ASHRAE TC 9.9 Technical Bulletin
Liquid Cooling: Resiliency Guidance for Cold Plate Deployments

OVERVIEW

Compute workloads continue to push for faster, more powerful, more efficient chips resulting in extreme chip power, lower temperature requirements, and broader use of liquid cooling. The loss of cooling can be catastrophic when supporting extreme chip powers.

This unchartered territory prompted ASHRAE TC 9.9 to approve research on Data Center Direct-to-Chip Liquid Cooling Resiliency (Work Statement #1972).

The extreme chip power is a design and operational challenge.

This technical bulletin introduces the primary concerns and addresses both design and operation recommendations for ensuring resiliency.

- **Primary Concern 1 (throttling)** – reduced computational performance due to temperature excursions within IT components.
- **Primary Concern 2 (rapid temperature rise)** – hardware damage

DESIGN RECOMMENDATIONS OVERVIEW

- **CDU -** Use a coolant distribution unit
- **Thermal Inertia -** Increase the system's thermal inertia
- **Redundancy -** Use active redundancy on TCS/FWS loop pumps
- **Modeling -** Perform transient modeling of the TCS/FWS systems
- **Installation -** Understand the installation planning requirements of the ITE
- **S-Class -** Consider a lower S-Class in a high-density zone for higher power ITE

OPERATION RECOMMENDATIONS OVERVIEW

- **Quality -** Maintain the quality of the cooling fluid in the TCS loop
- **Commissioning -** Perform hydraulic and thermal commissioning
- **Load Migration -** Consider IT load migration strategies
- **Software -** Limit step/block loading/unloading of software

DESIGN RECOMMENDATIONS

Design Recommendation 1 (CDU)
Use a Coolant Distribution Unit (CDU) to ensure demarcation between the Facility Water System (FWS) and Technology Cooling System (TCS) within the data center.

- It must be understood that the TCS and FWS are two different fluid loops in the data center with different requirements.
- The primary issue in the water quality for IT equipment water cooling systems has been a misapplication of water quality recommendations.
- The ITE manufacturer's coolant quality specifications must be followed to ensure a highly reliable TCS.

Design Recommendation 2 (thermal inertia)
Increase thermal inertia to avoid hardware thermal damage due to large load changes and power loss.

- Deploy uninterruptible power systems (UPS) that enable continued operation of the TCS upon a loss of facility utility power.
- Use thermal energy storage (TES) systems to provide cooling water supply in case a chiller plant or cooling source loses power. A stratified TES tank can continue providing cooling to the ITE through the transition to generator power.
- Implement buffer tanks (BT) for the storage of cooling water return (can be either FCW or TCS) to protect the cooling plant from a sharp spike in loads, which could shut down the ability to cool.
- Designs should be concerned about the rate-of-change during recovery as well as the rate-of-change during a failure.

Design Recommendation 3 (redundancy)
Incorporate active redundancy to maintain cooling during changeover from primary to redundant systems.

- Experience by ITE manufacturers has demonstrated as much as 27°C (48.6°F) jump in TCS temperatures due to the changeover of CDU pumps without active redundancy.
- In many cases, high-power device temperatures would not be able to ride through these events.
- Research[1] with medium power processors (relative to today's high-power ITE) indicates that the time between CDU pump failure and ITE processor throttling can occur in as little as 23 seconds.

Design Recommendation 4 (modeling)
Perform transient modeling to verify the performance of systems, products, and components that do not have empirical data from prior testing.

- Computational fluid dynamics (CFD) and flow network modeling co-simulation tools can model a combination of the fluid mass and thermal transport.
- These tools have proven valuable[2] in answering questions pertaining to complex transient operating conditions associated with liquid-to-chip cooling applications.

Design Recommendation 5 (installation)
Obtain and carefully follow manufacturer's requirements
(extreme chip power makes this critical)

- These requirements include (but are not limited to) TCS temperature requirements (S-Classes), flow rates, pressure drop, maximum operating pressure, fluid quality and wetted materials.
- Understand the heat capture ratio and ITE air temperature requirements (recommended and allowable range).
- ASHRAE TC 9.9 is planning to provide a new liquid cooling thermal template this year (2024) in the ASHRAE TC 9.9 Datacom Encyclopedia[3].
- Designers are strongly encouraged to utilize online, software-based planning tools provided by many ITE manufacturers.

Design Recommendation 6 (S-Class)
Consider lower S-Classes (lower inlet temperature) only in a high-density zone

- At the time of publication, the S-classes, corresponding to the maximum inlet temperature to the ITE, range between 30ºC and 50ºC (86ºF to 122ºF).
- Chip power trends could drive towards the need for colder temperatures and TC 9.9 anticipates revisions to the published S-Classes in the future, which will be published in the ASHRAE TC 9.9 Datacom Encyclopedia.
- Lower S-Classes should only be used in a high-density zone to not impact the energy efficiency of the entire site. This could be multiple temperature loops, cascade cooling, and other topologies to minimize the energy impact.
- A lower S-Class can provide additional ride-through time before ITE will initiate a throttling event.

OPERATION RECOMMENDATIONS

Operation Recommendation 1 (quality)
Ensure programs are in place to monitor coolant quality and filtration as directed by the ITE manufacturer

- Cooling fluid of poor quality can cause adverse effects in a liquid system, such as reduced cooling capacity, increased energy consumption, and premature equipment failure.
- Coolant monitoring programs for the TCS will probably differ from those in the FWS.

Operation Recommendation 2 (commissioning)
Use hydraulic and thermal commissioning to validate the design assumptions and systems integrations using thermal test vehicles (TTV) or the actual IT gear.

- Additional guidance on liquid-cooling commissioning is provided in Chapter 11.4 and Appendix B of Liquid Cooling Guidelines, available in the ASHRAE TC 9.9 Datacom Encyclopedia.

Operation Recommendation 3 (load migration)
Use load migration strategies that will work within the timeframe of the minimum server time-to-throttle based on worst-case failure of the resilient design.

- Upon a cooling failure, initiate a live migration of virtual machines or containers in carefully constructed groups to avoid network congestion and allow CPU throttling without significant performance impacts.

Operation Recommendation 4 (software)
Limit software step/block loading/unloading to avoid problems in infrastructure system responses (mechanical and electrical).

- The response time for some infrastructure equipment is much slower than the deployment of software workloads.
- Consider the deployment of workloads in steps/blocks to minimize the need for a rapid response of the infrastructure.

ASHRAE TC 9.9 will continue to solicit industry feedback, publish research, papers, and books, and establish forums for the continued feedback of empirical data and lessons learned from today's higher density liquid-to-chip applications.

REFERENCES

1. Alkharabsheh, et.al. 2018. "Failure Analysis of Direct Liquid Cooling System in Data Centers," Journal of Electronic Packaging, Vol. 140, June 2018.
2. Ellsworth. 2014. Flow Network Analysis of the IBM Power 775 Supercomputer Water Cooling System. IEEE ITherm.
3. ASHRAE TC 9.9 Datacom Encyclopedia: https://datacom.ashrae.org/.

ACKNOWLDGEMENTS

TC 9.9 would like to thank the following contributors for their time and expertise in getting this necessary industry guidance published.

- Tom Davidson (DLB Associates)
- Dustin Demetriou (IBM)
- John Gross (JM Gross Engineering)
- Mark MacDonald (Intel)
- Jason Matteson (nVent)
- David Moss (Strategic Thermal Labs)
- Dave Quirk (DLB Associates) – **Lead Author**
- Justin Seter (DLB Associates)
- Vali Sorell (Oracle)