



Scaling for the Future: Liquid Cooling's Role in AI Data Centers

By Mike Connaughton and Jacques Fluet

The TIA-942 standard has been a resource for global data center design and implementation for many years. The “c” revision (ANSI/TIA-942-C) was released in May 2024, and it covers a variety of aspects such as physical infrastructure, cabling, power, cooling, redundancy, and security.

Data centers must support increasingly higher rack power densities as compute-intensive artificial intelligence (AI) and machine learning (ML) workloads scale. At the same time, operators are aiming to reduce energy consumption, extend equipment life, lower costs, and comply with stringent regulatory requirements.

To meet these demands, many AI data center operators are considering liquid cooling solutions, which use water or dielectric fluids to transfer heat up to 20 times more efficiently than air. Despite these advantages, liquid cooling introduces a range of new design, operational, and infrastructure challenges.

This article explores why data centers deploy direct liquid cooling (DLC) and in-rack indirect liquid cooling (IILC) systems, outlines key opportunities and barriers, and addresses infrastructure considerations for cabling and coolant routing. It also covers space, structure, and scalability requirements, highlights essential monitoring parameters, and details safety and operational needs—including redundancy and resilience.

To meet these demands, many AI data center operators are considering liquid cooling solutions, which use water or dielectric fluids to transfer heat up to 20 times more efficiently than air.

ADVANTAGES OF LIQUID COOLING IN AI DATA CENTERS

Many liquid cooling deployments are driven by compute-intensive workloads that require higher performance and thermal design power (TDP) than conventional applications. Generative AI (GenAI), for example, relies on graphics processing units (GPUs) that deliver up to 12 times the processing performance of general purpose CPUs.

Beyond performance, data center operators increasingly prioritize energy efficiency, lower operational costs, and improved hardware reliability. Liquid cooling helps maintain optimal operating temperatures, reducing thermal-induced failures and extending equipment lifespan. It also supports heat reuse strategies by capturing and redirecting waste heat to nearby systems or buildings—advancing both sustainability and energy efficiency goals.

Liquid cooling further reduces facility-wide water consumption and decreases reliance on high speed fans. Immersion cooling, in particular, can optimize space by eliminating traditional hot and cold aisle layouts. Some configurations also support concurrent maintainability—allowing operators to service or replace rack-level components without taking the system offline. Although overall data center energy demands may increase over time, liquid cooling can improve total cost of ownership (TCO) by increasing thermal efficiency and reducing cooling-related energy consumption.

LIQUID COOLING ARCHITECTURES: DLC AND IILC EXPLAINED

Liquid cooling systems are generally categorized as either DLC or IILC. DLC uses a liquid coolant—typically water or dielectric fluid—to transfer heat directly from electronic components. It provides significantly higher efficiency than traditional air cooling, especially for infrastructure with high thermal design power (TDP) ratings.

There are two primary types of DLC configurations. In direct-to-chip (DTC) cooling, liquid is pumped

through cold plates attached to heat-generating components, with air cooling handling any remaining thermal load. In immersion cooling, entire servers are submerged in a dielectric liquid that absorbs and dissipates heat.

Both DTC and immersion systems can operate in single-phase or two-phase modes. Single-phase systems circulate liquid continuously through a heat exchanger without phase change, while two-phase systems use engineered fluids that evaporate during heat absorption and condense back into liquid form. Some hybrid systems combine DTC and immersion within a single deployment.

Unlike DLC, IILC removes heat from exhaust air at the rack or server level using liquid-based components such as rear door heat exchangers (RDHx). While IILC is less efficient than DLC, it significantly outperforms traditional room- or row-level air cooling and is easier to integrate into existing deployments. RDHx units can operate as standalone solutions or integrate with DTC systems for optimized thermal performance.

OPPORTUNITIES AND BARRIERS

Despite its benefits, liquid cooling deployments remain relatively limited compared to traditional air cooling, which offers proven reliability, lower initial costs, and broad industry familiarity. In contrast, liquid cooling typically requires higher upfront capital investment in specialized infrastructure and equipment, particularly in retrofit scenarios.

Adoption on a broader scale has been slow due to challenges, including a lack of standardized metrics, design best practices, and performance benchmarks for liquid cooling systems. While issues like skilled labor shortages and evolving supply chains affect the entire data center industry, they can complicate liquid cooling deployments in particular. Additionally, restrictions on refrigerants with high global warming potential (GWP) and dielectric fluids containing per- and polyfluoroalkyl substances (PFAS) chemicals introduce environmental and compliance complexities. Concerns about long term reliability and limited rack level redundancy further contributes to operators' hesitation to adopt this promising capability.

