Life Is On | Schneider Electric

# 10 Ways to Harness the Energy and Water Efficiencies of Direct Liquid Cooling

White Paper 211

Version 1

## Data Center Research & Strategy

by Maria A. Torres Arango
Wendy Torell
Stefano Lena
David McGlocklin

## Executive summary

As liquid cooling becomes more prevalent, there are growing uncertainties with how its implementation impacts sustainability. This paper explores 7 design factors and 3 operational choices for single-phase direct-to-chip liquid-cooled systems that affect energy use and water use as well as greenhouse gas (GHG) emissions and toxicity. We provide practical guidance to improve sustainability, supported by quantitative analysis, and discuss how future trends may further impact liquid cooling's sustainability.

RATE THIS PAPER ★★★★★

# Key takeaways

1. **Compared to air cooling, direct-to-chip liquid cooling reduces energy by 30% to 60% *without* the need for adiabatic cooling (water use).** Liquid cooling can save significant energy and water, depending on the operating temperature of the liquid cooling system and the optimization of the heat rejection architecture. Since energy and water use are key contributors to sustainability, liquid cooling aligns well with sustainability objectives.

2. **Liquid-to-liquid CDUs with dedicated high-temperature air-cooled chillers will maximize efficiency.** These design choices allow for maximum economizer hours and reduced mechanical cooling compared to shared and traditional systems.

3. **The coolant supply temperature to the IT equipment (TCS supply) is a significant sustainability driver.** Increasing the operating temperatures in the technology cooling system (TCS) loop by 20°C reduces energy use by 30% to 40%. For adiabatic heat rejection, it reduces water use by 40% to 60%. But it's essential to find the ideal operating temperature, balancing energy efficiency gains with business and downtime risks. Analysis showed diminishing energy reductions per degree at higher temperatures for a data center located in Paris.

4. **IT application primarily determines rack density, with higher densities typically leading to improved heat capture ratio and reduced energy consumption.** In our analysis, a data center with rack densities of 120 kW used 4% to 16% less energy than one with rack densities of 20 kW. These savings are driven by the assumption that the heat capture ratio increases at higher densities.

5. **Liquid-cooled data centers make heat re-use feasible and practical.** Air-cooled data centers expel low-grade heat, making it inefficient and often impractical to transfer to other applications. Liquid cooling, on the other hand, could expel higher-grade heat, which is far more useful for applications such as district heating or industrial processes. Significant challenges still exist for broad adoption, but this is a key enabler.

6. **Component selection impacts efficiency, lifespan, and overall carbon footprint.** Prioritizing components with high efficiency, a long life expectancy, and material compatibility helps reduce energy consumption and embodied carbon associated with manufacturing, spare parts, and system replacements. Vendors that offer transparent product environmental profile (PEP) documents – also known as environmental product disclosures (EPD) – enable informed comparisons.

7. **The operational sustainability of a liquid-cooled data center depends on an intelligent management and control scheme (software and hardware), proactive maintenance, and the ability to adapt to evolving IT demands.** A sustainable system must be able to respond dynamically to changing conditions while being simple to maintain and flexible enough to support future hardware.

# Introduction

As data centers increasingly deploy AI, liquid-cooled servers are becoming a necessity, not an option. This shift is driven by the escalating thermal design power (TDP) of next-generation GPUs and accelerators. In our Executive Report 3, *Optimizing AI Infrastructure: The Critical Role of Liquid Cooling*, we elaborate on this shift to liquid cooled servers in AI data centers, and discuss the importance of careful planning and a strong ecosystem of partners for successful deployment.

The sustainability impacts of liquid cooling are not as well understood as those of air cooling. This is because air cooling has been the primary cooling technology in data centers for decades. However, **liquid cooling *is* more sustainable than air cooling** (see **sidebar**).[1] The broader adoption of liquid cooling necessitates a more comprehensive examination of the design and operational levers that impact sustainability. Four main categories of sustainability should be considered:

- **Energy use** – The cooling system is generally the second-biggest data center energy consumer, next to the IT systems. As such, energy savings can be significant, making this a primary sustainability consideration for liquid cooling deployments. Energy use is driven by several design and operational choices, ranging from foundational decisions like location selection (which determines climatic conditions) to more specific design and operational factors like heat rejection architecture and IT inlet fluid temperature set points.
- **Water use** – Water use can vary significantly based on factors like heat rejection system design (e.g., evaporative cooling, adiabatic cooling), climate, and IT inlet fluid temperature. Since the liquid cooling loop itself is a closed-loop system, there is *no direct water consumption within the racks*. The broader data center's water footprint is still impacted by how that heat is ultimately rejected to the outdoors. In fact, because liquid cooling can operate at higher temperatures, the use of outdoor heat rejection water is reduced. Like energy use, water consumption is a direct, measurable environmental impact.
- **Greenhouse gas (GHG) emissions** – While the carbon footprint is influenced by decisions made throughout the cooling system's lifecycle (e.g., component selection, water use[2]), it is *primarily* driven by energy consumption (more so, if the grid relies heavily on fossil fuels).
- **Toxicity** – Different coolants have different potential for harmful effects on human health and the environment. While direct-to-chip liquid cooling is deployed as a closed-loop system, this does not eliminate the risk of exposure during maintenance, accidental leaks, or end-of-life disposal. Considering coolant toxicity is, therefore, a factor for safe operations and disposal.

As **Figure 1** illustrates, **energy and water** use represent the most significant and direct, measurable environmental impacts of liquid cooling in data centers. While GHG emissions and toxicity are also critical environmental considerations, their mitigation often stems from improvements in energy and water use or responsible material handling.

## Liquid cooling is more sustainable than air cooling

Some GPU-based servers will come in both air-cooled and liquid-cooled options. For small-scale, retrofit cases, where densities under 50 kW/rack are acceptable, the air-cooled option may provide simpler implementation. However, liquid cooling provides several sustainability benefits:

- **Liquid cooling can save 30% to 60% of energy** consumed, depending on liquid cooling operating temperature and optimization of the heat rejection architecture.
- **Liquid cooling can avoid the use of water altogether.** This is because the gains in energy savings by using liquid cooling significantly outweigh the savings from the use of adiabatic air cooling.
- **Savings in energy and water directly translate into carbon footprint savings.** The amount largely depends on the electricity emission factor of the data center location.
- **Liquid cooling allows for much smaller server footprint and therefore greater rack density.** For example, a liquid-cooled server that requires a 2U chassis could require 10-12U for the equivalent air cooled server. A smaller footprint and the resulting reduced rack count translates to a lower embodied carbon.

[1] Note, while the focus of this paper is on situations where liquid cooling is required, there is sometimes discussion around air vs. liquid in terms of sustainability differences. The sidebar on this page summarizes the key sustainability benefits of liquid cooling over air cooling.

[2] While evaporated water vapor from heat rejection has a negligible direct global warming impact compared to CO2 from energy use, the water's lifecycle itself carries an embodied carbon footprint due to energy used for pumping and treatment.

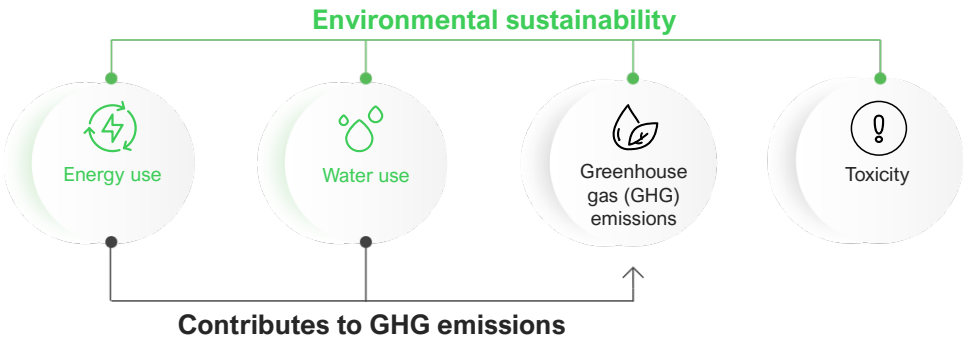



**Figure 1**

*Liquid cooling impacts four sustainability categories. Energy use and water use are the primary categories.*

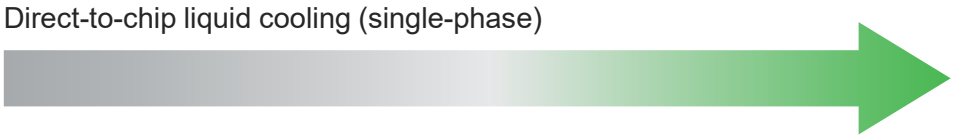
To make informed decisions, data center planners and operators must understand the nuances of liquid cooling and its impact on sustainability. These decisions should align with an organization's sustainability goals, priorities, and any regulatory compliance requirements. For example, in water-stressed regions, minimizing water usage might take precedence over energy consumption. Likewise, in a location where electricity has a high carbon intensity and water is plentiful, energy reduction may take precedence over water use.

