Cool at the source. Scale with confidence.

Executive Summary

More power per rack. More performance. New thermal rules.

Artificial intelligence (AI), high-performance computing (HPC), advanced analytics, and other demanding workloads are pushing rack-level power consumption to levels never seen before.

This unprecedented power draw can also lead to unprecedented heat generation, favoring at-source heat capture over high-volume airflow.

Liquid cooling delivers effective cooling at scale, trimming the cost of moving air.

While it used to be limited to CPUs and GPUs, it is now extending to SSDs for a reason: keeping drives at stable operating temperatures helps prevent throttling, improves performance consistency, and reduces the facility-level energy we use moving air instead of moving data. In short, liquid-cooled SSDs give operators a path to higher performance and lower energy use, both of which are now table stakes for modern AI infrastructure.

To quantify those efficiency gains beyond first principles and platform trends, it helps to look at measured outcomes from real deployments and published analyses. With that in mind, let’s explore an independent study to learn about the energy savings liquid cooling can provide.

Reducing operating costs through data center liquid cooling: A case study

Vertiv published a case study focused on the results of increasing the percentage of data center liquid cooling and the resulting reduction in total data center power.[1]

The results help demonstrate what innovators and early adopters already suspect: Liquid cooling is an effective technology for reducing power consumption as traditional cooling falters under modern, high-performance workloads.

Figure 1 represents the results of this study. The four tested cases (labeled “Study X,” “Study Y,”…) are on the x-axis. The blue line represents the increasing portion of the total data center cooling load served by liquid cooling, ranging from 0% in Study 1 to 74.90% in Study 4 (the remaining load was served by traditional air cooling).

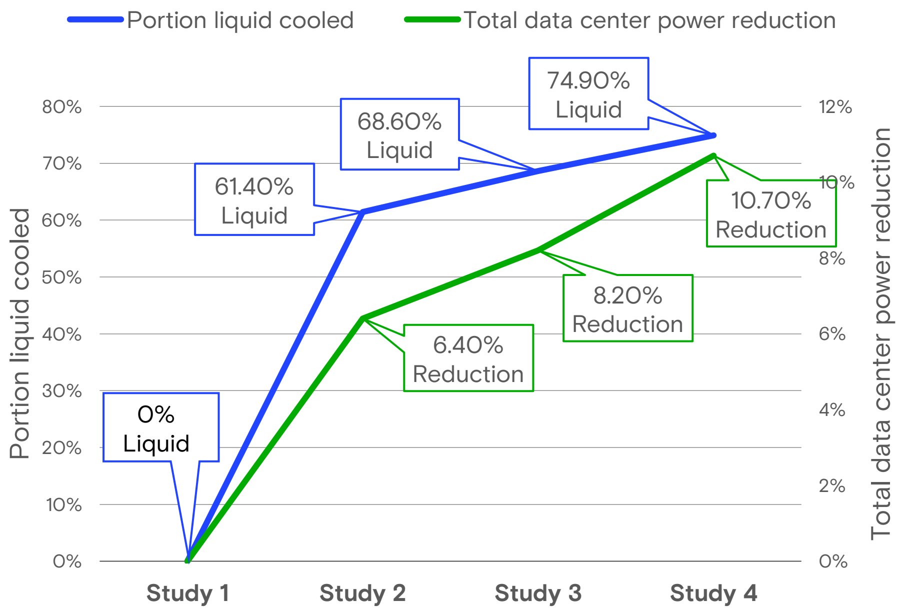

Figure 1: Effects of data center liquid cooling

The green line represents the corresponding reduction in total data center power consumed. The study found that increasing the percentage of the data center cooling load satisfied by liquid cooling showed a direct impact on overall power savings across the entire data center.

These findings underscore the role liquid cooling can play in enabling overall power reduction across data centers, making it an essential strategy for reducing data center power needs.

Why SSDs need liquid cooling now

Generative AI and GPU-dense architectures have redefined thermal envelopes. Top-end data center GPUs have moved from ~600W-class[2] to up to 1,200W.[3] Looking forward, these numbers could continue to climb as far as 600kW[4] per rack, a region where air cooling becomes both thermally ineffective and economically prohibitive at required airflow volumes.

