

Energy Efficiency Of Liquid-Cooled Data Centers

BY DUSTIN DEMETRIOU, PH.D., MEMBER ASHRAE; DAVID QUIRK, P.E., MEMBER ASHRAE; TOM DAVIDSON, P.E., MEMBER ASHRAE

Almost since inception, data centers have been cooled by air. The metrics that measure their efficiency are based on air cooling, and ASHRAE Standard 90.4-2022, *Energy Standard for Data Centers*,¹ which ASHRAE has adopted to determine the minimum acceptable infrastructure energy efficiency, uses equipment that assumes that 100% of the server heat is rejected by the air of a data center. Liquid cooling is not considered in establishing the minimum energy requirements. What are the inherent efficiency advantages of liquid-cooled vs. air-cooled equipment? Are some liquid-cooling technologies more efficient than others? And are today's energy standards using the best metrics to identify the improvements offered by liquid-cooling technologies?

With the exponential increase in the use of artificial intelligence (AI) software in data centers, liquid cooling can no longer be considered a “fringe” technology in the data center industry. Some data centers are already being built to handle primarily liquid-cooled information technology equipment (ITE). Energy consumption of U.S. data centers is also growing rapidly, increasing from approximately 1.9% to 4.4% of U.S. annual electric energy consumption from 2017 to 2023.² Further increases are expected; a significant portion of this increase will be liquid-cooled ITE. With this increased use comes the responsibility to better understand liquid-cooled servers and how the

infrastructure can most efficiently cool them.

How is Liquid-Cooled Server Efficiency Different Than Air-Cooled?

Before discussing the various energy metrics associated with liquid cooling, it is worthwhile to identify two different types of liquid-cooled servers entering the market and the differences between these liquid-cooled servers and air-cooled servers. Two types of liquid cooling are (a) those in which only the

Dustin Demetriou, Ph.D., is chair of the ASHRAE TC 9.9 IT subcommittee. David Quirk, P.E., is president and Tom Davidson, P.E., is a senior mechanical engineer at DLB Associates.

TABLE 1 Comparison of cooling modes for air-cooled, liquid-cooled conductive and liquid-cooled immersion servers.

COMPARISON CRITERIA	AIR-COOLED SERVERS	HYBRID DIRECT-TO-CHIP LIQUID-COOLED SERVERS	IMMERSION LIQUID-COOLED SERVERS
Heat rejection from server to space	Air, using onboard server fans	Forced convection heat transfer from processor cold plate to pumped liquid. Two-phase also possible	Natural or forced convection heat transfer using dielectric liquid
Cooling medium transport power source	Fans (pumps may also be used “downstream” for chilled water systems)	Fans for residual air cooling, pumps for liquid	Pumps
Primary cooling equipment	DX CRAC units or chillers; direct to atmosphere possible depending on ambient conditions (economizer cycle)	Chillers for low S-classes with economizer available; direct to atmosphere for high S-classes*	Chillers for low S-classes with economizer available; direct to atmosphere for high S-classes*
Heat rejection to atmosphere	Refrigerant condensers, cooling towers, dry coolers	Cooling towers or dry coolers	Cooling towers or dry coolers

*S-classes³ refer to the inlet fluid temperature to the server from the technology cooling system. They are specified as: S-30 (30°C, 86°F); S-35 (35°C, 95°F); S-40 (40°C, 104°F); S-45 (45°C, 113°F); S-50 (50°C, 122°F)

high heat flux components (e.g., processors) are liquid cooled, typically with a cold plate, and (b) those that are 100% liquid cooled with the server submerged in the liquid coolant. The former is often called conductive cooling or hybrid direct-to-chip liquid cooling, and the latter is often called immersion cooling. *Table 1* illustrates some of these differences, along with a comparison to the more prevalent air-cooled server.

