- **Debris/contamination in flow path**—leak

Typical Fluid Coupling Characteristics

A brief survey of commercially available quick disconnects was conducted to determine the typical characteristics as it pertains to different locations of use. Table 6 shows the results of that survey for the following system locations; cabinet or rack, chassis, drawer, and blade. Further detail and specification can be obtained from the fluid coupling suppliers to ensure proper selection and implementation. The termination styles vary from threaded connections such as national pipe thread (NPT), O-ring boss (ORB), or hose barb.

Flexible Hose Used in Water-Cooled Systems

The use of flexible hose within liquid cooling systems for ITE is highly beneficial for cost mitigation, ease of service, and design simplification; however, there are several key design considerations that should be included. These hoses can be used to connect server elements to a manifold, as seen in Figure 14.

Another common use for flexible hoses is to provide internal interconnect within the server as well. Figure 15 shows how flexible hoses are used to route between CPU cold plates and dual in-line memory modules (DIMM) cooling structures.

Maximum Temperature Ratings

The maximum temperature rating of the flexible hose should be carefully considered. The inlet water temperature plus the caloric temperature rise of the fluid should be utilized to determine the bulk fluid temperature. The resulting highest bulk temperature of the fluid should be used as a minimum temperature to consider. However, one must pay attention to all the surfaces that come in contact with the flexible hose. If this flexible hose is brought to the cold plate of a CPU, for example, the temperature of the contact surface should also be considered. This surface temperature will be at a higher temperature than the bulk fluid tem-



Figure 14 Flexible hose used to connect a manifold to a server blade in a system.

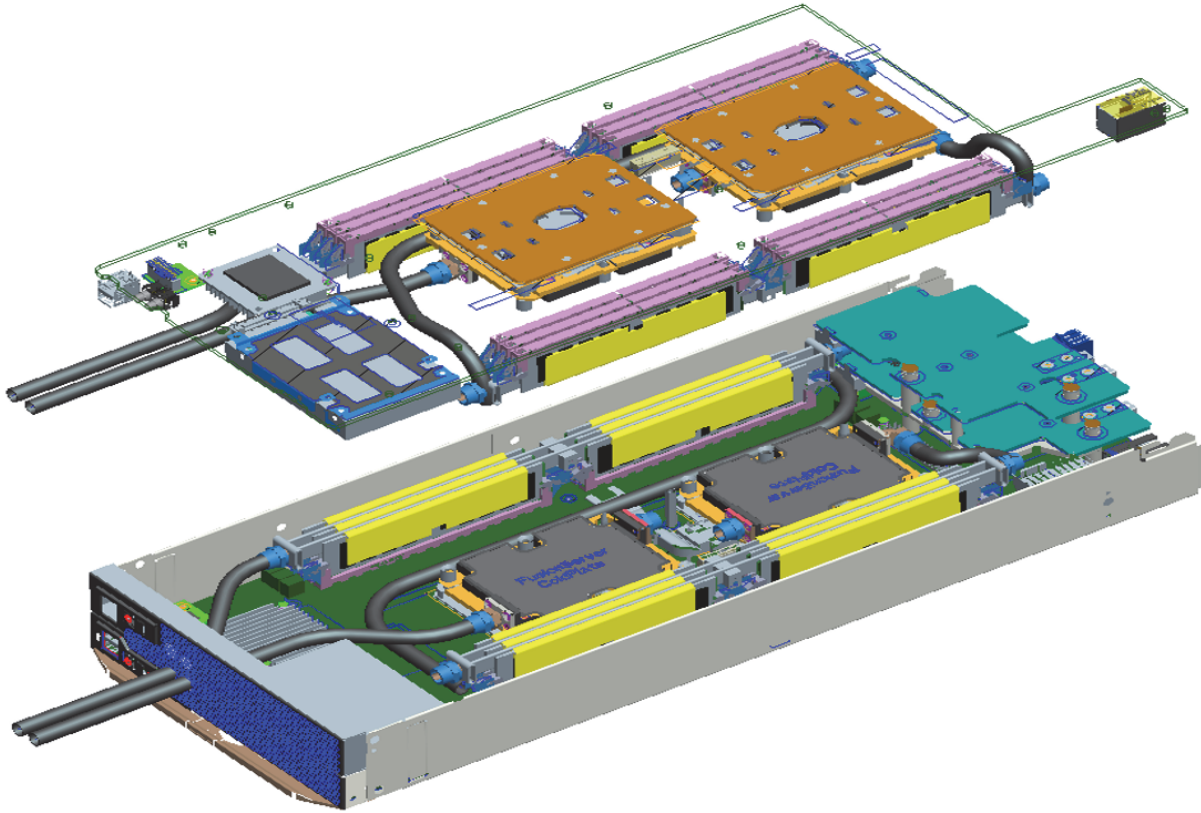


Figure 15 Flexible hose used to connect CPU cold plates and DIMM cooling structures in a server.

perature. The design must carefully determine the temperature rating of the hose based on the highest temperature seen by that hose.