AI servers can support multiple, high-performance, high-density SSDs inside a single server chassis. With these SSDs often drawing up to 25 watts each, the power (and, thus, heat) density per server can be substantial. Liquid cooling helps SSDs avoid thermal throttling during intensive write activity while reducing reliance on excessive airflow.

SSD cooling needs a solution that scales with power density without forcing disproportionate increases in chassis airflow. The following section explains how cold-plate liquid cooling achieves this by moving heat through a defined conduction path (SSD enclosure → thermal interface material → cold plate) and into a controlled single-phase coolant loop.

In the Uptime Institute Cooling Systems Survey 2024: Direct liquid cooling, their survey asked respondents

“At what IT rack power density do you think air cooling is so costly (or unable to meet cooling requirements) that the use of direct liquid cooling becomes necessary?’

29% of respondents reported that liquid cooling becomes necessary at just 20-29kW (only 20% cited a value less than 20kW)[5]

1. See “Quantifying the Impact on PUE and Energy Consumption When Introducing Liquid Cooling Into an Air-cooled Data Center” on veritiv.com for more details. (Note: PUE = power usage effectiveness).
2. See “GPU power and maximum number of trays in the enclosure” (lenovo.com).
3. See “AI’s Engine Room: Inside the High-Performance Data Centers Powering the Future” (equinix.com).
4. See “Immersion cooling systems: Advantages and deployment strategies for AI and HPC data centers” (vertiv.com).
5. See “Uptime Institute Cooling Systems Survey 2024: Direct liquid cooling” (intelligence.uptimeinstitute.com).

Cold-plate liquid cooling for SSDs: How it works

To extend liquid cooling to SSDs, cold plates (machined metal blocks with internal microchannels mounted to the SSD enclosure) are a practical option. A coolant (often water-glycol in a single-phase loop) circulates through the plate, extracting heat at the device boundary. Figure 2 represents an example. The thermal interface material (TIM) shown in light blue is sandwiched between the SSD enclosure (grey) and the cold plate (black). The TIM facilitates better thermal transfer between these two elements. Coolant (shown in blue) circulates through channels in the cold plate, extracting heat from the SSD. The warmed coolant is pumped through the liquid cooling system, where its heat is extracted; the cycle then repeats.

Before spring-loaded cold plates and blind-mate quick-disconnect manifolds, liquid-cooled devices (including SSDs, when attempted) relied on rigid, mechanically mounted cold plates with fixed coolant-tubing connections. These traditional designs required draining and disconnecting coolant lines to service or replace a device, making hot-swap impossible.

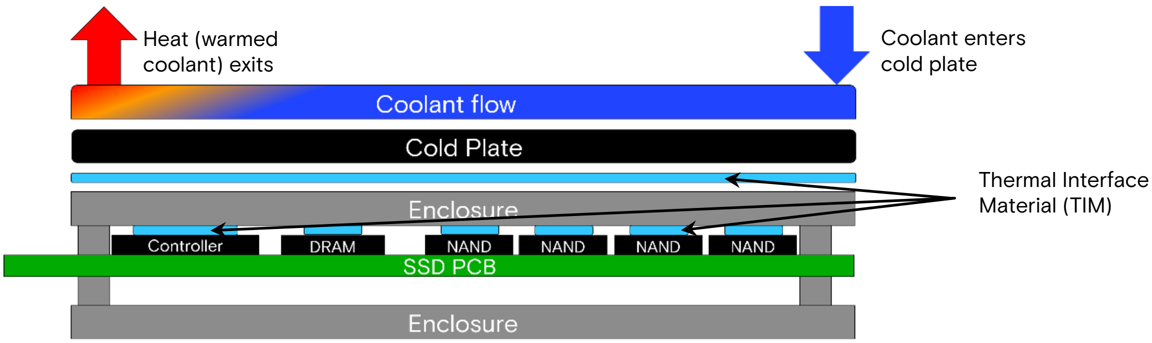

Figure 2: Simplified SSD cold plate implementation (example)

Thermal efficiency analysis: Liquid vs. air for SSDs

Air cooling depends on convective heat transfer, which is proportional to airflow velocity and the temperature differential; pushing enough air across densely packed SSDs and other components can be challenging. Since cold-plate systems place a high-conductivity metal and a fast-moving liquid directly at the heat source, they can significantly increase effective heat transfer compared with air.

