Five Reasons to Adopt Liquid Cooling

White Paper 279

Revision 0

by Paul Lin and Tony Day

Executive summary

IT equipment chip densities has been the commonly-discussed driver for adopting liquid cooling. But, there are four other reasons why data center owners should consider liquid cooling including low PUE targets, space constraints, harsh IT environment, and water restrictions. This paper describes these reasons. With this information, data center owners can make an informed decision on whether liquid cooling has advantages for their application.

RATE THIS PAPER ★★★★★

Introduction

Compared with conventional air cooling, liquid cooling provides benefits for data center owners such as higher energy efficiency, smaller footprint, lower TCO, enhanced server reliability, lower noise, etc. However, liquid cooling also has some drawbacks such as higher investment, retrofit or redesign of servers, new data center operation and maintenance skills required, etc. For more information on this topic, see White Paper 265, *Liquid Cooling Technologies for Data Centers and Edge Applications*. Data center owners need to determine if liquid cooling is the right solution that solves the challenges they're facing.

In this paper, we present five reasons to consider adopting liquid cooling:

1. Rising chip and rack densities
2. Pressure to reduce energy consumption
3. Space constraints
4. Water usage restrictions
5. Harsh IT environments

This paper describes each reason in detail and explains how liquid cooling addresses it.

1. Rising chip and rack densities

Changes in IT equipment technology have always been a primary driver in the development of infrastructure cooling solutions. Today, the demands of cloud, IoT, AI, and edge applications are once again resulting in IT technology changes which impact the supporting cooling infrastructure. We summarize some of these changes below:

- **CPU power consumption has increased.** Processor performance continues to improve as the number of cores increases along with processor power. This results in a corresponding increase in CPU heat flux and overall higher heat density within the server enclosure itself. Meanwhile, overclocking is also used to improve compute performance in certain applications like gaming and high-performance computing, which also leads to hotter chips (as shown in **Figure 1**).
- **Increasing use of high-power GPUs.** With its origins in rendering 3D gaming, the GPU is being used alongside the CPU to accelerate computational workloads in areas such as finance, analytics, AI, scientific research, and oil & gas exploration. A CPU has relatively few cores with significant cache memory that can handle a few software threads at a time. GPUs are composed of hundreds of cores that can handle thousands of threads simultaneously and have correspondingly much higher power consumption (as shown in **Figure 1**).
- **Lower latency requirements result in increased heat density.** As component performance increases, the interconnections between these components start becoming a bottleneck in terms of latency. To take full advantage of this improved performance, CPUs, GPUs, and other components on the board such as the memory chipsets, are moving closer together to reduce latency. Resulting in increased physical density and temperatures within the server.

All the above IT technology trends drive the increase in server power consumption and heat density. When a rack is heavily populated with these kinds of servers, the rack power density may be too high to be air-cooled. There are examples of extreme air-cooling such as floor tile fans that can push perimeter cooling systems up to 40 kW rack, or active rear door heat exchangers. Row-based cooling units have been used to provide up to 67 kW for supercomputer HPC installations. However, when these extreme applications are used for more than just a few high-density racks, the cooling system total cost of ownership (TCO) significantly increases. These applications also cause other issues such as high airflow velocities, extreme fan noise, and risk of equipment damage due to a cooling outage.

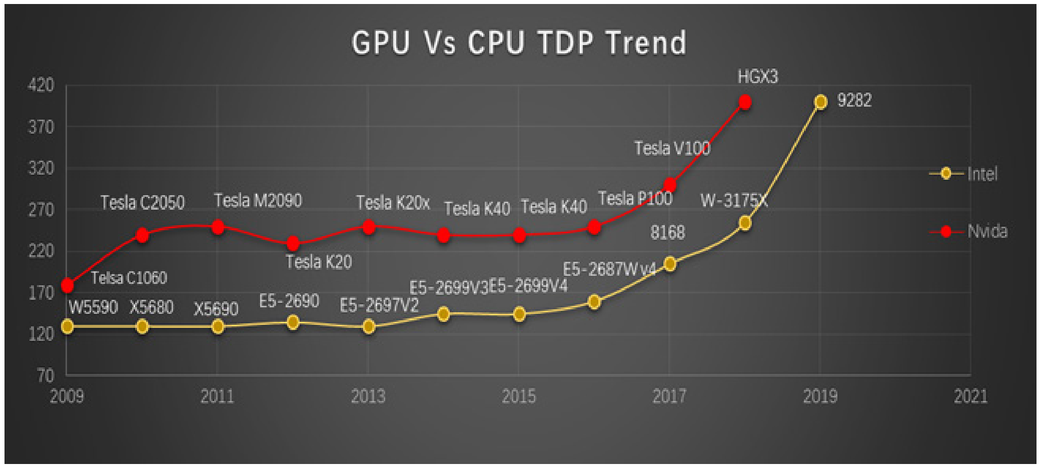

Figure 1
Thermal Design Power (TDP) is trending up

Source: Alibaba[1]

The Green Grid (TGG) suggests a range of 15-25 kW per rack as the limit for air-cooled racks "without the use of additional cooling equipment such as a rear door heat exchanger". Therefore, The Green Grid uses 20 kW per rack in its air-cooled examples within white paper 70, *Liquid Cooling Technology Update*. By contrast, liquid-cooled systems will comfortably deliver twice this capacity and scale to over 100 kW per rack. Liquid cooling provides more predictable thermal control without the over provisioning and airflow management needed in air cooling. Liquid cooling can also improve chip and hard drive reliability by providing a lower stable operating temperature & removing humidity issues.

