The Complexities of Shipping Liquid Cooling AI Systems: A Deep Dive

Ellie Chen, Meta Platforms Inc, Austin, TX, USA, elliecjchen@meta.com

Soumya Padmanabha, Meta Platforms Inc, Menlo Park, CA USA, spadmanabha@meta.com

Gautam Nandkishore Nayak, Meta Platforms Inc, Menlo Park, CA USA, gautamnayak@meta.com

Abstract— The demand for high-performance computing systems has increased exponentially. However, traditional air-cooling methods are no longer sufficient to meet the thermal management needs of these complex systems. Liquid cooling has emerged as a game-changer in AI systems, offering improved performance, efficiency, and reliability. As AI hardware becomes increasingly complex, with liquid cooling involved, it poses novel challenges for manufacturing and operation. Liquid-cooled hardware poses unique risks due to its freezing point, volume expansion properties, and material compatibility associated with coolant LC25. In particular, dealing with coolant in hardware products and addressing the challenges associated with shipping these systems and selecting reliable materials are frequently overlooked, leading to potential damage, downtime and in some cases affecting performance. This paper explores the complexities of landing liquid-cooled hardware products, highlighting the need for careful planning and validation to ensure safe and efficient shipment. Specifically, the use of coolant LC25 was found to pose unique challenges due to its freezing point and volume expansion risks. Meta has conducted Long Term Reliability (LTR) testing on liquid-cooled racks, identifying challenges related to material compatibility that can cause performance degradation tracked with performance factors such as coolant flow rate, system pressure drop, and cooling capacity. Root cause and failure analysis revealed oxidation of metal and the presence of silicone as a thread when used in liquid-cooled hardware under higher coolant temperatures, leading to expedited corrosion (metal oxidation) around the GPU cold plates. Air entrainment during handling combined with coolant blocks, can lead to leakage and coolant spillage. The findings emphasize the importance of properly handling liquid cooling AI systems in use with LC25.

Keywords—Liquid cooling AI systems, Coolant properties, LC25, Shipping and handling

I. INTRODUCTION

The rapid growth in high-performance computing demand, especially in AI, has outpaced traditional air-cooling capabilities. Liquid cooling is now critical for effective thermal management, enhancing performance, reliability, and energy efficiency [1].

Liquid cooling, historically used in applications like engine blocks, circulates coolant in a closed loop to transfer heat to ambient air via radiators. In AI hardware, server blade modules act as liquid cooling units where high-power GPUs and CPUs dissipate heat through water blocks with integrated cold plates and pumps. The coolant moves to heat exchangers or radiators assisted by fans. Cold plates, typically copper with microchannels, connect via hoses and valves. Racks assembled with these modules use cold and hot manifolds for coolant distribution, interfacing with facility pipelines or air-assisted racks for heat rejection.

Coolant selection for AI hardware varies, but water or glycol mixtures are most commonly used due to their natural antimicrobial properties, which help maintain system cleanliness and performance. DOWFROST™ LC25 is a preferred heat transfer fluid because of its high thermal efficiency under demanding workloads, superior corrosion protection for copper and other metals, freeze protection down to -10°C (14°F), and fluorescent yellow dye for easy leak detection [2].

Understanding liquid cooling implementation in AI hardware reveals several unique challenges [3]. Liquid cooling introduces complexities in manufacturing, operation, and logistics, categorized as follows:

Coolant Properties: LC25 presents risks due to its freezing point and significant volume expansion [4], which can cause physical damage during shipment or operation if not properly managed. Elevated temperatures could be another factor of degradation.

Shipping and Handling: Liquid-cooled AI hardware requires careful planning for shipping and handling to ensure successful deployment at data centers.

The following chapters will explore each risk in detail, present process solutions to address these challenges, and offer recommendations and best practices for implementing liquid-cooled AI hardware.

II. PROPOSED PROCESS SOLUTION FOR LC25 COOLANT PROPERTIES

Starting with **Coolant Properties**, LC25 is a proprietary coolant blend containing 25% propylene glycol, with strict material compatibility requirements to prevent long-term fouling and corrosion in technical water systems. A key concern is the volume change of LC25 with temperature fluctuations, causing structural risks from expansion and contraction. While flexible components like EPDM or stainless-steel hoses can accommodate some expansion, rigid parts such as copper tubes and cold plates are susceptible to damage if sufficient expansion allowance is not provided.