There are multiple metrics to track performance against these categories. In White Paper 67, *Guide to Environmental Sustainability Metrics for Data Centers*, we describe these metrics and provide recommendations for data centers at various stages of their sustainability journey.

## Key liquid cooling concepts

There are multiple ways to implement liquid cooling. In White Paper 265, *Liquid Cooling Technologies for Data Centers and Edge Applications*, we discuss the various liquid cooling methods, and in White Paper 133, *Navigating Liquid Cooling Architectures for Data Centers with AI Workloads*, we explore how these methods are applied in data centers.

The two broad liquid cooling categories are direct-to-chip and immersive. Single-phase[3] direct-to-chip is most prevalent in the industry today. As such, that's the focus of this paper. Figure 2 summarizes key concepts for this architecture.




**Heat capture within the server**

A cold plate, in direct contact with the server component (like silicon chips), transfers heat. The cold plate replaces air-cooled heat sinks.

**Heat exchange inside the data center**

A CDU pump circulates coolant through the cold plate. The component's heat is then carried to a CDU which transfers the heat to the outdoor rejection system.

**Heat rejection to outdoor ambient**

The heat is then rejected to the outdoor air through the heat rejection system's cooling coils (e.g., a chiller).

**Figure 2**

*Three main heat transfer elements in single-phase direct-to-chip liquid cooling*

In this paper, we present design factors and operational choices associated with single-phase direct-to-chip liquid-cooled systems that impact environmental sustainability (with specific focus on energy and water use). We then provide practical guidance on how to reduce the consumption of these resources. Last, we discuss how future trends may affect liquid cooling's impact on sustainability.

[3] Single-phase refers to the coolant remaining in a liquid state throughout the entire cooling process.

# 7 design factors impacting sustainability

Decisions made during a liquid cooling system's design directly influence its sustainability performance. A comprehensive view of the liquid cooling architecture, from chip to chiller, is critical because component interdependencies impact sustainability.

We present **seven core design factors** that influence your liquid cooling deployment's sustainability (**Figure 3**). For each, we define the choice and illustrate its potential impact on relevant sustainability categories with practical use case examples and guidance. While our focus is on levers affecting sustainability, it's important to recognize that these decisions also impact other crucial factors, including cost and availability.

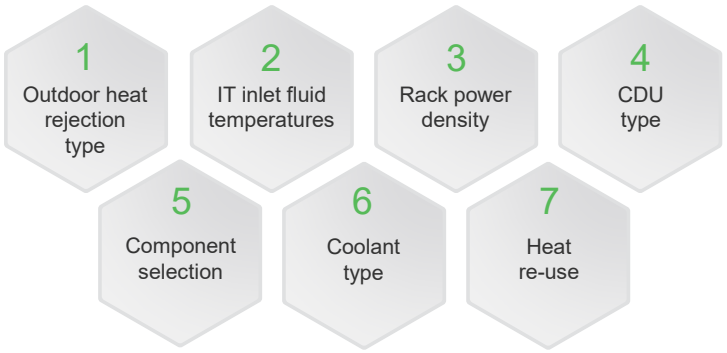



**Figure 3**

*Seven key design factors impacting sustainability*

**Figure 4** is a simplified diagram of a single-phase direct-to-chip liquid cooling design that utilizes liquid-to-liquid coolant distribution units (CDUs). It shows two main elements that represent the entire cooling system from chip to chiller:

- **Technology cooling system (TCS)** – This includes cold plates within the IT equipment, piping, manifolds, CDUs, and pumps, all in the white space or technical corridors. This is a closed-loop system.
- **Facility water system (FWS)** – This includes pumps, piping, and the outdoor heat rejection system. Choices in the IT space including operating coolant temperature and CDU specifications impact performance of the FWS's energy and water consumption.

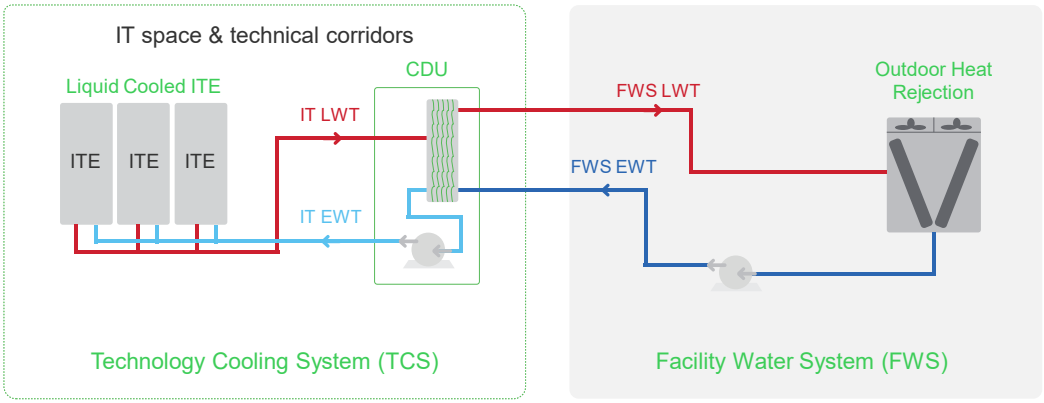



**Figure 4**

*Representative liquid cooling architecture illustrating key elements of direct-to-chip liquid cooling.*

*ITE: Information technology equipment; EWT: Entering water temperature; LWT: Leaving water temperature*

## Design factor 1: Outdoor heat rejection type

Although the outdoor heat rejection system is not necessarily different for air-cooled versus liquid-cooled IT loads, it drives the facility's energy and water consumption. Key aspects to consider include the choice of chiller type, whether to opt for adiabatic cooling or not, and whether to deploy a dedicated chiller or shared system.

**Air-cooled chillers with free cooling are often more sustainable than water-cooled chillers**. The choice of outdoor heat rejection technology has a significant impact on the energy and water consumption of a liquid-cooled data center. Air-cooled chillers with free cooling (economizer coils) represent a sustainable approach by employing a closed-loop system that eliminates the need for continuous water evaporation. These systems utilize ambient air to reject heat, much like a car's radiator. As ambient air temperatures drop, these chillers rely less on their energy-intensive compressors to reject the heat. Instead, they run in economizer mode.

**In contrast, water-cooled chillers paired with cooling towers are highly water intensive.** Although typically more energy efficient, they rely on the continuous evaporation of water to reject heat, consuming significant volumes of water year-round. This water consumption also requires pumping and treatment, adding to the overall energy footprint. In fact, as we discuss in White Paper 132, *Economizer Modes of Data Center Cooling Systems*, cooling towers, on average, consume approximately 43 liters per minute for every 1 MW of cooling capacity. This translates to 22,629 cubic meters per MW per year. High water usage has a negative environmental impact that extends beyond the consumption of a scarce resource; it also impacts a data center's carbon footprint. Every cubic meter of water consumed generates an average of 23 lbs (10.6 kg) of carbon emissions from the energy required for its pumping, distribution, and treatment.[4]

**The benefit of adiabatic cooling depends on location, resources, and priorities**. This optional feature is used to increase the energy efficiency of air-cooled heat rejection systems, particularly in warm, dry climates. This process involves misting water into the airflow entering the condenser. As the water evaporates, it cools the air. This cooler air removes more heat from the refrigerant, thereby reducing the compressor's energy consumption. This approach allows for more economizer hours, especially during hot periods. **Unlike cooling towers, adiabatic systems use water only as needed to boost performance**. This strategic use of water makes it a more water-efficient option than a traditional cooling tower, as it is only activated when ambient temperatures exceed a specific threshold. White Paper 227, *How Higher Chilled Water Temperature Can Improve Data Center Cooling System Efficiency*, discusses how adiabatic cooling is used and quantifies its benefit.