2. Pressure to reduce energy consumption

As more data centers are built (including cloud, regional, and edge), more energy is consumed by the data center industry, which was about 1% of total global energy consumption in 2017[2]. This high energy use is prompting regulations requiring data center PUEs (power usage effectiveness) to be below a certain value to reduce energy consumption. For example, Shanghai requires that new data center PUEs be below 1.3 and existing data center retrofits be below 1.4. Data center energy expense is becoming an increasingly larger portion of the total cost of ownership (TCO). After IT energy, cooling system energy tends to be the next biggest energy consumer. As a result, data center owners are eager to reduce the cooling system energy consumption.

It's important to note that we normally use PUE as a metric to evaluate the efficiency level of a data center, however, for data centers with liquid cooling, we think energy consumption is a better metric. With liquid cooling, some or all of the IT equipment fans are removed which might lead to worse PUEs despite the fact that the overall electric bill decreases. For immersive liquid cooling, the removal of server fans can lower server energy consumption by approximately 4-15%, depending on the required airflow (8% is typically used for energy analysis). For new or retrofit applications, liquid cooling can also reduce the energy consumption by replacing more energy-intensive air-cooling systems. A workgroup consisting of various US national

[1] https://www.opencompute.org/files/Immersion-Cooling-for-Green-Computing-V1.0.pdf

[2] https://www.iea.org/tcep/buildings/ict/

laboratories and Intel derived a new metric to bring visibility to this issue and iTUE and TUE[3] were created. As of the writing of this paper, use of these metrics is limited, but gaining attention as liquid cooling adoption increases.

3. Space constraints

While air cooling systems can deal with relatively dense IT loads (e.g. deployments running 67+ kW per rack exist), the cost and complications of doing so, increase directly in proportion with the increasing IT load. Although rack count goes down as density increases, the ratio of physical space dedicated to cooling equipment grows higher. This is also true for overhead cooling devices that don't occupy floor space (they sit on top of the rack) because their cooling capacity is generally limited to 20 kW/rack.

Liquid cooling provides an opportunity to reduce the overall data center space for a given IT load. Not only is the white space compacted, but since warm water cooling can be used, reduction of chiller systems and supporting switchgear reduces the required grey space. In a 100% liquid-cooled deployment, hot & cold isle requirements go away, and IT could theoretically be placed almost anywhere. As the demand for IT deployments in urban areas, high rise buildings, and at the edge increase, the need for placement in constrained locations will increase. Liquid cooling can provide the compaction needed to enable compute in places air cooling cannot support.

4. Water usage restrictions

Cooling towers and other evaporative cooling techniques are popular heat rejection solutions for hyperscale data centers because of their high efficiency and large cooling capacity. However, these use the evaporation of water as a heat rejection mechanism, which consumes significant amounts of water (as shown in **Figure 2**). Water usage increases the operation costs, but also in some regions with limited water resources, the local AHJs[4] are putting more pressure on data center owners to reduce water usage. The Green Grid introduced a metric called water usage effectiveness (WUE) to address water usage in data centers in its white paper 35, *Water Usage Effectiveness (WUE™): a Green Grid Data Center Sustainability Metric*.

Figure 2
Cooling towers use the evaporation of water for heat rejection

[3] iTUE (IT-power usage effectiveness); TUE (Total-power usage effectiveness)

[4] Authority having jurisdiction

There are several ways to reduce water usage in data centers including IT load reduction, improve cooling system efficiency, optimize cooling tower operation, etc. Among these ways, liquid cooling can reduce or eliminate water usage from a cooling system design. Because most liquid cooling techniques can use warm water directly to the IT (up to 45°C/113°F), simple dry coolers can be used in most climates to reject the heat.

5. Harsh IT environments

As more IT equipment is deployed at the edge of the network, the requirement to place it in non-ideal environments is increasing. Examples such as IoT in manufacturing facilities and distribution centers can present challenges. The environment is often harsh both in terms of airborne contaminants and the quality of the power supply (**Figure 3** shows two examples). Cost effective standard IT systems are often deployed in these scenarios, sometimes resulting in lower reliability than expected. As IT becomes more integrated with production processes, any downtime affects the output of the manufacturing plant itself.

Oil & gas facility — Automotive manufacturing plant

Figure 3
Two examples of harsh IT environments

Containment in ruggedized enclosures with integrated air conditioning offer one solution, but can be large, costly, and inefficient. Liquid cooling provides and alternative approach in these environments. Immersive liquid cooling isolates servers from the environment. Removing the fans from the IT equipment provides protection against airborne contaminants in harsh environments such as heavy industrial manufacturing plants.