To illustrate the potential magnitude of this difference, consider an example server with 32 NVMe™ SSDs, each drawing its maximum consumption of 25 watts (0.800 kW total). We'll estimate air and liquid required flow differences to cool this group of SSDs using two examples: Case A, where the data center ambient temperature is 11.1 °C different from the SSDs, so ΔT = 11.1 °C (i.e., the SSDs are 11.1 °C warmer than the air surrounding them). ΔT is a temperature difference, not an absolute temperature. 11.1 °C and 11.1 K, both expressed as a difference in temperature, are numerically the same, while 25 °C = 298 K is an absolute temperature and is not used here. We'll use the data to estimate the electricity (in watts) needed to reach those values.

The necessary fluid flow (air, glycol-water, or other coolant) to remove a given amount of heat is calculated using:

Heat removed (rate): $\dot{Q} = m \times c_p \times \Delta T$[6] and **Mass flow:** $m = \rho \times \dot{V}$[7]

where:

Formula element		Units	Formula element		Units
$\dot{Q}$	Heat rate	Watts (=Joules/seconds)	$\dot{V}$	Flow rate	m^3/s (volumetric flow)
$\dot{m}$	Mass flow	kg/s (volumetric)	ρ	Fluid density	kg/m^3
c_p	Heat capacity	J/(kg·K)	ΔT	Temperature difference	Same units as temperature
C	Celsius scale	Celsius	P	Pressure difference	Pascals
K	Kelvin scale	Kelvin	J	Energy unit	Joules

Table 1: Heat removal rate and mass flow formula units

6. Air properties at ~25 °C, 1 atm: ρ = 1.184 kg/m³ and cp = 1.007 kJ/kg·K (me.psu.edu).
7. M. Bahrami, Simon Fraser University, The First Law of Thermodynamics: Control Volumes.

Air flow requirements for Case A | ΔT 11.1 °C

In this example, we will use the following assumptions and present calculations for clarity (calculations will not be shown in subsequent airflow analysis).

- Ambient air temperature T = 25 °C and an allowed air temperature rise across the SSDs ΔT = 11.1 °C
- Air properties at ~25 °C, 1 atm: $\rho = 1.184\ kg/m^3$ and cp = 1.007 kJ/kg·K (me.psu.edu)

Air flow requirements calculations

1. $\dot{Q}$ = 0.800 kW = 0.800 kJ/s (a unit conversion for simplicity)
2. ΔT = 11.1 °C = 11.1 °K (unit conversion)
3. Since $\dot{Q} = \dot{m} \times cp \times \Delta T$, $\dot{m} = \dot{Q} \div (cp \times \Delta T)$
   $= 0.800\ kJ/s \div (1.007\ kJ/kg{\cdot}K \times 11.1\ K)$
   $= 0.800 \div 11.1777\ kg/s$
   $\approx 0.0716\ kg/s$
4. Since $\dot{m} = \rho \times \dot{V}$, $\dot{V} = \dot{m} \div \rho$ (identity), $\dot{m} \approx 0.0716\ kg/s$ (from above), and $\rho = 1.184\ kg/m^3$ then,
   $\dot{\mathbf{V}} = 0.0716\ kg/s \div 1.184\ kg/m^3$ and $\dot{V} \approx$ **0.06045** $\mathbf{m^3/s}$ to maintain ΔT = 11.1°C

Air flow requirements for Case B | ΔT 8.3°C

Using the same process as for air cooling for Case A, with a desired ΔT of 8.3°C, we find $\dot{\mathbf{V}} = 0.0957\ kg/s \div 1.184\ kg/m^3$ = **0.08084** $\mathbf{m^3/s}$, the air flow needed to maintain ΔT = 8.3°C.