**A dedicated high-temperature chiller can be optimized to minimize energy AND water use**. A dedicated system can be tailored to the specific requirements of liquid-cooled IT equipment, enabling a more efficient design. White Paper 227 discusses and quantifies the benefit of operating at higher temperatures. In contrast, a shared system that serves both air-cooled and liquid-cooled loads, introduces complexities. These different cooling loads often require different temperatures. For example, in a retrofit case, a system designed to support new liquid-cooled servers (which can operate at higher chilled-water temperatures) along with the existing air-cooled servers (which require lower chilled-water temperatures) must be configured for the lowest chilled-water temperature. This often results in the entire system running less efficiently and at a colder temperature than what is needed for the data center's liquid-cooled infrastructure. This leads to suboptimal energy use compared to a dedicated, purpose-built system. See White Paper 210, *Direct Liquid Cooling System Challenges in Data Centers*, for further discussion on the implications of shared versus dedicated cooling systems.

[4] Water Intelligence, *The carbon footprint of water*

## Design factor 2: IT inlet fluid temperatures

Liquid-cooled IT devices have specific IT inlet temperature requirements necessary for effective heat removal. These conditions are specified by the IT vendor as ranges. For example, the Dell PowerEdge XE9640 Installation and Service Manual provides guidance that "we require 32°C primary water temperature, but it may be possible to go higher under certain circumstances".[5] Another example is the Lenovo ThinkSystem SR665 V3, which specifies water inlet temperatures of 40°C, 45°C, or 50°C, depending on the flow rate.[6] Staying below the specified maximum is generally necessary for proper IT performance and to maintain the warranty. In addition, service level agreements (SLAs) are stacking up with compounding requirements, leading to conservative lower operating temperatures.

Within this specified range, a higher IT inlet temperature reduces the cooling plant's energy consumption, driven by decreased compressor energy. This also increases the potential for valuable waste heat reuse (discussed in Design factor 7: Heat re-use), which becomes more viable at higher temperatures.

The chip's heat density (watts/cm$^2$) and TDP (watts) directly influence the cooling requirements specified by the IT vendor. When chips generate more heat in a smaller area, they necessitate a lower temperature at the cold plate interface to transfer that heat away efficiently. A lower temperature at the cold plate interface, in turn, requires a lower cooling water temperature throughout the liquid cooling loop. Figure 5 illustrates ASHRAE's trend toward lower IT supply temperatures as TDP rises. This creates a challenge when aligning to sustainability objectives, as energy increases at these lower temperatures.

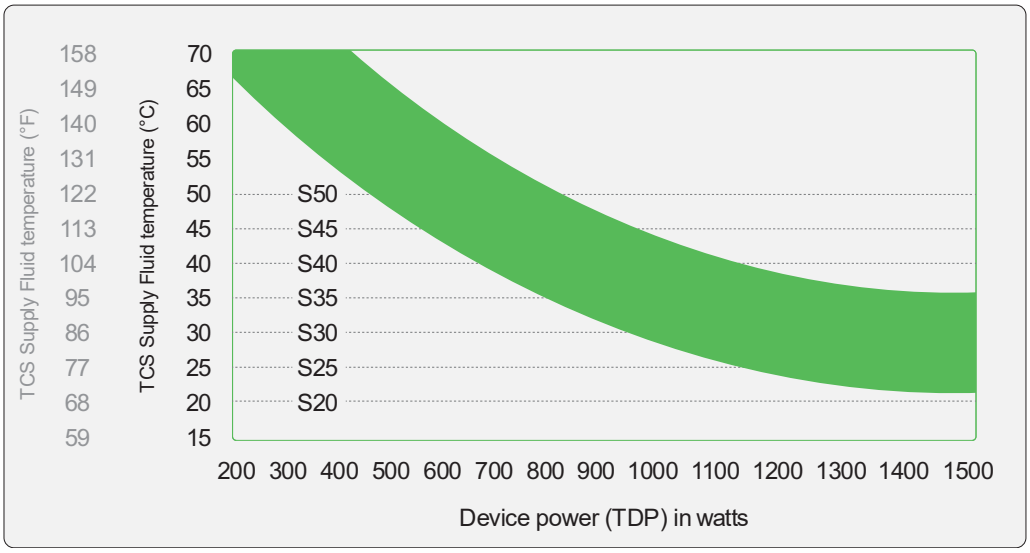



**Figure 5**

*Trend line of TCS supply temperature (IT inlet) as TDP rises*

*Source: ASHRAE*

Vendor collaboration is helping to address these challenges. Organizations such as the Open Compute Project (OCP) are defining a long-term roadmap to address topics including reliability and energy efficiency. Their paper emphasizes setting a minimum temperature for a data center's cooling system is essential for viable, long-term infrastructure investment.[7] Through this work, vendors are aligning on 30°C as the target **minimum fluid temperature**, highlighting a shared commitment to advancing technologies such as silicon packaging to maintain this temperature as an industry standard.

[5] Dell, *PowerEdge XE9640 Installation and Service Manual*

[6] Lenovo, *ThinkSystem SR665 V3 System Configuration Guide*, page 12

[7] OCP, *30°C Coolant – A Durable Roadmap for the Future*

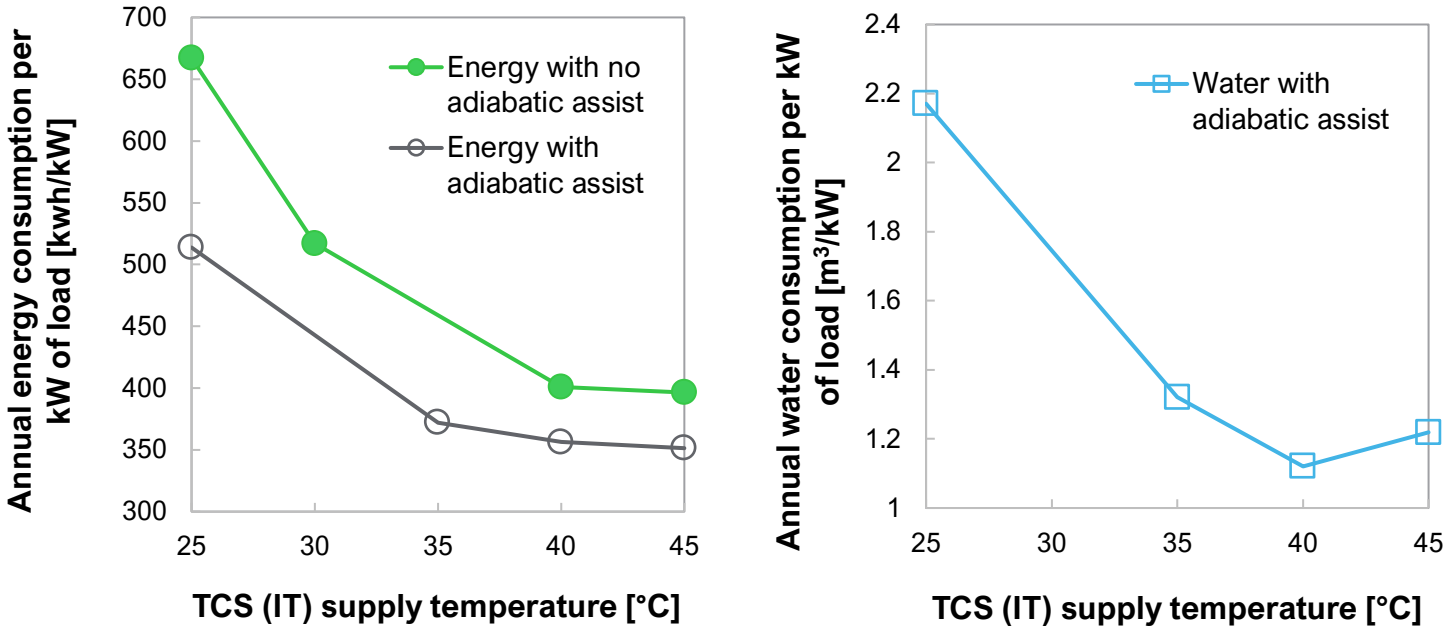
**Risk tolerance for the application should dictate the operating setpoint**. There is no one right answer to what fluid temperature to operate at within the device's specified range. From a sustainability perspective, allowing higher fluid temperatures can lead to significant energy savings.