Additional benefits of liquid cooling

While the five reasons above are seen to drive the adoption of liquid cooling, the following benefits of liquid cooling are worth mentioning:

- Minimal heat added to room
- Elimination of fans
- Waste heat recovery
- Less layout complexity (hot/cold aisle & containment)
- Less dependent on geographic climate

Minimal heat added to room

Direct to chip liquid cooling can remove 70-80% of IT heat from the IT space to the outdoor environment, while immersive liquid cooling can remove over 95% of the heat. As a result, it's easier to achieve a comfortable working environment in the IT space as minimal IT heat is added to the room. In contrast, for data centers with traditional air cooling, to improve energy efficiency, the IT inlet and return temperatures are increased, which could result in heat stress to individuals working in the hot aisle for long periods of time. For more information on this topic, see White Paper 123, *Impact of High-Density Hot Aisles on IT Personnel Work Conditions*. Liquid cooling allows for increased density while efficiently maintaining a comfortable working environment.

Elimination of fans

The elimination of fans not only saves 4-15% of energy consumption as discussed above, but also eliminates the occupational health risks caused from fan noise. With traditional air-cooled systems, as rack density increases, the associated increase in facility fan noise not only impacts operational and maintenance staff, which may violate OSHA regulations or ISO7243 guidelines[5], but also increases the risk for IT disk drive failures. Many server designers now find themselves getting very close to the maximum occupational health limits of fan noise associated with servers. Data center operators will face the same health issue due to fan noise from cooling devices.

For liquid cooling, the removal or reduction of server fans and removal or reduction of IT space cooling (e.g. CRAH units or CRAC units), provides a very low noise signature. This is of particular benefit in military and commercial office applications and prevents the occupational health risks from fan noise.

Waste heat recovery

Hot water (over 30°C/86°F) can be used to remove the heat from the chips for direct-to-chip liquid cooling or the heat in the dielectric loops for immersive liquid cooling. The hot return water provides practical recovery of waste heat and energy reuse. The recovered heat can be used for facility or district heating, which can reduce the operational expense and the facility's overall carbon footprint, especially in cold climates. This is very difficult to achieve with air cooling and is an attribute of liquid cooling that municipalities hope to leverage.

Less layout complexity (hot/cold aisle & containment)

For data centers with air cooling, the best practice for airflow management is to have the racks in a cold/hot aisle layout. Containment systems are then used to contain the cold or hot aisle, to separate the cold and hot airflow, in order to save cooling energy and eliminate hot spots. For data centers with immersive liquid cooling, the server and CRAH/CRAC fans are eliminated, therefore there is no concern for airflow management, and no need for cold/hot aisle layout and containment. The rack layout is simplified and is dependent only on the piping paths and maintainability, allowing for configurations currently not imagined.

Less dependent on geographic climate

Air-cooled economizer solutions used to cool IT equipment are dependent on a specific geographic location for their effective deployment. Because liquid cooling uses warm water – up to 45°C/113°F, full economization can be achieved in most parts of the world. This provides advantages in standardization and high efficiency in both regional data centers and edge deployments.

---

[5] OSHA (Occupational Safety & Health Administration) Technical Manual section III, Chapter 4 ISO (International Organization for Standardization) 7243, "Hot environments – Estimation of the heat stress on working man based on WBGT index"

Conclusion

Liquid cooling is most often viewed as a technical solution for high performance computing and high-density applications. But the advantages liquid cooling brings to both large data centers and edge deployments may help overcome many of the challenges data center designers and operators face today. Reducing energy consumption in all geographies, reducing water usage, deployment is space constrained and harsh environments are reasons we expect to see liquid cooling adopted more broadly in the IT industry in coming years.

About the authors

Paul Lin is a Research Director at Schneider Electric's Data Center Science Center. He is responsible for data center design and operation research and consults with clients on risk assessment and design practices to optimize the availability and efficiency of their data center environment. Before joining Schneider Electric, Paul worked as the R&D Project Leader in LG Electronics for several years. He is now designated as a "Data Center Certified Associate", an internationally recognized validation of the knowledge and skills required for a data center professional. He is also a registered HVAC professional engineer. Paul holds a master's degree in mechanical engineering from Jilin University with a background in HVAC and Thermodynamic Engineering.

Tony Day is a Director of Business Development within the IT Division's Office of the CTO at Schneider Electric. He is responsible for working with clients to develop integrated data centre design solutions. He is a Chartered Architect and mechanical engineer with experience in both industrial engineering and construction industries. Prior to joining Schneider Electric, he worked both in professional private practice and within commercial organisations with a focus on construction of highly serviced environments including manufacturing, telecommunications, computer rooms, financial trading floors and data centers.

RATE THIS PAPER ★★★★★

Resources

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

Impact of High-Density Hot Aisles on IT Personnel Work Conditions Conditions
White Paper 123

Browse all
white papers
whitepapers.apc.com

Browse all
TradeOff Tools™
tools.apc.com

Contact us

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

Data Center Science Center
dcsc@schneider-electric.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