Summarizing:

Case A | ΔT 11.1°C: $\dot{\mathbf{V}} \approx$ **0.06045** $\mathbf{m^3/s}$

Case B | ΔT 8.3°C: $\dot{\mathbf{V}} \approx$ **0.08084** $\mathbf{m^3/s}$

The liquid cooling flow requirements are calculated similarly.

Liquid cooling for Case A | ΔT 11.1 °C

We will use an example water-glycol (liquid) mix: "DOWFROST™ LC 25," a propylene glycol-based heat transfer fluid.[8] This fluid has ρ= 1031.5 kg/m^3 and cp (specific heat) = 3.88 kJ/kg·K. Since we are cooling the same set of SSDs, we also know:

1. $\dot{Q}$ = 0.800 kW = 0.800 kJ/s (same equation as used for air cooling)
2. ΔT = 11.1 °C = 11.1 °K
3. $\dot{m} = \dot{Q} \div (cp \times \Delta T) = 0.800\ kJ/s \div (3.88\ kJ/kg{\cdot}K \times 11.1\ °K)$
   $= 0.800 \div 43.068\ kg/s$
   $= 0.01858\ kg/s$
4. $\dot{\mathbf{V}} = \dot{m} \div \rho = 0.01858\ kg/s \div 1031.5\ kg/m^3$
   $= \mathbf{1.8008{\times}10^{-5}\ m^3/s}$

Liquid cooling Case B | ΔT 8.3°C

Using the same process as liquid cooling, Case A and using ΔT 8.3°C:

$\dot{\mathbf{V}} = 0.02484\ kg/s \div 1031.5\ kg/m^3$ so $\dot{\mathbf{V}} = \mathbf{2.40814 \times 10^{-5}\ m^3/s}$

Table 2 combines these calculated flow values into a single flow reference:

Cooling type	Case	Required flow across the SSDs
Air	A (ΔT 11.1°C)	$\dot{V} = 0.06045\ m^3/s$
Liquid	A (ΔT 11.1°C)	$\dot{V} = 1.8008{\times}10^{-5}\ m^3/s$
Air	B (ΔT 8.3°C)	$\dot{V} = 0.08084\ m^3/s$
Liquid	B (ΔT 8.3°C)	$\dot{V} = 2.40814{\times}10^{-5}\ m^3/s$

Table 2: Air and liquid cooling flows at evaluated ΔT (air cooling assumes 25 °C inlet temperature)

8. See "Engineering and Operating Guide for DOWFROST™ LC fluids. In that guide, Table 2 ("Physical properties of DOWFROST™ LC 25 Heat Transfer Fluid" (lentusllc.com).

We can use these flow values to estimate the electricity required by each cooling method for each case.

Air cooling: Estimating watts needed to reach the required air flow
We estimate the power required to achieve the necessary airflow using the air power equation $P\_air = \Delta P \times \dot{V}$.[9] We also know that air flowing through the server encounters resistance from the server bezel and other internal components, resulting in a pressure difference (ΔP, measured in Pascals; this is not the case with liquid cooling). Since servers differ, we'll use an estimate of ΔP = 250 Pa (mdpi.com).

- Air, Case A
  - $\dot{V} = 0.06045\ m^3/s$
  - $P\_air = \Delta P \times \dot{V}$
    $= 250\ Pa \times 0.0604483\ m^3/s$
    $= 15.1121$ W for 32 SSDs
- Air, Case B
  - $\dot{V} = 0.08084\ m^3/s$
  - $P\_air = \Delta P \times \dot{V}$
    $= 0.0808333\ m^3/s$
    $= 20.2083$ W for 32 SSDs

Since we are interested in actual power, we'll use the above values combined with a unitless efficiency conversion ("η") to convert electrical input into airflow (the American Society of Heating, Refrigerating and Air-Conditioning Engineers, ASHRAE calls this "wire-to-air efficiency" and notes that fan, motor, and controller all have losses) using $P\_elec = P\_ air \div \eta$).[10]