Conversely, opting for lower fluid temperatures provides a thermal safety buffer from **overshoot (instances where the coolant delivered temporarily exceeds its desired or setpoint temperature)**. This overshoot may be caused, in part, by highly transient power demands. For expensive AI servers, this overshoot can trigger performance throttling or emergency shutdowns. Given that any downtime or degraded performance can represent a business risk, higher energy and water use may be a justifiable trade-off based on your reliability and performance requirements.

We performed simulations[8] to quantify the impact of temperature on energy and water use for a hypothetical data center in Paris, France. Figure 6 illustrates the correlation of increasing TCS inlet temperature and decreasing energy and water consumption. However, for both non-adiabatic and adiabatic cooling, raising the temperature above a certain point provides diminishing returns. While a 20°C increase (from 25°C to 45°C) yields a substantial overall energy reduction, the most significant savings are realized in the lower temperature increments.

**Figure 6**

*Simulation analysis conducted to quantify the impact of temperature on energy (left) and water (right). Data shown corresponds to a 40 kW/rack configuration in Paris, France. Source: Schneider Electric Innovation and Strategy Cooling Team*




[8] Only simulation data points validated by Schneider Electric cooling experts were included in the results.

The location of the data center impacts the shape of this curve. The curve would be shifted higher (resulting in more energy use) for higher outdoor design temperatures, and with a modified slope based on the point where free cooling kicks in. In Appendix Figure A2, we demonstrate the increased energy in Dubai compared to Paris, and depict the steeper slope between 30°C and 40°C. Also, significant gains would be obtainable, even in the higher IT supply temperature range, due to further utilization of free cooling. In contrast, in a cooler climate, where free cooling can be achieved year-round, the curve would shift downward (reducing energy consumption), and the point of diminishing returns would occur at a lower IT supply temperature.

For the simulated data center in Paris at 40 kW/rack, with non-adiabatic cooling, a 20°C temperature increase leads to an approximate 40% reduction in energy. The bulk of this savings comes from the first 5°C increase (from 25°C to 30°C), which accounts for approximately half of the total reduction. In contrast, the final 5°C increase (from 40°C to 45°C) provides a minimal 1.1% reduction. This suggests that the energy savings from increasing the temperature beyond a certain threshold may not justify the associated risks to reliability and downtime.

With adiabatic cooling, where water consumption is introduced for energy savings, a 20°C TCS temperature increase results in an approximate 30% reduction in energy, alongside a significant 40% to 60% reduction in water usage. The principle of diminishing returns still applies, with a decrease in savings per degree as temperatures rise. Again, the location impacts the magnitude of the energy and water savings.

Let's look at the energy and water cost for this adiabatic cooling system assuming a 10 MW Paris data center with an electricity cost of €0.18 per kWh and a water cost of €4.10 per $m^3$.

- Increasing the TCS temperature from 25°C to 45°C translates to an energy savings of €292k and a water savings of €39k. Of these savings:
- €255k of energy and €35k of water savings are achieved by increasing the temperature from 25°C to 35°C
- €37k of energy and €4k of water savings are realized when increasing the temperature from 35°C to 45°C

Given the trade-off of performance improvement and risk, the general **best practice is to start with a conservative (lower) operating temperature within the specified range and not to run at the maximum temperature**. This provides a thermal buffer from the start. Over time, you can evaluate the system's performance and measure any temperature overshoot that occurs due to the workload's transient power demands. With this data, you can then gradually and safely raise the fluid temperature in a controlled fashion to a higher, more energy-efficient setpoint that still ensures reliability without risking performance throttling or unplanned shutdowns.

Collaboration with IT and physical infrastructure vendors will help refine your design set points. Schneider Electric created reference designs in collaboration with NVIDIA to support their latest GPUs and SuperPODs. These provide a blueprint for a reliable and validated AI deployment. For example, the reference design for NVIDIA GB300 NVL72[9] specifies a TCS supply temperature of 40°C.

[9] Schneider Electric, *Reference Design 110* (IEC) and *Reference Design 111* (ANSI) for NVIDIA DGX GB300

## Design factor 3: Rack power density

Rack power density is primarily driven by the AI application. AI training models face latency constraints that drive the number of accelerators in a rack, leading to higher rack densities in AI clusters. Inference models are often more compressed, therefore rely on far fewer accelerators, and allow a lower density. See White Paper 110, *How 6 AI Attributes Change Data Center Design*, for more details on accelerator network communication and its impact on density.

The question is how does this rack density influence sustainability of the cooling plant? Figure 7 shows a slight improvement in energy consumption as density increases from 20 kW/rack to 120 kW/rack, for both non-adiabatic and adiabatic chiller design. We attribute this improvement to the increasing heat capture ratio (percentage of cooling by liquid vs. air) with increasing rack densities. This relationship is generally true because as chip TDPs go up, they become a greater portion of the total server power. Since liquid cooling is more efficient than air cooling, this results in a *decrease* in overall cooling plant energy.

Additionally, although *not quantified* in the analysis, there are other considerations related to pump power that correlate with density. At higher rack density, greater cooling capacity is needed in each rack. With higher rack capacity, the required coolant flow rate often increases, necessitating a corresponding *increase* in pump power for a given pipe size (diameter and length). In contrast, at higher rack densities, the footprint can be smaller for a given total data center capacity, and therefore, pipe lengths can be shorter. This translates to a *decrease* in pump power. Whether a data center sees an increase or decrease in pump power depends on several factors, including whether it's a retrofit or new build, the added capacity, and changes to the pipe design.

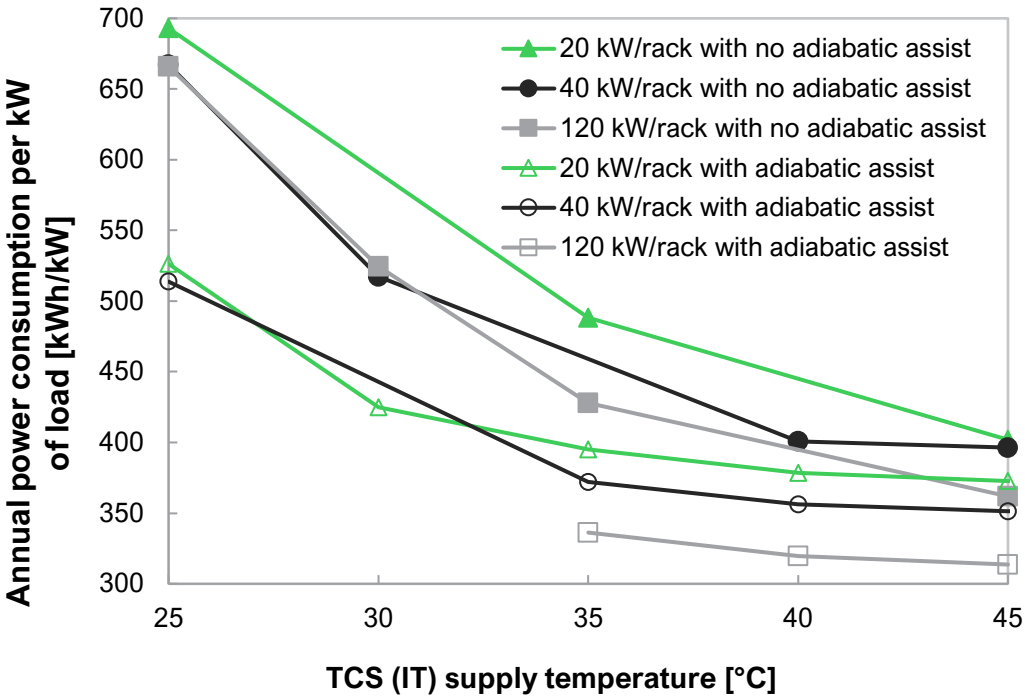



**Figure 7**

*Energy simulation analysis conducted to quantify the impact of rack density, for both adiabatic and non-adiabatic chiller design.*

*Source: Schneider Electric Innovation and Strategy Cooling Team*

The net result is that energy does improve with rack density, but it is not a significant driver. As previously stated, the IT application is most indicative of the density deployed. Density does, however, have an impact on other factors that influence the carbon footprint. For instance, higher rack density means a lower embodied carbon from piping and core & shell per kW of load.