Two significant ways in which liquid cooling can be more efficient than air cooling are:

- Unit heat transport energy for liquid heat rejection is typically less than for heat rejection by air, greatly aided by the difference in the specific heat of the two (water is 1.00 Btu/lb·°F or 4.184 J/g·°C vs air at 0.24 Btu/lb·°F or 1.005 J/g·°C).⁴ The differential can be even greater with two-phase flow.
- Approach temperatures with liquid-to-liquid heat exchangers are typically smaller than for either air-to-air or for air-to-water, increasing potential economizer hours. In many cases, the heat can be rejected without use of a chiller for most or all hours in a year.

Regarding the first bullet, for a given processor energy draw, air-cooled servers have the highest peripheral component energy use (fans), followed by hybrid (fans, but to a lesser degree, and pumps, typically external) and finally immersion-cooled servers having the lowest peripheral energy use (pumps, typically external).

Complicating the measurement of liquid cooling savings via a widely-adopted metric—power usage effectiveness (PUE)—is the fact that the reduction in air transport energy occurs both within the server and external to it (i.e., both in the numerator [the total data center energy use] and the denominator [the ITE energy use] of the PUE equation). If the percentage energy reduction of the two is the same, the PUE will remain the

same even though the energy consumption of the facility drops significantly. This is not a good situation for an energy efficiency metric to be in!

For the second bullet, the main impact will be reduced chiller hours of operation. The higher the S-class, the greater the potential for increased economizer hours. The energy savings of liquid cooling is closely tied to the heat rejection pathway and the local climate. Among the heat rejection pathways for the technology cooling system (TCS) liquid options are:

- TCS liquid-to-air heat exchanger in the cooling distribution unit (CDU), with heat rejected to the data center space;
- TCS liquid-to-cold liquid heat exchanger in the CDU, with heat rejected to the chilled water loop;
- TCS liquid-to-warm liquid heat exchanger, with heat rejected to the dry cooler or cooling tower; and
- TCS liquid-to-hot liquid heat exchanger, with heat reclaimed for another use (such as office heating).

As one might expect, rejecting heat from a processor to a liquid coolant, and then using air to reject it back into the data center, is a less inefficient process than rejecting it to a liquid. But if a liquid-cooled server is installed in an air-cooled data center without any chilled or condenser water, it may be the only short-term option. On the other side of the spectrum, a data center that can operate at S40, S45 or S50 has the potential to reject CDU heat directly to the ambient environment without using a chiller. Supporting high S-classes (S40, S45, S50), while energy efficient, may not always be possible given future ITE chip power density trends.³

Heat reclamation is also an opportunity. A supercomputer cluster in Germany implemented waste heat reuse sufficient to reheat about 40,000 m² (430,500 ft²) of office space and to drive an adsorption

chiller for rear-door heat exchangers, capturing the non-liquid-cooled heat and eliminating the need to add computer room air conditioning/computer room air handler (CRAC/CRAH) to the data center.⁵ Another example is a data center park in Stockholm, Sweden, where data center waste heat is used to help heat a 12 TWh/yr city-wide district heating network.⁵

What Energy Metric Could Be Used to Measure the Efficiency Difference Between Air-Cooled And Liquid-Cooled Systems?

In recognition that the PUE metric might not be able to reflect the increased efficiency of a data center with liquid-cooled servers, several entities developed and published a metric called TUE, or total-power usage effectiveness.⁶ Although not in common use, it was developed more than a decade ago. It is linked to PUE in a simple way through an additional metric called the iTUE, or IT power usage effectiveness. Whereas PUE has the total data center energy use in the numerator and the ITE energy use in the denominator, iTUE has the ITE energy use in the numerator and the ITE compute energy use in the denominator. $TUE = iTUE \times PUE = \text{total data center energy use} / \text{ITE compute energy use}$.

How does the TUE fare in the case above where PUE can be misleading? In the case of concurrent reduction in both server fan energy and infrastructure fan energy, rather than staying about the same, like the PUE, the TUE clearly drops, as both sets of fans are located in the numerator of the metric, but not in the denominator.