Potential damage resulting from coolant thermal expansion should be considered in the full rack's design optimization. For instance, if the temperature changes from a controlled room setting (20°C) to exposure near chipsets (up to 70°C), you'll need to calculate the expected expansion:

- Volume of water LC25 in the system:
  $Vs = 50.00$ liters
- Specific volume of LC25 at lower temperature:
  $v1 = 0.00097828\ m^3/kg$
- Specific volume of LC25 at higher temperature:
  $v2 = 0.00100624\ m^3/kg$
- Temperature differential:
  $\Delta t = 50^\circ C$

Calculating Mass:

- Convert liters to cubic meters:
  50.00 liters = $0.05\ m^3$
- Mass in the system:
  Mass = $0.05\ m^3 / 0.00097828\ m^3/kg \approx 51.11\ kg$

Change in fluid volume due to thermal expansion:

$$\Delta V_{fluid} = Vs \times (v2/v1 - 1)$$
$$= 50.00 \times (0.00100624 / 0.00097828 - 1)$$
$$= 50.00 \times (1.0286 - 1)$$
$$= 50.00 \times 0.0286$$
$$= 1.43 \text{ liters}$$

The system should therefore accommodate at least an extra 1.43 liters (about 2.86% of the fluid volume) for thermal expansion.

During early system validation, it is recommended to fill the system to 100% capacity but drain 5% prior to shipment (resulting in a 95% fill, or 2.5L less for a 50L system). After completing robustness testing—which should include thermal cycling, shock and vibration, and leak tests—the system may transition to full 100% fill operation. These tests ensure the system can endure stresses from coolant volume expansion.

These guidelines apply only to liquid-phase coolant and do not account for expansion during freezing, which poses greater risks due to unpredictable ice expansion. Systems containing LC25 coolant must be shipped under temperature-controlled conditions to prevent freezing and potential damage to cold plates. While copper cold plates are generally durable, components with complex features such as soldered joints or skived fins may be more susceptible to leaks or damage under freezing conditions.

To better understand the properties of LC25 coolant at elevated temperatures beyond volume expansion effects, Meta conducted Long Term Reliability (LTR) testing on liquid-cooled racks with LC25 heated to 50°C around the cold plate. Performance metrics tracked included coolant flow rate, system pressure drop, and cooling capacity [5]. The initial 90-day test showed a reduced flow rate (**Figure 1**) and increased system pressure drop (**Figure 2**), while the calculated cooling capacity remained relatively stable with minimal changes (**Figure 3**). The test was then aborted for inspection.

Several cold plates revealed silicone fibers on the copper surfaces (**Figure 4**), which may cause performance degradation and promote corrosion by hindering proper contact between the metal and the corrosion inhibitors in LC25. Root cause and failure analysis identified that polydimethylsiloxane (PDMS), a silicone polymer used as a rubber gasket between the cold plate cover and base, had broken down. This indicates that elevated temperatures not only pose risks of volume expansion but also accelerate the breakdown of certain silicone structures, impacting the coolant properties and overall system performance. Further study on material compatibility with LC25 is needed.