## Design factor 4: CDU type

The type of CDU has a significant impact on a data center's sustainability. Two key considerations are:

- choosing between liquid-to-liquid (L-to-L) and liquid-to-air (L-to-A) CDUs
- choosing the best CDU form factor

**L-to-L CDUs are more energy efficient**. They transfer heat from the TCS loop to the FWS loop. The FWS loop then transfers the heat to the chiller plant. This method is more energy-efficient at removing heat, making it the preferred choice for new builds in high-density environments. **L-to-A CDUs are less energy efficient**. They reject heat from the TCS loop directly to the data center air. However, it requires an additional heat exchange from the room's CRAHs/CRACs[10] to the chiller plant/condenser. This approach is often used in smaller-scale deployments or retrofits. L-A CDUs don't require a separate piping connection to a chilled water plant, simplifying the plumbing and installation.

**CDU placement, whether in-rack or centralized (floor mount), also impacts sustainability**. Floor-mount CDUs offer placement versatility, as they could be in the IT space or service corridors. However, the further they are placed from the IT equipment, the more pumping power they require. Rack-mounted CDUs, which are dedicated to a single rack, offer the sustainability advantage of reduced TCS pumping energy by minimizing coolant travel distance for small scale deployments. However, they become less energy-efficient as a deployment scales up, since operating many individual CDU pumps is less efficient than a single, larger system (i.e., floor-mount CDU). In addition, a rack-mounted CDU must deliver high capacity in a very small space (inside the rack). Therefore, the heat exchanger in this CDU is smaller and works at a higher approach temperature compared to a floor mount CDU, lowering the efficiency of the cooling system. Rack-mounted CDUs also add complexity in the maintenance needed during operation, as there are now many more TCS coolant loops that need to be monitored and put in a maintenance schedule. This impacts the embodied carbon of the system. White Paper 133, *Navigating Liquid Cooling Architectures for Data Centers with AI Workloads*, provides more depth into the trade-offs of form factor and CDU type.

## Design factor 5: Component selection

Component selection should consider these key sustainability factors: efficiency, lifecycle, quality, and material compatibility. These contribute to the data center's total lifecycle carbon footprint as we describe next.

The **efficiency of liquid cooling components** like CDUs and chillers influence how much energy is consumed. This is often specified as a Coefficient of Performance.[11] Component efficiency can be a significant GHG controlling lever.[12] For the CDU[13], the heat exchanger and the pump(s) are two key components that drive its efficiency.

[10] CRAH stands for computer room air handler (used with chiller cooling architectures), and CRAC stands for computer room air conditioner (used with direction expansion or DX cooling architectures).

[11] The Coefficient of Performance (COP) informs how the component's cooling capacity compares to its energy consumption. The larger the COP, the better the performance and the lower the energy impact.

[12] Depending on the local electricity emissions factor.

[13] Your CDU selection must also include the **CDU net cooling capacity** and **TCS external pump head available**, which ultimately determine the amount of IT equipment that can be cooled at the CDU rated TCS flow rate.

- **Heat exchanger** – In the CDU domain, the approach temperature is the difference between the temperature of the coolant leaving the CDU (entering the IT servers), and the temperature of the facility's chilled water entering the CDU from the FWS loop. For a given capacity, the lower the CDU approach temperature, the higher you can specify the chilled water temperature. Higher chilled water temperature increases economizer hours and reduces chiller compressor energy, provided the chiller is dedicated to your liquid cooling load and can be optimized. State of the art CDU approach temperatures range from 2˚C (3.6˚F) to 4˚C (7.2˚F). Note, it's important that CDU capacity and approach temperature are always considered together. Operating the same heat exchanger at a higher CDU approach temperature would provide higher cooling capacity.
- **Pump(s)** – The pump efficiency is the ratio of the hydraulic power delivered to the electrical power consumed by the pump. More efficient pumps reduce overall energy consumption. Selecting a CDU properly sized for the specific thermal load prevents the use of oversized pumps that waste energy.

**Life expectancy of components** is another lever on lifecycle $CO_2$. The reliability of parts heavily influences the lifespan of a system. Robust, high-quality components with longer life expectancies reduce the frequency of replacements. This impacts lifecycle $CO_2$, as it minimizes the embodied carbon associated with manufacturing, transportation, and disposing of replacement parts.

**Quality, type, and compatibility of materials** of cooling components, like the CDU, TCS piping, and cold plates, also influence overall sustainability. This is because high-quality and compatible materials ensure less contamination of the TCS and better cooling performance. This leads to less frequent coolant replacement, and/or prevention of premature failure of TCS and IT components. Premature failure of components resulting in equipment replacement contribute to increased Scope 3 GHG emissions and increased expenses. The choice of materials also impacts Scope 3 emissions. Using materials with low embodied emissions, such as green metals,[14] reduces the carbon footprint from manufacturing. Furthermore, selecting materials that are easily recyclable or reusable at the end of their life promotes a circular economy, minimizing waste and the energy needed to create new components. Look for vendors that are transparent about the lifecycle embodied carbon of their products and make product environmental profile (PEP) documents available. These documents are also referred to as environmental product declarations (EPDs).

While we have focused on the CDU in the discussion above, the entire cooling system design from the chip to the chiller determines the overall cooling plant efficiency. You should assess the total annual energy consumption against alternative cooling designs for a given cooling capacity.

## Design factor 6: Coolant type

The choice of coolant for a single-phase direct-to-chip liquid cooling system involves three primary considerations: **thermal performance, operational simplicity, and vendor requirements**. The two main options for coolants in the TCS loop include deionized water and propylene glycol (PG) solution:

- **Deionized water** – Deionized water offers the highest thermal capacity and conductivity, making it the most efficient coolant. However, it requires additives to inhibit biogrowth and corrosion, which presents challenges in maintaining the proper water chemistry. It therefore requires ongoing, precise

[14] https://en.wikipedia.org/wiki/Green_metals

management of its additives. As a result, it is not commonly specified as an approved coolant by IT vendors.

- Propylene glycol (PG) solution – PG is a common coolant mixture. While various concentrations can be used depending on the application, a 25% propylene glycol to 75% deionized water ratio (PG25) is the most common industry standard for the TCS loop. Due to its lower specific heat capacity, higher viscosity, and lower thermal conductivity compared to deionized water, a PG25 solution requires a higher mass flow rate and greater pump power to achieve the same heat rejection. Although PG solutions are less efficient than pure deionized water, they are pre-mixed solutions that contain the necessary corrosion and biogrowth inhibitors.[15] PG25 has been demonstrated to be biostatic[16]. The pre-formulated PG solution is designed to be a stable, simple, all-in-one fluid that provides long-term protection and reduces the maintenance burden. Note, mixing of PG25 solutions from different vendors is not allowed and will likely result in voidance of warranties. This is especially critical during commissioning and/or addition of components to the liquid cooling systems.

Neither deionized water nor propylene glycol contributes to global warming as a direct greenhouse gas. The environmental impact associated with their use in a data center is related to the energy consumption of the cooling system and the manufacturing process of the fluids themselves.

In short, while deionized water offers superior thermal performance, its need for constant chemical management and the associated risks make it a less practical choice. PG25 has become the industry standard mainly due to its biostatic nature[17], making it the most vendor-approved option for protecting sensitive IT equipment. Always consult the documentation provided by the server manufacturer (e.g., Dell, HPE, Supermicro) and, by extension, the chip manufacturer (Intel, AMD, NVIDIA) if their guidelines are separate. Failing to adhere to their coolant specifications carries risks, including hardware damage, system failure, and voided warranties. OCP also provides guidelines on deploying PG25 solutions in the TCS loop.[18]

## Design factor 7: Heat re-use

Because the shift from air to liquid cooling allows for higher coolant temperatures, we can recover heat more efficiently and cost-effectively. Governments and other stakeholders are encouraging data centers to capture and repurpose waste heat streams from their liquid cooling loops for uses like district heating and agriculture.[19] This turns a wasted byproduct into a valued asset, improving the data center's energy-and-carbon-neutrality profile.