- **η = 0.25** (very conservative, aligns closely with an axial-fan under load)[11]
- **η = 0.30** (reasonable "typical" example for restricted small-fan airflow)[12]
- **η = 0.40** (peak efficiency for axial propeller)[13]

Watts (P_elec) demand for air cooling cases A and B for these values of η are summarized in Table 3

Case	η	P_elec
A (ΔT 11.1°C)	0.25	60.45 watts
	0.30	50.37 watts
	0.40	37.78 watts
B (ΔT 8.3°C)	0.25	80.83 watts
	0.30	67.36 watts
	0.40	50.52 watts

Table 3: Estimated electrical power for air cooling, with example efficiency factors

Liquid cooling: Estimating watts needed to reach the required fluid flow

We estimate the pump electrical power required to achieve the necessary coolant flow using the required coolant flow (from Table 2), the power delivered to the coolant, the pump input shaft power efficiency, and the coolant loop pressure drop (ΔP). We'll use an example ΔP = 14 kPa as suggested by Chilldyne.[14]

Hydraulic power (P_hyd) is the rate of doing flow work, and pressure is energy per unit volume, so multiplying the pressure difference (Pa) by the volume flow (m^3/s) yields watts. Hydraulic power delivered to the coolant fluid can be calculated as $P\_hyd\ (W) = \Delta P\ (Pa) \times \dot{V}\ (m^3/s)$.[14] Using values of $\dot{V}$ from Table 2, we see:

- **Case A:** $P\_hyd = 14\ kPa \times 1.8008\times10^{-5}\ m^3/s \approx 0.252$ W
- **Case B**: $P\_hyd = 14\ kPa \times 2.40814\times10^{-5}\ m^3/s \approx 0.338$ W

9. See "Fan Efficiency, an Increasingly Important Selection Criteria (engineeringtoolbox.com).
10. See "How to improve energy efficiency of fans for air handling units" (rehva.eu).
11. See "Fan Efficiency, an Increasingly Important Selection Criteria" (mouser.com).
12. See "Fan Efficiency, an Increasingly Important Selection Criteria" (mouser.com).
13. See "Fan Efficiency, an Increasingly Important Selection Criteria" (mouser.com).
14. See "Improving Liquid Cooled Server Efficiency with Larger Connectors and Design Improvements" (chilldyne.com). This value for η and all values in this document are intended as illustrative examples only and may not reflect actual products or implementations.

Finally, we assume that pumps are not 100% efficient and consider three example pump efficiency values (η) for illustrative purposes. In this analysis, pump efficiency η is defined as hydraulic power (P_hyd) imparted to the coolant divided by electrical input power to the circulation pump (P_elec).[15] Hence, electrical power (P_elec) is calculated as P_elec = P_hyd ÷ η.

- **η = 0.25** (a conservative pump efficiency assumption)[16]
- **η = 0.40** (a mid-case for data center coolant circulation)[17]
- **η = 0.60** (a planning value supported by guidance for typical centrifugal pump efficiencies around 55% for small pumps and 70% for large pumps)[18]

The liquid-cooling power required for cases A and B for these η values is summarized in Table 4.

Case	η	P_elec
A (ΔT 11.1°C)	0.25	1.01 watts
	0.40	0.63 watts
	0.60	0.42 watts
B (ΔT 8.3°C)	0.25	1.35 watts
	0.40	0.85 watts
	0.60	0.56 watts

Table 4: Estimated electrical power for requirements for liquid cooling, with example efficiency factors

At-source heat removal cuts power and scales where air cooling cannot

After working through real thermodynamic equations and realistic efficiency assumptions, the payoff is clear. Table 5 below summarizes the calculated electrical power needed to cool the same set of SSDs using traditional air cooling and cold-plate liquid cooling.