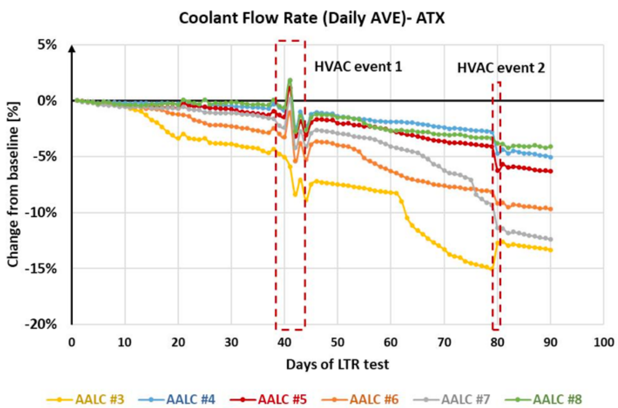

Figure 1: Coolant flow rate of all AALCs dropped by 4% to 13%

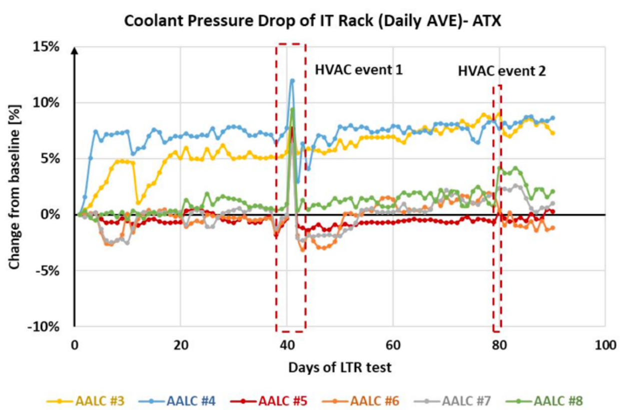

Figure 2: System pressure drop increase

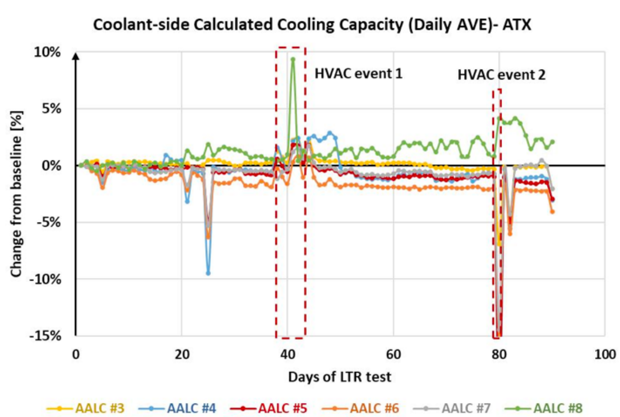

Figure 3: The coolant-side calculated cooling capacity is relatively stable without many changes

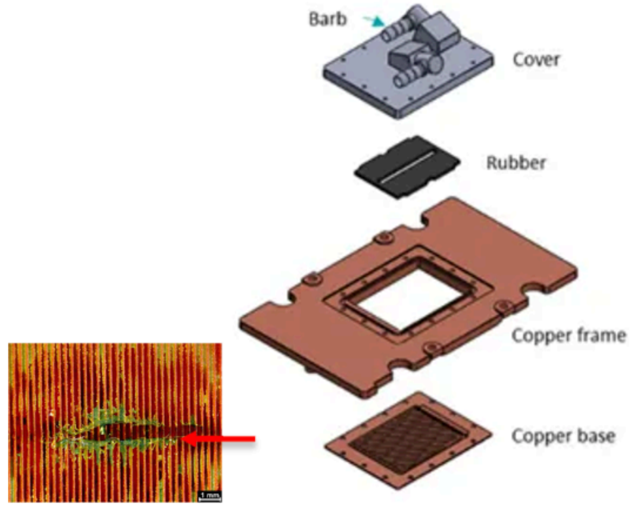

Figure 4: Silicone fibers on copper cold plate

III. PROPOSED PROCESS SOLUTION FOR SHIPPING AND HANDLING

Liquid-cooled hardware may be shipped without coolant, with partial coolant, or fully filled with coolant—commonly referred to as Ship Dry, Ship Semi-Dry, or Ship Wet, respectively. The choice of shipping method affects transportation requirements, pre-shipment preparations, post-arrival handling, and inventory management. Changing the shipment option has far-reaching impacts, but there is no perfect solution yet due to limited experience and data. Below, the trade-offs and risks associated with each option are outlined:

Ship Dry:

- **Definition**: Ship without liquid or after removing as much coolant as possible; typically charged with nitrogen to avoid corrosion.
- **Pros**: No temperature control needed during transport/storage; minimal pre-ship steps.
- **Cons**: Requires multiple preparation steps both before and after shipping, resulting in longer processing times (approximately 2-3 hours per rack), and expensive fill rigs are required at data centers.
- **Impact**: While dry shipping can accommodate a variety of coolants used by customers, it necessitates significant handling—including filling and draining—at data centers.