Repurposing waste heat reduces the energy and carbon emissions that would have otherwise likely come from fossil fuel sources. The industry uses the metric "energy reuse factor (ERF)," the ratio of reused energy to total data center energy consumption. Regulations and government incentives are also expected to increase, driving further adoption of heat re-use designs.

---

[15] PG25 also benefits from anti-freeze protection, which is necessary for the FWS loops, but not for the TCS loop. This is because it is located indoors and the temperature inside a data center is carefully controlled and kept well above freezing, so there's no risk of the coolant freezing.

[16] Duggan et al., *Reviewing the Evidence of Antimicrobial Activity of Glycols*, Apr. 2024, Journal of Applied Microbiology

[17] ASHRAE, *Liquid Cooling Guidelines for Datacom Equipment Centers*, 3rd Edition, Chapter 7

[18] Open Compute, *Guidelines for Using Propylene Glycol Based Heat Transfer Fluids in Single Phase Cold Plate-Based Liquid Cooled Racks*

[19] Yuan et al., *Waste Heat Recoveries in Data Centers: A Review*, Oct. 2023, Renewable and Sustainable Energy Reviews

# 3 operational practices impacting sustainability

Beyond the cooling system design, the day-to-day operational practices of a data center's liquid cooling system influence its sustainability. The goal is to maximize energy efficiency and minimize environmental impact throughout the data center use phase. This section explores three key operational practices (**Figure 8**) that directly affect sustainability outcomes:

**Figure 8**

*Three operational practices impacting sustainability*




## Operational practice 1: Management and control scheme

Pipes within the IT server racks are new for most IT professionals, bringing uncertainty on how to control and manage an unfamiliar set of components in their cooling architecture. Therefore, **control of the CDUs which directly feed the servers must be simple** to reduce risk to operations, but also to optimize sustainability performance. Having an intuitive interface to display key metrics like fluid temperature, flow rate, and pressure, presented as clear visuals, provides IT staff with immediate feedback without requiring specialized liquid cooling expertise.

Beyond on-board CDU controls, it is important to consider the broader cooling plant. Interoperability between various management systems delivers communication and coordination of response. This allows visualization and tracking of resource use for a more efficient operation. Systems to consider integrating include:

- building management system (BMS)
- energy and power management software (EPMS)
- supervisory control and data acquisition software (SCADA)
- data center infrastructure management system (DCIM)
- IT cluster management software (e.g., NVIDIA mission control software)

In large AI factories, effectively managing a liquid-cooled data center requires a coordinated, multi-layered control system that prioritizes energy efficiency. This system, often centered on a PLC (programmable logic controller) manager, acts as a central brain to ensure operations are both resilient and efficient. Controls Reference Design 1, *AI Infrastructure and Liquid-cooled Controls Design*, provides an example blueprint to an effective control system design, including how to coordinate this management in a hybrid air-cooled/liquid-cooled environment.

A crucial aspect of this involves creating **a digital twin** of the cooling system. This virtual model enables operators to make informed, sustainable decisions throughout every phase, from planning and design to real-time operations. This digital twin provides real-time insights into temperature fluctuations, coolant flow rates, and energy consumption, giving operators a comprehensive view of the system's performance. The ability to dynamically optimize these parameters directly leads to an improved power usage effectiveness (PUE) and water usage effectiveness (WUE), which are key metrics for measuring sustainability.

As the adoption of liquid cooling continues to scale, we expect further advancements, standardization, and convergence on best practices for the management and control scheme. For example, in the early days of liquid cooling, pumps are run at constant speed, but with growing confidence, this granular data allows for analysis and dynamic control. Thus, you can dynamically adjust the operating parameters based on the real-time compute load. This precise tuning of cooling would ensure that no pump energy is spent on cooling idle or low-load servers (wasted energy).[20]

## Operational practice 2: Maintenance practices

To ensure optimal performance and sustainability of a direct liquid cooling system's TCS loop, consistent maintenance is essential. Specifically, these maintenance practices should include the following:

- Coolant quality monitoring – The best practice is to take a baseline snapshot on day zero and conduct frequent measurements during the initial weeks of operation. This tracks the stabilization of the fluid's chemistry, including its pH, conductivity, and turbidity (indicative of debris/particles suspended in the coolant). Once stability is achieved, continuous monitoring for any deviation or drift is necessary, especially when new equipment is integrated. A sudden change in fluid chemistry, including a rise in turbidity, can indicate contamination or corrosion and should be investigated immediately.[17] By maintaining optimal coolant quality, you not only prevent equipment damage but also extend the life of the fluid and components. This minimizes the environmental impact associated with frequent fluid replacement and hardware disposal.
- Leak detection and inspection – Regularly inspect all lines, connections, and components for any signs of leaks. Early detection prevents damage to expensive IT equipment and reduces the need for coolant replacement. Pressure measurement is the best way to guarantee immediate detection of suspect leaks, followed by visual inspection to find the actual leak.
- Filter cleaning and replacement – Filters within the CDU must be regularly inspected and cleaned. It goes without saying that clogged filters restrict coolant flow, forcing pumps to work harder, which consumes more energy and can lead to pump failure.
- Component inspection – Check system components like pumps for proper function and signs of wear. Additionally, use pressure drop versus flow data to identify if the heat exchanger is performing optimally. Monitoring this through software in real time and analyzing historical values helps to detect wear or clogging. Some systems have built-in redundancy, which allows for maintenance on individual components without disrupting the entire system.
- Spare kit availability – It is recommended that parts frequently changed or serviced during liquid cooling maintenance are kept on site to minimize the high carbon impact of overnight rush air delivery. This also prevents operation disruption and added cost.

These practices contribute to sustainability by improving energy efficiency and extending the lifespan of both the cooling infrastructure and the IT hardware it supports. Predictive maintenance, often leveraging AI-powered digital twin models, can further improve sustainability by anticipating potential failures before they occur.

[20] Additionally, having data on the system's performance would allow for prediction of maintenance needs and proactive notification of system errors to facilitate timely correction of issues.

### Operational practice 3: Adapting to evolving IT requirements

Evolving IT with higher TDPs, and a trend towards lower TCS supply temperatures, creates a challenge for data centers. Future-proofing cooling infrastructure often means installing oversized components like pipes and chillers. While this potentially saves energy in the long run, it creates a sustainability paradox by increasing embodied carbon and "stranded capacity" from day one. The central challenge is to scale for future needs without compromising present-day sustainability.

The solution is to adopt a modular and phased approach. Instead of building for maximum capacity upfront, data centers can deploy foundational infrastructure with the ability to scale gradually, such as pouring cement pads for future chillers. This strategy minimizes initial embodied carbon and capital expenditure, allowing the facility to adapt to future chip evolution.

A sustainable path forward requires a fundamental shift in operational practice. By shifting from a "build for maximum capacity" mindset toward a flexible, adaptive, and modular strategy, we can align infrastructure deployment with actual IT demand, delivering efficiency without compromising environmental responsibility. White Paper 210, *Direct Liquid Cooling System Challenges in Data Centers*, addresses these future-proofing considerations in greater detail.

## How future trends in liquid cooling will impact sustainability

The evolution of liquid cooling technology will continue to improve data center sustainability. Future trends are focused on more efficient, intelligent, and integrated systems to manage the growing heat demand while minimizing environmental impact. We discuss future trends for physical infrastructure in White Paper 110, *How 6 AI Attributes Change Data Center Design*. Below are five areas of advancement that we anticipate:

- Coolant – As discussed earlier, PG25 and deionized water are the primary coolants used today. But sustainable dielectrics may replace these if they are able to increase the heat transfer efficiency and allow for higher chip TDPs. These next-generation fluids aim to achieve low global warming potential (GWP), low toxicity, as well as eliminating harmful chemicals like PFAS (per- and polyfluoroalkyl substances)[21].