Case	Air cooling (watts)	Liquid cooling (watts	Liquid cooling power reduction[19]
A (ΔT 11.1°C)	37.78 to 60.45	0.42 to 1.01	98.89% to 98.33% reduction
B (ΔT 8.3°C)	50.52 to 80.83	0.56 to 1.35	98.89% to 98.33% reduction

Table 5: Estimated electrical power for air cooling, at example efficiency factors

What emerges is not a marginal improvement, but a fundamental efficiency shift—one that explains why liquid cooling changes facility-level power, fan energy, and power usage effectiveness (PUE) rather than merely relocating heat. For readers evaluating dense AI and storage platforms, these numbers provide a concrete answer to an important question: what do you gain, in energy terms, by moving heat with a liquid rather than air?

Implementation considerations for SSD cold-plate cooling

Integrating liquid cooling at the SSD level requires attention to several practical design and operational details. The following outlines the key mechanical, thermal, and facility factors to consider for a successful cold-plate deployment.

SSD form factors and mechanics: Due to their inherent thermal density, SSD-dense enclosures (such as those supporting E1.S form factors) are best. Cold plates should support compression, durable TIMs, and dripless quick disconnects.

15. See "Pump Efficiency Explained" (engineeringexcel.com).
16. See "Improving Liquid Cooled Server Efficiency with Larger Connectors and Design Improvements" (chilldyne.com).
17. See "What is the efficiency of a centrifugal pump?" (wilo.com).
18. See "What is the efficiency of a centrifugal pump?" (wilo.com).
19. Ranges use the lowest and highest values for each case from Table 3 (air cooling) and Table 4 (liquid cooling), and are calculated as %reduction = ((1 – (liquid cooling value / air cooling value)) for each case. For SSDs in this example. Other designs or components may yield different results

Facility impacts: Because direct liquid cooling (DLC) extracts most of the heat into liquid, operators can raise supply air temperatures, reduce fan speeds, and target a lower power usage effectiveness (PUE).[20] Hyperscalers lead adoption, with co-locations adapting for AI tenants.

Criteria	Air Cooling	Liquid (Cold-Plate)
Heat removal path	Ambient airflow across bays; indirect	At-source via cold plate; direct
Steady-state SSD temperatures	Higher, with larger deltas at bursts	Lower, tighter deltas under load
Fan energy	Higher: More noise, cost, and wear	N/A
Serviceability	Simple hot-swap	Hot-swap via quick disconnects
Facility cooling load[21]	Highly dependent on room cooling effectiveness	Reduced room load

Table 6: Comparing air and cold-plate liquid cooling

Why this transition is inevitable for AI

As AI pushes the performance limits of SSDs, power consumption will increase, driving the need for more advanced, scalable cooling solutions. Air cooling may not scale physically or economically, making liquid cooling a density / performance enabler that governs useful compute output. With cold-plate cooling on SSDs, hot-swap manifolds, and denser server designs, storage becomes an active participant in liquid-cooled architectures. The business case blends operational expenditure savings, environmental / social / governance (ESG) goals, and density gains, while the engineering case delivers thermal headroom and consistent results under sustained I/O.

20. PUE compares total facility power (everything the building uses: IT + cooling + power distribution losses + lighting, etc.) to IT equipment power (servers, storage, networking). It indicates how many watts you must buy at the utility meter to deliver 1 watt to IT. See "What Is PUE (Power Usage Effectiveness) and What Does It Measure?" (vertiv.com).
21. See "Characteristics and Risks of Emerging Large Loads" (nerc.com) for additional background on facilities load and how it refers to the necessary infrastructure and power required to remove the heat from the facility once it has been removed from the SSDs in these examples.

micron.com/ssd

©2026 Micron Technology, Inc. All rights reserved. All information herein is provided on as "AS IS" basis without warranties of any kind, including any implied warranties, warranties of merchantability or warranties of fitness for a particular purpose. Micron, the Micron logo, and all other Micron trademarks are the property of Micron Technology, Inc. All other trademarks are the property of their respective owners. Products are warranted only to meet Micron's production data sheet specifications. Products, programs and specifications are subject to change without notice. Dates are estimates only. Rev. A 03/2026 CCM004-1681249710-11862