Ship Wet:

- **Definition**: Ship hardware fully filled with coolant; requires temperature-controlled transport to avoid coolant freezing.
- **Pros**: Low risk of damage from coolant volume expansion; minimal pre-shipment steps.
- **Cons**: Necessitates temperature control during transport/storage; risk of damage due to volume changes if freezing occurs.
- **Impact**: Shipping Wet makes it practical to deliver liquid-cooled products without major retrofits or significant post-rack setup efforts at data centers. However, the added coolant weight increases overall shipping mass.

Ship Semi-Dry:
- **Definition**: Ship with partial coolant and partial air; purges some liquid and exposes interiors to air.
- **Pros**: Minor additional pre- and post-shipment effort.
- **Cons**: Short shelf life (<3 months due to nitrogen pressure retention); higher risk for trays than racks because of increased vulnerability.
- **Impact**: Partial draining does not guarantee protection against pipe bursts caused by ice plugs. Corrosion inspections over time are required, as vaporized water dilutes LC25's inhibitors and accelerates corrosion. Meta has observed corrosion and Verdigris oxidation on copper cold plates (**Figure 5**) with the Ship Semi-Dry method.

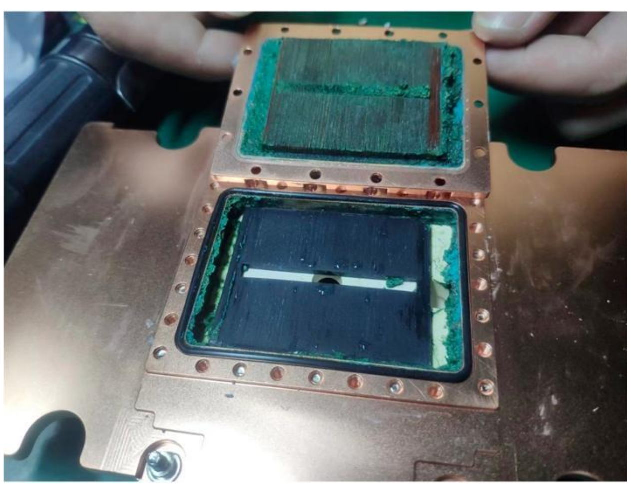

Figure 5: Verdigris oxidation on copper cold plate

Both Ship Dry and Ship Wet are standard shipping methods. Ship Wet enables AI hardware to be delivered safely with minimal post-shipment handling, with the filling process typically managed by the hardware provider. Ship Dry, which eliminates coolant for easier transport, was also evaluated. Two onsite filling approaches were assessed: one with an external filling machine and another using a facility pipeline (where available).

Filling with External Filling Machine:
1. **Vent Nitrogen Charged System**: Nitrogen maintains positive pressure, slowing moisture ingress during transport. Air vents atop rack manifolds release excess gas.
2. **Connect Fill Rig & Set Coolant Flow**: The rig is connected to the rack manifold for recirculation, expelling trapped air/nitrogen from trays/blades. High flow rates improve air removal but should not exceed the maximum operating flow rate.
3. **Unseat & Reseat Trays/Blades**: This is done bottom-to-top as trapped air rises, allowing efficient recirculation per tray. Bottom trays are generally heavier due to longer fluid paths.
4. **Create Coolant Closed Loop**: The entire rack is recirculated to remove air.
5. **Use Vacuum Filling Machine (Optional)**: This equipment can reduce or eliminate recirculation, quickly extracting trapped air compared to manual venting.

Filling with Facility Pipeline:
1. **Vent Nitrogen Charge and Open-Air Vents**: Install external air bleeders with visual indication near the top of the rack manifold or ensure an inline bleeder is present in the facility pipeline above the rack.
2. **Flow Setter & Initial Fill**: Begin with a low flow rate for steady filling.
3. **Flush at Nominal Flow**: Once filled, set to nominal flow rate and flush until all visible air blocks are cleared from the bleeder.

In addition to the coolant filling process, detailed procedures should be established to minimize relative air entrainment during handling. Trapped air, combined with coolant blocks, can lead to leakage and coolant spillage. Since the system operates at elevated temperatures, heated air and coolant tend to push air out of the system, often carrying coolant along with the air block. Systems experiencing leakage during operation typically require shutdown and servicing, resulting in downtime.