Maximum Pressure Ratings

The flexible hose is most likely the most susceptible component in the water-cooled system to high pressure. As discussed in the previous section, the operating and burst pressure should be determined and sized as to give appropriate factors of safety when using flexible hoses.

Flammability Requirements

The flammability requirements for the system must be maintained when selecting the flexible hose. The designers company, local geographies, or customer sites may have very specific requirements regarding the flammability requirements of the products in use. Typically, the flexible hose should meet a minimum rating of UL94-V1 (UL 2013).

Hose Material

Because flexible hose is in contact with the fluid for a long time, compatibility with the working fluid must be considered when selecting the hose material. In addition, low hose permeability must be taken into account to ensure a significant system life. Suggested materials when using deionized water or a mixture with ethylene glycol include PTFE, EPDM, NBR, FEP, FKM, polyurethane, and nylon. It is better to consult with the fluid supplier to get more information on the selection of hose materials.

Hose Reinforcement

Another design consideration when using flexible hose is the use of reinforcements in the hose material. Some flexible hoses come with a form of reinforcement that enhances the pressure ratings of the hose, as seen in Figure 16. It is very important to understand if that reinforcement ever comes in contact with the fluid. There may be a protective layer on the inside of the hose that could fail with puncture, stress, or pressure and result in a non-desired material coming in contact with the internal fluid. An increase in hose bend radius is another disadvantage of adding hose reinforcement.

Bend Radius

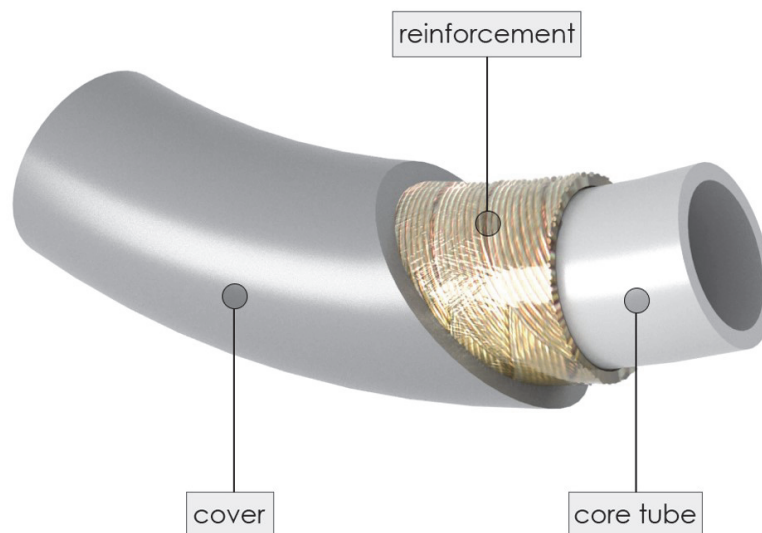


Figure 16 Flexible hose reinforcement.

The bend radius of a flexible hose can vary greatly and is typically the smallest radius that the hose can be bent, as seen in Figure 17. It is important to design to the manufacturers' recommended bend radius. This ensures that the hose does not deform or become structurally compromised. Each manufacturer will have a specific bend radius for specific flexible hose types.

Termination

The termination of the flexible hose can be a critical element to their successful use. There are several popular styles of termination and connection to other components: hose barbs, pagoda joints, compression fittings or sleeves, quick release, and push-on styles.

Fluid Interaction

The flexible hose must be materially compatible with the working fluid and surface materials it comes in contact with. In addition, the permeability of the hose should be well understood and must be low enough such that the fluid will not significantly move through the hose material. Finally, the flexible hose material may absorb the chemicals being used to provide anticorrosion protection. It is



Figure 17 Example of a flexible hose bend radius.

strongly recommended that a long-term study of the hose material and working fluid be conducted to look for these issues.

Thermal Expansion

Flexible hose materials tend to experience greater thermal expansion than hard tube/pipe materials. This expansion and contraction can gradually loosen a flexible hose at the fitting connection points. This may suggest applying double hose clamps in some cases. Additionally, certain hose materials may become more brittle and less suitable for internal liquid pressure cycling at lower temperatures. To ensure liquid solution reliability in deployment, temperature cycling should be conducted beyond the supported operating and nonoperating temperatures for the ITE.

Puncture Prevention

Sharp edges from server sheet metal, poor quality control on barb fittings, liquid-born metallic shavings left in an attached piping system, and other nearby mechanical structures are just a few examples of elements within a liquid data center environment that might potentially puncture the sidewall of a flexible hose. Poorly machined edges on barb fittings can create small cuts on the inside of flexible hoses that weaken the material's ability to withstand maximum rated pressure and may present fatigue failure in the future. It is recommended that shielding and/or hose reinforcement be implemented in and around flexible hose installations to prevent unintentional puncture.