Nevertheless, liquid cooling is positioned for significant growth. According to Grand View Research, cloud and hyperscale operators are expected to lead the market, which is projected to grow at a compound annual growth rate (CAGR) of nearly 24 percent over the next decade.¹ Additionally, a recent survey suggests that more than one-third of enterprise data centers plan to implement some form of liquid cooling by 2026.² As adoption accelerates, operators will continue to evaluate technical, operational, and regulatory considerations to guide deployment strategies.

INFRASTRUCTURE CONSIDERATIONS FOR CABLING AND COOLANT ROUTING

Integrating liquid cooling into AI data center rows and racks requires careful cabling infrastructure planning. Space must be allocated for liquid distribution components, including piping, manifolds, and airflow paths. In DTC systems, racks must accommodate both liquid cooling manifolds and rear airflow, while cable trays should allow space for cooling system pipes. Racks may require added depth for RDHx to support thicker rear doors to route cables and manifolds—especially when combined with DTC systems.

Immersion cooling introduces additional operational requirements. All cabling must exit from the top of the rack, and any portion submerged in coolant must be chemically compatible with the fluid. Material compatibility is critical for both cabling and coolant. Cabling considerations include sheath and connector integrity, mechanical durability, electrical or optical performance, signal integrity, and compliance with flame rating and labeling standards. Fluid-related factors span potential contamination from cable materials, impact on thermal performance, and effects on system characteristics such as viscosity, filtering, and pump reliability.

SPACE, STRUCTURE, AND SCALABILITY

Liquid cooling systems require dedicated space for circulation and distribution, affecting both floor loading and ceiling-mounted (hanging) capacity in AI data centers. To accommodate future growth, system design should prioritize flexibility, maintainability, and modularity.

Design considerations vary across DTC, immersion, and RDHx implementations, with each introducing distinct structural demands. DTC adds weight from piping, liquid lines, pumps, and cooling distribution units (CDUs). Immersion systems impose the greatest load due to the volume of dielectric fluid and supporting infrastructure, while RDHx increases floor loading through liquid weight, piping, and heavier rear door assembly.

All three systems require coordinated pipe and cable routing to maintain accessibility. Quick-coupling connectors and shutoff valves are essential to support maintenance and expansion without disrupting live operations. Leak management is also critical: DTC and RDHx systems typically use drip trays, sealed couplers, and leak sensors, while immersion systems may incorporate double-walled tanks for fortified containment.

Pipe sizing should account for future capacity, and CDUs in DTC, immersion, and RDHx deployments must scale accordingly. Filtration requirements vary by system and should align with allowable particle size to ensure long-term thermal performance and reliability.

KEY LIQUID COOLING MONITORING PARAMETERS

Real-time monitoring helps maintain the performance, reliability, and safety of liquid cooling in AI data centers. Key operational parameters include:

- **Temperature and Volume:** Monitor liquid temperature at system ingress and egress points—at the rack manifold for DTC, at the rack for immersion, and at the rear door for RDHx. Volume tracking varies by system. DTC systems monitor volume to support server additions or removals, immersion systems track fluid level in the tank, and RDHx systems measure total liquid volume.

- **Pressure and Quality:** Measure liquid pressure at the pump across all three systems. Fluid quality is critical, as impurities can clog DTC cold plates, cause failures in immersion systems, and reduce heat exchange efficiency in RDHx.
- **Leak Detection and Environmental Monitoring:** All liquid cooling systems require leak detection to prevent fluid loss and equipment damage. Air quality monitoring for two-phase vapor leaks applies to DTC and immersion systems but is not required for RDHx. Dew point monitoring—tracking both temperature and humidity—is important across all systems to prevent condensation.

SAFETY AND OPERATIONAL REQUIREMENTS

Liquid cooling deployments introduce distinct safety and operational requirements that vary by system type. Each approach requires updates to standard operating procedures, maintenance protocols, and emergency plans, along with specialized training for AI data center support personnel.

Operators must monitor liquid and pipe temperatures across all systems. Coolants should be non-toxic, recyclable, and nonflammable to minimize environmental and health risks. To reduce slip hazards from leaks or drips, system designs should incorporate containment strategies and appropriate floor treatments.

Weight is another critical consideration in DTC systems, which must account for the weight of individual servers. Immersion cooling, particularly with horizontal racks, requires provisions for safely lifting and removing heavy servers during maintenance. RDHx systems require structural support and handling procedures to manage the added weight of rear door assemblies.

Personal protective equipment (PPE)—such as gloves, safety glasses, and protective clothing—is essential across all systems, with added emphasis on immersion cooling, where coolant handling and server access are more frequent. Operational procedures must also address new server and rack configurations, as well as the coolant distribution infrastructure.