Another instance in which TUE is a better measure of liquid-cooled data center efficiency is it is agnostic to CDU location and its associated technology cooling system (TCS) pumps. The pump location will impact the iTUE and PUE, but they are offsetting impacts that exactly cancel out in the TUE calculation. TUE may look like the winner as the go-to energy efficiency metric for liquid-cooled data centers. But one “minor” question remains: can the TUE be measured?

TUE Measurement

From a practical standpoint, for a metric to be used, it needs to be “easily” measured. As indicated earlier, the TUE is equal to the product of the PUE and the iTUE.

* The heat capture ratio (HCR) is the amount of heat that is rejected from the IT equipment to the TCS loop. HCR can vary from ~50% to >90% depending on the ITE implementation.

The PUE can definitely be measured, and in fact The Green Grid and ASHRAE have written an entire book on the subject.⁷ If server manufacturers can measure the iTUE—and ideally start reporting this value through the data center infrastructure management (DCIM) system—we are all set. For immersion liquid-cooled servers, this is probably the only way to obtain this information. ASHRAE’s *Advancing DCIM with IT Equipment Integration*⁸ requires ITE that meets the Tier 2 level to report the “system cooling subsystem power,” “power supply unit (PSU) input power” and “PSU output power,” which represent many metrics needed to calculate iTUE.

For direct-to-chip liquid cooling, DCIM measurements are still recommended, but another possible route, which could be obtained by metering the liquid-cooling infrastructure with Btu meters, would be to assume that the ITE compute energy consumption is equal to the liquid-cooled heat rejection. With knowledge of the entire facility load available from a building demand meter or equal and the ITE compute energy obtained from the Btu meters, a fairly good estimate of TUE could be obtained without even dividing the TUE into its component PUE and iTUE values. This method would be appropriate for ITE that has a high heat capture ratio.*

A Theoretical Case Study

A recent theoretical case study⁹ did a rigorous analysis of PUE, iTUE and TUE for a 1.7 MW 100% air-cooled data center converted to a 25%–75% mix of air cooling and water cooling. The authors looked at many cases, but a typical one had a 10.2% reduction in total energy (TUE), but PUE reduction was less than 4%. This illustrates the caution with which PUE should be used as a metric for savings associated with conversions to liquid cooling. In this example the authors did not consider the additional infrastructure energy savings that could have come from rejecting the liquid cooling loop water directly to a dry cooler during favorable ambient conditions rather than to a year-round chilled water loop.

The Ultimate Metric

Both TUE and PUE metrics focus on the data center’s “overhead” power. With recent advances in power distribution efficiency and the advent of liquid cooling, “overhead” power is poised to become quite small, with marginal benefits to further reduction. The ITE compute energy is the driving factor of efficiency in this scenario.

Ultimately, data centers need a performance metric that looks at the amount of useful work coming from the ITE workloads. The European Energy Efficiency Directive¹⁰ and its Delegated Act¹¹ requires reporting of ITE Capacity (information and communication technology capacity for servers, C_{serv}), a measure of the server's work capacity. This is not a new idea, just one that has proven difficult to implement in practice, particularly for comparative use, given the diversity of IT workloads. There are many examples of workloads and benchmarks focusing on ITE energy efficiency. A common one is the Server Efficiency Rating Tool (SERT),¹² which is a collection of 11 worklets that are aggregated into an efficiency score. Servers that receive the U.S. EPA Energy Star rating have met a minimum SERT score. ASHRAE Datacom Book 12, *Server Efficiency—Metrics for Computer Servers and Storage*,¹³ describes many other metrics in use.

Energy Standard/Code Implications

ASHRAE Standard 90.4-2022, *Energy Standard for Data Centers*, has a mechanical load component (MLC) and electrical loss component (ELC). The MLC and ELC represent minimum cooling and uninterruptible power supply (UPS)/electrical distribution infrastructure efficiency values, respectively, which are typically calculated during the design phase when the details of servers are unknown (though it probably will be known if the servers are to be air- or liquid-cooled). The ELC is a maximum electrical efficiency loss value, and the calculation is relatively straightforward. The MLC calculation is dependent on climate and the opportunity for economizer use.