Clarifications of the FCS and TCS Loops Described in *Liquid Cooling Guidelines for Datacom Equipment Centers*, Second Edition

In *Liquid Cooling Guidelines for Datacom Equipment Centers*, second edition (ASHRAE 2014), there are four loops described. The first and closest to the heat source is a DECS, next is the TCS, then the FWS, and finally the CWS. In practice, the DECS terminology is rarely used and the TCS loop is given much greater usage. The primary difference between the two loops is the physical boundary (the DECS being fully inside the electronics rack). When the DECS loop uses water, there is no other significant difference. Note also there are full chapters dedicated to the TCS loop (Chapter 6) and the FWS loop (Chapter 5), but not the DECS loop. In general, the data in Chapter 6 is largely applicable to a DECS loop (if using water) which is essentially a TCS loop inside of the rack boundary. The main reason for referring to a separate DECS loop is when that loop contains a fluid different from water, e.g., a proprietary heat transfer fluid.

The above-referenced chapters have also been a source of confusion through their occasional misapplication. This is described in the above section on water quality and the FWS and TCS guidelines. There have been instances where these have been misapplied. For example, when the facility owner is told their building-level cooling system (essentially an FWS) should meet the water quality guideline table for the TCS loop from Chapter 6 by mistake, when they are likely already at or near the guidelines in Chapter 5, the wisdom of going to liquid cooling is called into question. Conversely, applying Chapter 5 water (FWS) to a very fine pitch microchannel cold plate on a CPU could lead to plugging, corrosion, or other failure. Neither are good, and misapplying Chapters 5 and 6 to the wrong loops can be catastrophic.

Another common pitfall is using Table 5.1 or Figure 5.3 in Chapter 5 too literally. These provide information about typical installations and methods to make a water that meets W1–W5. There is no requirement to use the typical equipment called out as W_n at any level. The engineer should in fact work to provide the water temperatures called out with the most efficient and lowest total cost of ownership (TCO) method possible. If a site can produce W2 water using cooling towers alone there is no requirement to have a chiller, even though a chiller is listed as typical for W2. Over-applying these guidelines can cost money.

Future Work and Additional Resources

There are other ongoing activities outside of ASHRAE focused on liquid cooling of ITE. The following are two of particular note:

- **The Energy Efficient High-Performance Working Group (EEHPCWG)**—The EEHPCWG is a volunteer organization of primarily high-end computing and simulation practitioners. This group has a subteam working on commissioning. Liquid cooling has shown to be more quickly adopted into high-performance computing than enterprise or cloud workloads and because of this the working group has an extensive user base to draw upon for lessons learned and best practices. The EEHPCWG web-sites are:
 - <https://eehpcwg.llnl.gov/index.html>
 - https://eehpcwg.llnl.gov/infra_lcc.html
- **The Green Grid (TGG)**—The Green Grid has a long history of collaboration with TC 9.9 and that organization is currently working on select liquid cooling topics, specifically IT design enabling multi-refresh liquid cooling and a total cost of ownership (TCO) calculator for liquid cooling data centers. TGG's web site is:
 - www.thegreengrid.org

Immersion cooling continues to be an area of interest, although adoption rates have been slow. There are generally two classes: first mineral oil or synthetic oils (polyalphaolefin [PAO]), or engineered fluids. In the case of oils, there remain minor material compatibility issues that all parties need to be aware of and address in the implementation. In the case of engineered fluids there have been studies which showed significant material issues for some fluids and other engineered fluids having global warming potential values that are so high they become nonstarters from an environmental perspective. The TC is considering a broader coverage of immersion cooling in the forthcoming third edition of *Liquid Cooling Guidelines for Datacom Equipment Centers*.

References

- ASHRAE. 2014. *Liquid cooling guidelines for datacom equipment centers*, second edition. Atlanta: ASHRAE.
- ASHRAE. 2018. *IT equipment power trends*, third edition. Atlanta: ASHRAE.
- ASME. 2016. ASME B31.3, *Process piping*. New York: American Society of Mechanical Engineers.
- EU. 2011. *Directive 2011/65/EU of the European Parliament and of the Council of 8 June 2011 on the restriction of the use of certain hazardous substances in electrical and electronic equipment*. Brussels: European Union.
- IEC. 2005. IEC 60950-1:2005, A1+A2:2013, *Information technology equipment—Safety—Part 1: General requirements*. Geneva: International Electrotechnical Commission.
- IEC. 2014. IEC 62368-1:2014, *Audio/video, information and communication technology equipment—Part 1: Safety requirements*. Geneva: International Electrotechnical Commission.
- IEC. 2018. IEC 62368-1:2018, *Audio/video, information and communication technology equipment—Part 1: Safety requirements*. Geneva: International Electrotechnical Commission.
- ISO. 2008. ISO 16889:2008, *Hydraulic fluid power—Filters—Multi-pass method for evaluating filtration performance of a filter element*. Geneva: International Organization for Standardization.
- UL. 2012. The road toward a new hazard-based standard: IEC 62368-1. Northbrook, IL: Underwriters Laboratories.
- UL. 2013. UL 94, Standard for tests for flammability of plastic materials for parts in devices and appliances. Northbrook, IL: Underwriters Laboratories.