  On the longer horizon, other technologies such as nanofluids are being investigated to further enhance the fluid's thermal conductivity and heat transfer capabilities. However, to be practical, these technologies will need to overcome challenges for their long-term stability and adoption.
- Two-phase liquid cooling – For future ultra-high compute-density scenarios, we believe direct-to-chip two-phase cooling[22] will be adopted for its high heat transfer efficiency. However, its adoption is currently limited by refrigerants with high GWP and toxicity, in addition to regulatory risks associated with PFAS. As the technology matures, it's important to note that single-phase systems remain highly effective and can support racks over 1 MW.
- Advanced analytics and dynamic cooling – The integration of AI and advanced analytics will enable dynamic cooling optimization. This allows a system to autonomously manage and adjust cooling parameters in real-time, matching cooling output more precisely to the IT load. The result is a further reduction in energy consumption as well as a reduction in manual tasks.

[21] Bonner et al., *Understanding PFAS Concerns for Two-Phase Cooling of Data Centers*, Feb. 2024, Data Center Frontier

[22] A two-phase system utilizes a cooling medium that changes phase from liquid to vapor as it absorbs heat and condenses back into liquid before recirculation. This state change enables it to carry away more heat energy than a single-phase system, at the same facility water temperature.

- **Waste heat recovery** – Future cooling designs will increasingly focus on harnessing waste heat to achieve greater sustainability. Technology developments in waste heat recovery will enable a higher energy reuse factor (ERF), such as pairing liquid cooling with Organic Rankine Cycle (ORC) systems.[23]
- **Server architecture evolution** – There is space for innovation in the server architecture to increase the efficiency of the cooling system. We anticipate servers leveraging 100% liquid cooling (meaning 100% heat capture ratio). This can be achieved through fanless designs, extended cold plate coverage within the server, and photonic[24] components that don't generate heat. These changes reduce the need for energy-intensive air cooling, further boosting the efficiency and sustainability of data centers. Moreover, as obstacles in liquid cooling are overcome and reliability of liquid cooling systems improve, it is likely that liquid cooling becomes the preferred choice from its higher efficiency, even for rack densities where air cooling is still possible.

## Next steps

In this paper, we discussed the main drivers of liquid cooling design and operation that influence data center energy and water consumption. If your upcoming data center project involves liquid cooling, we recommend you take the following next steps:

1. **Engage and plan early with an ecosystem of partners**. Start your cooling architecture planning simultaneously with your IT planning. Collaborate with technology partners, including IT vendors, cooling specialists, systems integrators, and service partners. They can help plan, design, implement, and maintain your liquid cooling architecture, in alignment with your sustainability goals.
2. **Be clear on your sustainability priorities before designing your data center**. Gain clarity from leadership, compliance teams, and sustainability teams on any regulations, governance requirements, and business targets. Consider the specific geographical characteristics of the chosen site, as local regulations, climate, and access to resources will significantly impact design choices. Since some liquid cooling design and operational factors trade off business and downtime risks with sustainability, it is crucial to go into the design fully informed.
3. **Think about day-to-day operational practices during the design phase**. The choices you make in the liquid cooling design, such as coolant selection and management and control scheme, directly impact how sustainable the data center will be.
4. **Stay connected on the continued advances in liquid cooling technologies.** Developments in various areas of direct-to-chip liquid cooling offer great promise in further increasing cooling efficiency and reducing environmental impacts.

In addition to the steps suggested here, White Paper 188, *10 Data Center Sustainability Improvements Quantified by Environmental Impact*, quantifies the carbon footprint of a data center across site selection, design, and operation. Use that paper to identify the most significant contributors to GHG emissions and to assist in optimizing your data center design and operation.

[23] Corigliano et al. *Turning Data Center Waste Heat into Energy: A Guide to Organic Rankine Cycle System Design and Performance Evaluation*, Jul. 2024, Applied Sciences

[24] Speciale Invest, *Photonics: A Deep Dive into the Future of Light-Speed Computing*

## About the authors

**Dr. Maria A. Torres Arango** is a Research Analyst in Schneider Electric's Data Center Research & Strategy group. She investigates technology, materials, and infrastructure to guide data center strategies. Her current work focuses on energy storage and liquid cooling technologies. She also examines market forces driving technology advancement and participates in developing tools that highlight tradeoffs or selection in critical data center decisions. Maria's former expertise includes materials design and optimization, as well as fundamental studies on materials synthesis processes using X-ray characterization at the National Synchrotron Light Source II, Brookhaven National Laboratory. A lifelong learner, Maria holds a BS in aeronautical engineering from Universidad Pontificia Bolivariana, Colombia, and a MSc in aerospace engineering and a PhD in materials science and engineering from West Virginia University.

**Wendy Torell** is a Senior Research Analyst in Schneider Electric's Data Center Research & Strategy group, bringing 30 years of data center experience. Her focus is analyzing and measuring the value of emerging technologies and trends: providing practical, best practice guidance in data center design and operation. Beyond traditional thought leadership, she championed and leads development of interactive, web-based TradeOff Tools. These calculators help clients quantify business decisions, while optimizing their availability, sustainability, and cost of their data center environments. Her deep background in availability science approaches and design practices helps clients meet their current and future data center performance objectives. She brings a wealth of experience across Schneider Electric's broad portfolio and with the market at large. She holds a BS in mechanical engineering from Union College and an MBA from University of Rhode Island. Wendy is an ASQ Certified Reliability Engineer.

**Stefano Lena** Stefano Lena has spent two decades in researching, developing and delivering heating and cooling systems. With an MS in electronic engineering, he has experience across the cooling portfolio. Starting in product development, key account management, and into marketing, he built his data center competency. In 2020, he joined the Schneider Electric Cooling Innovation and Strategy team, concentrating on the liquid cooling program. Today, he is a subject matter expert at the forefront of this emerging technology, with additional research focused on systems-level flow management.

**David McGlocklin** has spent 25 years working in the IT and data center cooling business. In his work as a sustaining engineer, he most recently led the new product design and lab teams before joining the Cooling Innovation & Strategy team. He received his BS in mechanical engineering from Georgia Tech. Prior to joining Schneider Electric, he developed transport refrigeration for Carrier Transicold. He is very active in ASHRAE and AHRI, serving on multiple committees and is the current chair of SSPC127.

## RATE THIS PAPER ★★★★★

This document is to be considered as an opinion paper presenting general and non-binding information on a particular subject. The analysis, hypothesis and conclusions presented therein are provided as is with all faults and without any representation or warranty of any kind or nature either express, implied or otherwise.

## Resources

Optimizing AI Infrastructure: The Critical Role of Liquid Cooling
**Executive Report 3**

Guide to Environmental Sustainability Metrics for Data Centers
**White Paper 67**

How 6 AI Attributes Change Data Center Design
**White Paper 110**

Economizer Modes of Data Center Cooling Systems
**White Paper 132**

Navigating Liquid Cooling Architectures for Data Centers with AI Workloads
**White Paper 133**

10 Data Center Sustainability Improvements Quantified by Environmental Impact
**White Paper 188**

Direct Liquid Cooling System Challenges in Data Centers
**White Paper 210**

How Higher Chilled Water Temperature Can Improve Data Center Cooling System Efficiency
**White Paper 227**

Liquid Cooling Technologies for Data Centers and Edge Applications
**White Paper 265**

Browse all white papers and executive reports
**Data Center Research & Strategy Papers**

Cooling Economizer Mode PUE Calculator
**TradeOff Tool 11**

Browse all TradeOff Tools™
**Data Center Research & Strategy Tools**

**Note**: Internet links can become obsolete over time. The referenced links were available at the time this paper was written but may no longer be available now.

## Contact us

For feedback and comments about the content of this white paper:

Schneider Electric Data Center Research & Strategy
dcsc@se.com

If you are a customer and have questions specific to your data center project:

Contact your Schneider Electric representative at
www.apc.com/support/contact/index.cfm

© 2025 Schneider Electric. All rights reserved.



Life Is On Schneider Electric

# Appendix

We conducted a simulation with the objective of quantifying the impact of key design factors (cooling type, rack density, IT coolant supply temperature) on annual energy and water consumption. Below are the details of the methodology, assumptions, and results.