While MLC comes from a complex design calculation, at its heart it is similar to PUE in that the denominator of the equation is ITE energy use, not the “core” ITE compute energy. Reducing fan energy inside the boundary of a server, as liquid cooling does, will not be reflected as a reduction in MLC. Moving the transport power for the heat rejection medium from inside to outside the server boundary (fans inside the server vs. CDU pumps outside) also represents a challenge for the MLC. The Standard 90.4 committee is planning to share its MLC calculations with the successful bidder of ASHRAE's 1972-WS Work Statement, which will be taking a closer look at the relative efficiency of air-cooled and liquid-cooled data centers under several

conditions. As such, this is evolving research.

Summary

In most cases, the switch from air-cooled to liquid-cooled servers will reduce data center facility energy consumption. Measuring this reduction, however, is challenging using the traditional PUE metric, and similar issues are expected in using Standard 90.4-2022's MLC metric. ASHRAE TC 9.9 has proposed a research project (1972-WS) to, among other things, compare the efficiency of cooling air-cooled servers to liquid-cooled servers in various climates and under various heat rejection topologies, using various metrics. This data could be useful to the Standard 90.4 committee as it decides how to handle possible changes to the data center energy code based on increased use of liquid cooling and to the data center industry at large as the migration toward liquid-cooled servers evolves.

References

1. ANSI/ASHRAE Standard 90.4-2022, *Energy Standard for Data Centers*
2. Shehabi, A., S.J. Smith, A. Hubbard, A. Newkirk. 2024. “2024 United States Data Center Energy Usage Report.” Report LBNL-2001637. Lawrence Berkeley National Laboratory. <https://tinyurl.com/9xy87689>
3. ASHRAE. 2024. “Liquid Cooling Guidelines for Datacom Equipment Centers, 6 Liquid-Cooling Infrastructure Considerations for Technology Cooling Systems.” ASHRAE TC 9.9 Datacom Encyclopedia. <https://datacom.ashrae.org>
4. ASHRAE. 2009. *Best Practices for Datacom Facility Energy Efficiency, 2nd Edition*. Atlanta: ASHRAE.
5. ASHRAE. 2024. “Design Considerations for Datacom Equipment Centers. 15 Energy Efficiency.” ASHRAE TC 9.9 Datacom Encyclopedia. <https://datacom.ashrae.org>
6. Patterson, M., S. Poole, H. Chung-Hsing, D. Maxwell, et al. 2013. “TUE, a new energy-efficiency metric applied at ORNL's Jaguar.” *28th International Supercomputing Conference, ISC 2013*. https://doi.org/10.1007/978-3-642-38750-0_28
7. ASHRAE. 2014. *PUE: A Comprehensive Examination of the Metric*. Atlanta: ASHRAE.
8. ASHRAE. 2019. *Advancing DCIM with IT Equipment Integration*. Atlanta: ASHRAE.
9. Heydari, A., B. Eslami, V. Radmard, F. Rebarber, et al. 2022. “Power usage effectiveness analysis of a high-density air-liquid hybrid cooled data center.” *ASME 2022 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems*. ASME IPACK2022-97447.
10. European Union. 2023. “EU 2023/1791.” European Union. <https://tinyurl.com/3j3pvubk>
11. European Union. 2024. “EU 2024/1364.” European Union. <https://tinyurl.com/bpwuzsda>
12. Fogle, R. 2024. “How to Measure Server Efficiency with SERT.” U.S. Environmental Protection Agency. <https://tinyurl.com/4fxzcf9j>
13. ASHRAE. 2015. *Server Efficiency: Metrics for Computer Servers and Storage*. Atlanta: ASHRAE. ■