



Next Generation Data Center Thermal Management: Liquid Cooling

Providing the most efficient and sustainable data center cooling process for optimum equipment performance, critical data protection and cost savings



ENGINEERING YOUR SUCCESS.

Abstract

The interest and adoption of liquid cooling as the preferred approach for thermal management of data centers is being propelled by many simultaneous driving forces. The recent focus on energy efficiency, increasing chip densities and the latest IT advances such as IoT, cloud and harsh edge computing is forcing Data Center managers to take a hard look at best practices for optimizing performance and reducing costs of critical applications.

Introduction

With today's ever-increasing profusion of electronics into remarkably diverse markets, one constant continues to pose an ongoing challenge across most platforms – the generation of heat created by new technologies and processes. Thermal management is especially crucial to ensure optimum operation and performance of all electronic applications, regardless of complexity or scope. Data centers, in particular, consisting of enormous quantities of data processing equipment, have experienced a substantial surge worldwide primarily triggered by the explosion of big data, cloud computing, AI, IoT, edge applications and more recently as a result of the massive work-from-home shift. As companies in virtually every industry begin collecting, analyzing and transporting data, data centers have become vital to driving their success. The heat generated by the extensive electronics found at the core of these centers must be consistently and effectively controlled not only to protect critical data, but for maximum equipment performance, energy efficiency and sustainability. As a result, liquid cooling

is quickly being adopted by data centers as the ideal methodology for maintaining efficient operating environments, especially by hyperscale centers or those with outdated, power-hungry infrastructures using traditional cooling methods.

Liquid Cooling vs. Traditional Cooling Methods


For decades, thermal management of data centers has primarily focused on the combined use of raised floors and air-based cooling systems to deliver cold air to servers. Computer Room Air Conditioners (CRAC) or Computer Room Air Handlers (CRAH) help pressurize the space beneath the raised floor, and perforated floor tiles allow the cold air to enter the space, usually in front of the servers. These systems work to separate hot air from cold, with control of the data environment based almost solely on achieving a specific, desired room temperature.

Many data facilities, consisting of racks of servers and computing equipment with extensive power requirements and extreme heat susceptibility, also incorporate

the use of a hot aisle/cold aisle containment approach to control temperatures. With this practice, multiple server racks are lined up in alternating rows with cold air intakes facing one direction and hot air exhaust vents facing the other.

While these legacy methodologies work well in concept, the desired results are inconsistent and often result in hot spots throughout the data center. These hot spots are caused by airflow problems rather than inadequate cooling capability. As cooling systems operate simultaneously, residual airflow is created, which in turn creates the unwanted hot spots not only within the environment but within the racks and across other equipment as well.

The soaring demand for centers to accommodate the exponential growth of data via more equipment and power requirements, including the use of multi-core chips and other advanced technologies that continually drive temperatures upward, is rapidly making traditional cooling methods obsolete. Increased energy costs in the face of such rising exigencies are forcing data centers to take a closer look



at more efficient, space-saving alternatives for thermal management.

The Many Advantages of Liquid Cooling

Liquid cooling technology was originally conceived in the 1960's, however, it is rapidly gaining major traction again as the ideal process for use within data centers and other industries facing thermal management obsolescence, particularly now that the associated risk of electronics failure due to fluid leaks have been mitigated through innovative leak prevention advancements.

Since liquids have a much higher capacity to capture heat by unit volume – water is at least ten times more efficient than air – liquid cooling is already an accepted process for many HPC applications such as

mining or gaming but is now being widely recognized as an effective method for hyperscale data centers or any environments supporting the use of the latest chips with extremely high computer power. Liquid cooling significantly improves heat dissipation to thwart rising temperatures and subsequent equipment degradation, thereby substantially reducing associated IT hardware maintenance and energy costs.


At the core of Direct to Chip liquid cooling is the Central Distribution Unit (CDU), designed to pump fluids to achieve cooling directly at the primary heat source – the chip. A coolant, either water or dielectric, flows from rack manifolds vertically aligned with the server rack via hoses connected to server chassis or IT equipment. Fluids then come into direct contact with the heat source as it is pumped through fluid lines connected to cold plates covering the server chip (or other designated electronic

components). The flowing fluid absorbs the heat across the chip and carries it to a chiller designed to lower the fluid temperature to the optimal level.

With more than 40% of data center power usage attributed just to cooling efforts, removing heat via liquid cooling has been proven to significantly improve energy efficiency. By removing heat directly at its source, liquid cooling has also been shown to significantly improve processing time and capacity as well as equipment performance, reliability, and life span. While also using only a small percentage of the space required for traditional thermal management technologies, redirected heat to dry coolers or cooling towers may also be reused elsewhere, offering further cost savings and efficiencies.

This innovative thermal management process, however, is not without its own set of challenges

– leak prevention around critical electronics, system pressure maintenance, energy efficiency and space constraints all must be efficiently managed to satisfy the cooling requirements of sensitive electronics and electrical connections.



Leak Tight Connections

It is true that liquid cooling for electronics poses some inherent risks, including dangerous fluid leaks that can result in serious failure of critical equipment and applications. Quick connect couplings are essential components for continuous operation of liquid cooling systems. Most commonly found on CDUs, rack manifolds, water lines and on the backs of servers and IT gear, couplers offer easy serviceability via quick disconnect of equipment from the liquid cooling loop. System leaks and spills occur during connecting and

disconnecting and can also be caused by inferior couplings that are not intended for use in cooling applications. Installing high-quality, stainless steel or plated brass couplings will not only help prevent leakage but will also eliminate chipping associated with plastic components that ultimately deteriorate and enter the cooling system, wreaking havoc. Similarly, the use of valves, seal compounds and materials specifically designed for liquid cooling will help provide optimum system performance and trouble-free maintenance.

The extensive line of Parker Hannifin couplings and seals ensure a leak tight connection for critical data center applications. As the global leader in quick disconnect couplings, Parker's dedicated [Quick Coupling Division](#) offers compact, robust components for liquid cooling systems featuring flush-face, non-spill valving, higher flow rates, and low-pressure drop capabilities to provide maximum resistance to mechanical stresses and fluid loss prevention. Through collaboration and validation with industry leaders and customers, Parker Hannifin products are designed, manufactured, and tested to ensure reliable, dripless connections that are data-proven to be durable for the life of the equipment. In addition, Parker products

are tested via rigorous clean room assembly and helium-leak check quality control measures. As a result, data centers are able to operate at maximum efficiency for increased energy savings.

Conclusion

Proven to be more efficient and cost-effective than traditional air-based cooling, liquid cooling provides specific component cooling to an even greater degree with fluid. Today's evolving data centers are benefiting from liquid cooling systems that offer high-quality dripless connectors designed to prevent fluid loss, reduce pressure drops and accommodate a wide range of sealing materials across varying temperatures. These thermal management hardware advancements are dramatically improving high-computing environments, offering a scalable solution with optimized performance over the long life of their systems. In addition, liquid cooling thermal management is helping data centers reduce critical applications costs and realize maximum energy efficiency and sustainability.

To learn more about Parker's Quick Coupling Division and the critical role of dripless connectors in liquid cooling systems for optimized thermal management, [click here](#).

