



Why liquid cooling matters now

How rising density is reshaping the modern data center



The heat is on

If your racks are running hotter, your PUE is creeping higher, or your GPU workloads are slowing down, you're not alone.

Across the industry, AI and high-performance computing are rewriting the rules of thermal management. What used to be manageable with conventional air systems is now pushing data centers to their operational limits.

Why traditional cooling can't keep up

As GPU-driven workloads intensify, racks that once drew 10–20 kW are now **reaching 50, 80, even 150+ kW**. At those power levels, air cooling alone can no longer efficiently remove heat without overprovisioning space or power—two things most facilities can't afford to waste.

Meanwhile, sustainability pressures are rising. **Cooling systems can account for up to 40% of total data center energy consumption**, and organizations are being asked to meet ESG and efficiency goals without sacrificing uptime or performance.

This is no longer just a facilities issue—it's a business challenge. The next generation of infrastructure will belong to the organizations that master **how to keep performance high while energy and space footprints stay low**.

"In the recent survey results, the industry average PUE of 1.56 reveals a continuing trend of inertia, while many **recent builds consistently achieve a PUE of 1.3—and sometimes much better.**"

— *Uptime Institute,
Global Data Center Survey*

The moment it starts to matter

For many IT and data center leaders, the realization doesn't come all at once—it starts as a series of pain points.

Perhaps a new AI training workload starts triggering thermal alerts, or your GPU servers throttle back performance to stay within temperature limits. Maybe your colocation space is maxed out, but adding another cabinet means costly expansion. Or your sustainability report calls out cooling inefficiencies as a top driver of excess energy use.

These are the signs that traditional air systems are approaching their limits.

The workloads that power innovation—from AI inference to edge analytics—demand more from your infrastructure than air cooling alone can deliver.

That's where liquid cooling enters the picture. Once considered niche, it's now becoming the practical solution for anyone balancing performance, density, and sustainability.

“ The world's most advanced workloads demand the world's most efficient cooling.



How liquid cooling works (and why it's different)

Liquid cooling isn't new, but its role in today's data centers has never been more critical.

At its core, it's a simple concept: rather than moving heat through air, it transfers heat through liquid—a medium that's **20–25 times more efficient** at heat exchange.

According to the [International Energy Agency](#), cooling systems can consume anywhere from **7% to 30% of a data center's total electricity**, depending on efficiency and design.

That efficiency opens the door to new cooling designs:

- **Rear-Door Heat Exchangers (RDHX):** Add-on liquid-cooled doors that remove heat as it leaves the rack—ideal for hybrid or retrofit environments.
- **Direct-to-Chip cooling:** Coolant flows through small cold plates attached to CPUs and GPUs, capturing heat right at the source.

- **Immersion cooling:** Entire servers are submerged in a dielectric fluid that absorbs and dissipates heat efficiently, allowing for ultra-high-density workloads.

These approaches share a common goal: to maximize compute power while minimizing energy use, footprint, and risk.

In practice, liquid cooling allows operators to support workloads of 150+ kW per rack with stable performance and better efficiency.

**Cooling is no longer overhead.
It's a strategic decision.**

The business case: What it solves

Every organization's challenges look different—but most share common pain points.

Here's how liquid cooling directly addresses them:

Your challenge	How liquid cooling helps
Power density is rising	Supports high-density racks efficiently (up to 150+ kW)
Thermal throttling limits workloads	Maintains stable chip temperatures for peak performance
Physical space is constrained	Increases compute density, deferring costly expansions
Energy bills keep climbing	Cuts cooling energy use by 25–40%
Sustainability goals are tightening	Lowers PUE, reduces water usage, and enables heat reuse

It's not just about keeping equipment cool—it's about keeping innovation possible. Liquid cooling creates space for growth, both operationally and sustainably.

"Full implementation of liquid cooling produced an **18.1% reduction in facility power** and a 10.2% reduction in total data center power compared to 100% air cooling."

— Vertiv + NVIDIA study


Debunking common myths

As more enterprises and colocation providers adopt liquid cooling, a few lingering misconceptions remain.

Let's clear them up:

Myth	Reality
"It's too complex."	Modern systems are modular and pre-engineered for quick deployment.
"It's too risky."	Sealed systems and dielectric fluids eliminate leak concerns.
"It's only for hyperscalers."	Scalable options now make it practical for enterprises and colocation facilities.
"It's too expensive."	Lower operating costs and delayed expansion quickly offset initial investment.

Liquid cooling has evolved from experimental to essential—and it's easier to adopt than ever before.



"U.S. data center electricity usage climbed from 58 TWh in 2014 to 176 TWh in 2023 and is estimated at 325–580 TWh by 2028."

— U.S. Department of Energy

When and how to get started

Transitioning to liquid cooling doesn't require a complete data center overhaul.

Many organizations start small—focusing on their densest or most demanding workloads—and expand over time.

Begin by understanding where the pressure points are.

Audit your rack densities, especially those exceeding 30–40 kW. Evaluate your PUE and cooling efficiency to see where targeted improvements will have the most impact. Many operators begin with hybrid deployments, integrating rear-door or direct-to-chip systems alongside existing air infrastructure.

From there, the transition becomes modular and scalable.

Each step improves efficiency, extends capacity, and helps future-proof your data center against the rising thermal demands of AI, HPC, and next-generation workloads.

Liquid cooling isn't about replacing what works—it's about optimizing what's next.

“Built for sustainability: Commissioned Hillsboro 4 Phase 1 and reached substantial completion on Atlanta 5 Phase 1, each with a design PUE ≤ 1.4 and design WUE of zero.”

— *Flexential ESG Report*

The future flows in liquid

AI, automation, and real-time analytics are driving the need for more compute power—and more cooling innovation—than ever before. Liquid cooling is emerging as the key to balancing performance, density, and sustainability in this new era of infrastructure.

Those who embrace it early will not only reduce energy and space constraints—they'll gain an advantage in speed, reliability, and long-term cost efficiency.

"As workloads evolve, the smartest infrastructure leaders are building cooling strategies that evolve with them."

Liquid cooling is no longer futuristic—it's here, scalable, and shaping the next wave of high-performance data centers.

Ready to learn what it takes to become liquid-ready?

[Explore More](#)