The air vent position in the rack should be designed at the top to reduce the risk of coolant escaping through the air vent ports. If multiple air vents are installed to expedite air removal, all vents can be opened during operation. However, during draining—especially after operation—air may be drawn into the system. As explained in the coolant expansion section, coolant contraction occurs when the temperature decreases after system shutdown.

During draining, it is recommended to close the lower air vents and leave only the top air vents open. This reduces air entrainment in the middle section of the system, allowing air at the top to escape efficiently during the next operation. When using a vacuum filling or draining machine, all air vents should be closed to minimize air introduction during the process. Visual explanation is shown in **Figure 6**.

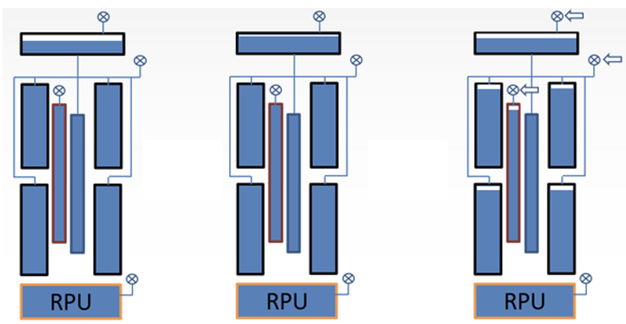

Figure 6: Air at the top of the rack (left); Air escaping during operation with coolant expansion (middle); Air trapped in the middle of the rack via air vents (right)

IV. CONCLUSION AND FUTURE WORKS

This comprehensive examination of the challenges associated with shipping liquid-cooled AI systems underscores the critical importance of robust material selection, chemical compatibility, and rigorous reliability testing. The documented failure modes—particularly metal oxidation and silicone degradation near GPU cold plates—highlight the need for heightened diligence in both component design and material choices.

Challenges related to coolant LC25, including risks from its freezing point and potential volumetric expansion, emphasize the need for precise process controls and targeted shipping protocols. By implementing process enhancements and conducting meticulous validation, this study demonstrates that sustained reliability in liquid-cooled AI hardware is achievable.

These findings reinforce the importance of an industry-wide strategy that prioritizes careful planning, exhaustive testing, and proactive mitigation of chemical and mechanical risks. Adopting the recommendations outlined will enable stakeholders to confidently ensure safe, reliable, and efficient deployment of advanced liquid-cooled AI systems—laying a robust foundation for scalable, high-performance computing infrastructure.

As for the next steps, this research will be focusing on the material compatibility of the coolant LC25 to gain deeper insight. The coolant health will be tracked throughout the test—parameters included copper inhibitor (Tolytriazole, TTZ) levels, pH, reserve alkalinity, and more. A 9-month LTR test has been conducted, and the analysis is currently on going. The corresponding solutions will also be proposed based on the analysis.

REFERENCE

[1] Heydari, A., Shahi, P., Radmard, V., Eslami, B., Chowdhury, U., Saini, S., ... & Rodriguez, J. (2022, October). Liquid to liquid cooling for high heat density liquid cooled data centers. In International Electronic Packaging Technical Conference and Exhibition (Vol. 86557, p. V001T01A007). American Society of Mechanical Engineers.

[2] Pichumani, P. S., Yan, T., & Dillon, E. (2024, October). Failure Analysis of Liquid Cooling Setup for Hyperscale Datacenter Infrastructure. In ISTFA 2024 (pp. 305-311). ASM International.

[3] Chan, A., Nguyen, D., Chen, J., Chen, C. C., & Brooks, M. (2023, March). Coolant Considerations for Liquid-Cooling. In 2023 39th Semiconductor Thermal Measurement, Modeling & Management Symposium (SEMI-THERM) (pp. 1-5). IEEE.

[4] Hang, Y., Zhang, W., Piette, G., Pichumani, P. S., Tseng, J., Mariam, F. A. I., ... & Mithani, N. (2024, May). Long-Term Reliability Test on An Air Assisted Liquid Cooling System. In 2024 23rd IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm) (pp. 1-6). IEEE.

[5] Yunus Cengel and John Cimbala. Fluid Mechanics: Fundamentals and Applications (2024).