## Methodology

We examined the energy and water consumption of several data center scenarios using a proprietary simulation software. Here are the steps involved:

1. **Identifying scenarios.** We identified ten unique scenarios with varying cooling types (e.g., with/without adiabatic and liquid vs. air cooling) and densities. See **Table A1**.
2. **Selecting cooling components**. We selected air-cooled chiller models based on capacities and other operating conditions, as shown in **Table A1**. Chiller quantities vary by temperature, since temperature drives capacity per unit. A Motivair by Schneider Electric MCDU50 CDU is used in all liquid cooled scenarios. Quantities vary by density. For all air-cooled loads, we used fan wall model FWCV40L2A in varying quantities based on the capacity needed.
3. **Running simulations**. For each scenario, we varied TCS coolant supply temperatures from 25°C to 45°C using proprietary simulation software.
4. **Analyzing results.** Once simulations were completed, we plotted results to identify trends and computed absolute and percentage savings in energy and water use. We present results in the **sidebar** (introduction section) as well as **Figures 6** and **7**, and discussion throughout this paper.

**Table A1**

*Scenarios and chiller selection*

| Cooling type | | Density | Heat capture | Chiller model |
|---|---|---|---|---|
| Non-adiabatic | Air-cooled IT | 20 kW | 100% air cooling | XRAF6422A |
| | | 40 kW | | |
| | Liquid-cooled IT | 20 kW | 70% liquid cooling | XRAF4242A-EHT |
| | | | 30% air cooling | XRAF6422A |
| | | 40 kW | 75% liquid cooling | XRAF4242A-EHT |
| | | | 25% air cooling | XRAF4812A |
| | | 120 kW | 85% liquid cooling | XRAF4242A-EHT |
| | | | 15% air cooling | XRAF3212A |
| Adiabatic | Air-cooled IT | 20 kW | 100% air cooling | XRAF6422A |
| | | 40 kW | | |
| | Liquid-cooled IT | 20 kW | 70% liquid cooling | XRAF3242A-EHT |
| | | | 30% air cooling | XRAF6422A |
| | | 40 kW | 75% liquid cooling | XRAF3242A-EHT |
| | | | 25% air cooling | XRAF4812A |
| | | 120 kW | 85% liquid cooling | XRAF3242A-EHT |
| | | | 15% air cooling | XRAF3212A |

## Assumptions

In addition to the attributes noted in Table A1, below are key assumptions that impacted the design of each scenario:

- Data center IT capacity analyzed at 10 MW, with results normalized to energy and water use per kW of capacity
- Data center location is Paris, France, with max outdoor temperature of 44.1°C (20-year max + 4°C)
- All air-cooled chillers have economizer modes
- TCS coolant flow rate is 1.5 lpm per kW of liquid-cooled load
- Piping variations are not considered. The pumps run at their nominal 100% constant speed

## Air-cooled IT vs. liquid-cooled IT comparison

The gains in energy savings by going from air cooling to liquid cooling significantly outweigh the energy gains of adding adiabatic cooling to an air-cooled architecture. Table A2 summarizes the magnitude in savings for the rack densities simulated.

**Table A2**

*Comparison of savings by changing architectures*

*This demonstrates the biggest savings comes from the shift from air to liquid cooling*

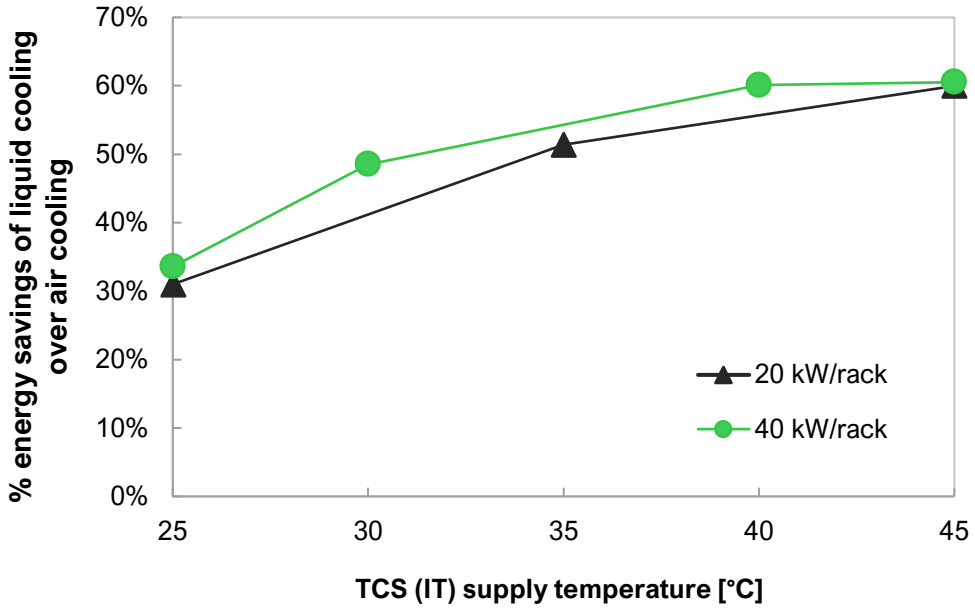
| Cooling architecture change | 20 kW / rack | 40 kW / rack |
|---|---|---|
| From non-adiabatic air cooling to adiabatic air cooling | 6.2% | 1.7% |
| From non-adiabatic air cooling to non-adiabatic liquid cooling (TCS supply from 25°C – 45°C) | 31.0% to 60.0% | 33.6% to 60.5% |
| From non-adiabatic liquid cooling to adiabatic liquid cooling (TCS supply from 25°C – 45°C) | 7.2% to 24.1% | 11.4% to 23.0% |

In addition, Figure A1 illustrates how the percent of energy saved increases as the TCS coolant supply temperature increases.

**Figure A1**

*% saving in energy consumption for liquid cooling compared to 100% air-cooled IT, using non-adiabatic outdoor heat rejection system.*

*Source: Schneider Electric Innovation and Strategy Cooling Team*



## Impact of location

As mentioned in the discussion for Design factor 2: Fluid IT inlet temperatures, the energy decreases at higher TCS supply temperature. The exact shape of the curve, the magnitude of energy consumption, and where diminishing returns occur is dependent on the climate and free cooling utilization. Figure A2 compares a data center in Paris with one located in Dubai to illustrate this point. Here, the diminished returns are the result of the chiller's upper limit of 40 °C.

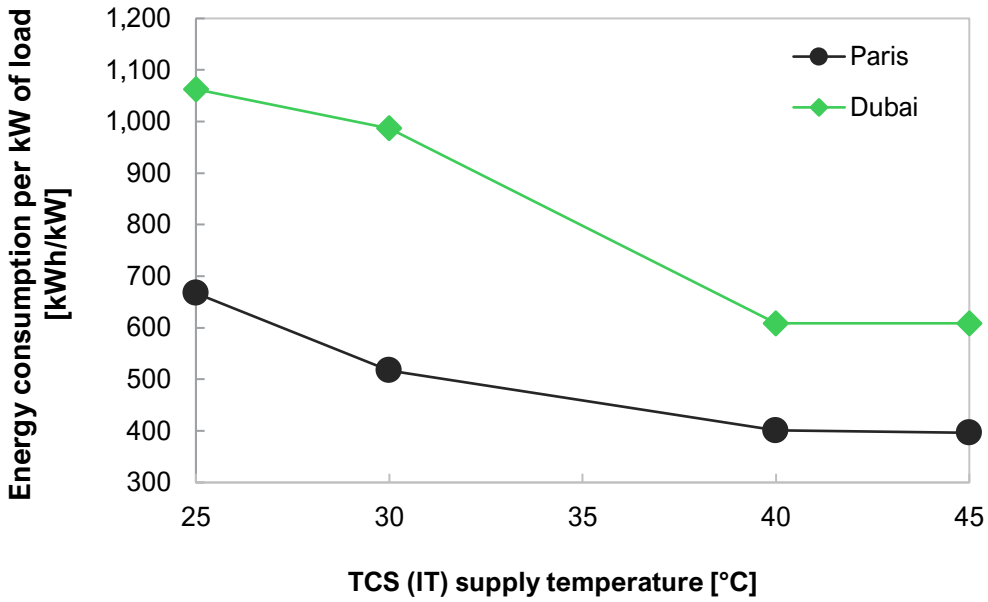



**Figure A2**

*Energy of liquid cooling in Paris vs. Dubai, to illustrate how climate impacts the energy curve*

*Source: Schneider Electric Innovation and Strategy Cooling Team*