REDUNDANCY AND RESILIENCE

Redundancy is a foundational design requirement in liquid-cooled AI data centers, where new components and rack-level failure points must be addressed from the outset. Across DTC, immersion, and RDHx systems, redundancy should extend to new components: pumps, distribution pipes, CDUs, and heat exchangers. These systems are typically backed by dual power feeds, uninterruptible power supplies (UPS), and emergency power protocols to ensure availability. CDUs and pumps are often configured for load sharing to enable seamless failover during maintenance or component failure.

DTC systems require additional safeguards due to minimal thermal inertia and high heat flux. They rely on redundant rack manifolds, server couplings, dual pumps, and heat rejection units with automatic switchover to maintain continuous coolant flow. In immersion and RDHx deployments, redundancy is typically built into CDUs and fluid distribution paths.

Liquid cooling enables AI data centers to support increasing power densities driven by GPU-based workloads, while reducing energy use for thermal management, extending equipment lifecycles, and advancing sustainability and regulatory goals.

Redundancy configurations vary by workload criticality. Inference workloads—latency-sensitive and customer-facing—typically use 2N configurations with fully mirrored power and cooling paths. Training workloads may adopt N+1 setups, balancing fault tolerance with cost efficiency. In high-density AI training clusters, where heat flux is extreme and thermal transients can exceed 1°C per second, even brief cooling disruptions can degrade performance or damage equipment.

To improve both resilience and operational flexibility, many AI data centers are adopting hybrid cooling strategies and modular infrastructure. Hybrid deployments—such as RDHx paired with split air/liquid systems—allow for maintenance without interrupting workloads. Modular CDUs, immersion tanks, and piping loops support scalable growth while integrating redundancy at the component level. Together, these strategies help liquid-cooled environments achieve power usage effectiveness (PUE) ratings as low as 1.05+ – 1.07+, well below the average of 1.5+ for air-cooled systems.

In addition to rack-level planning, liquid cooling influences broader aspects of data center design. Higher power densities may require increased capacity from primary power sources and distribution paths. Heat rejection infrastructure—such as dry coolers or heat exchangers—may require additional external space. During system failover events, rapid temperature rise in the technology cooling system (TCS) requires active redundancy to ensure a safe, seamless transition.

CONCLUSION

Liquid cooling enables AI data centers to support increasing power densities driven by GPU-based workloads, while reducing energy use for thermal management, extending equipment lifecycles, and advancing sustainability and regulatory goals.

However, successful deployment requires careful planning across infrastructure, operations, and safety systems. Many of the considerations outlined in this article remain vendor-specific and lack standardized, widely adopted solutions. As with any emerging technology, early adoption and pilot deployments are key to shaping best practices and identifying optimal design frameworks.

As the liquid cooling ecosystem evolves, new standardized and scalable solutions will emerge, and TIA will incorporate these best practices in future editions of the TIA-942 standard. In the meantime, AI data center operators must strategically weigh trade-offs to ensure they are positioned to adopt and benefit from more resilient, efficient, and cost-effective technologies.

AUTHOR BIOGRAPHIES: Mike Connaughton is currently Senior Product Manager for Leviton Network Solutions and has 30+ years of experience with optical fiber cabling. He is responsible for strategic data center planning, technical account support, and alliances. Mike received his BSEET degree from Wentworth Institute of Technology (Boston, MA) in 1990 and has been involved in optical fiber cable engineering ever since. He can be reached at mike.connaughton@leviton.com.

Jacques Fluet has more than 30 years of experience in telecommunications, including leadership roles at Nortel and Ericsson. He has extensive experience in global product introduction projects, leading diverse teams in product development, verification, and customer trials. As TIA's former Director of Data Center Program, Jacques contributed to technology programs related to 5G, service assurance, smart buildings, and data centers. He can be reached at jfluet@tiaonline.org.

SOURCES:

1. *Hyperscale Computing Market Size, Share & Trends Analysis Report 2023 – 2030*, Grand View Research
2. Tobias Mann, *More than a third of enterprise data centers expect to deploy liquid cooling by 2026*, The Register

LIMIT MANHOLE ACCESS

with **DURASHIELD**



- Strong and durable security
- Easily installed in seconds
- All stainless steel construction
- Requires registered t-Key to unlock
- Installed and removed without bending over


The Intimidator
Line Of Security Products

McGard  mcgard.com/durashield // 888-888-9192

Scan to watch a DuraShield® Video 

McGard, The Intimidator, and DuraShield are registered trademarks of McGard LLC. © 2023 McGard LLC

ISO 9001
CERTIFIED QUALITY
MANUFACTURING SYSTEM
REGISTERED BY NSF-ISR



ADRIAN

MAXIMIZE YOUR TRUCK'S POTENTIAL WITH ADRIAN

With an Adrian truck topper, your truck can work as hard as you do.

TAKE YOUR UPFIT TO THE NEXT LEVEL

ADRIAN.COM

