INTELLIGENCE UPDATE

Consensus and confusion in liquid cooling maintenance

Jacqueline Davis 2 Jun 2026

Over the past few years, the buildout of high-density data center space for generative AI hardware has led to the deployment of direct liquid cooling (DLC) at an unprecedented scale. This represented a growth of DLC technology out of smaller, specialized supercomputing applications at a rate that outpaced the development of standards and best practices for operating liquid-cooled data center environments. The chief growing pain for DLC was fragmentation in technical details: in system designs, coolant chemistry, resiliency approaches and maintenance procedures.

Accumulated experience from these early years of large-scale DLC has started to converge and produce some identifiable trends in operator preferences to address these problems. Data center operators and IT OEM partnerships tend to favor water cooling systems (specifically with 25% propylene glycol, PG25) over other DLC types for the next few years, though expectations for the longer term are less clear. Recent DLC deployment has trended toward larger (hundreds of kilowatts or megawatt-scale) coolant distribution units (CDUs), often in redundant arrangements to serve rows of IT racks, and operated by facility staff. This suggests growing confidence in one approach to divide maintenance responsibilities between facilities and IT teams, though it is not the only approach being used.

These developments suggest that liquid cooling may be moving beyond its most fragmented early days. Stronger consensus and broadly applicable guidance are likely still years away, especially since megawatt-scale liquid cooling deployments are still concentrated in relatively few facilities. For most operators that oversee any DLC, much of it is still in the stage of research or proof-of-concept. When smaller operators proceed with production DLC deployments, they may do so with different workloads and resiliency objectives, limiting the extent to which they can apply lessons learned from the experiences of hyperscalers.

Coolant chemistry is a compromise

Cooling manufacturers sell several types of cold plate and immersion systems, each using a variety of IT coolants. Manufacturers typically define coolant chemistry early in design, as it involves balancing several trade-offs between thermal performance, pumping costs, chemical compatibility, electrical resistance and the need for coolant monitoring and treatment.

IT vendors, hyperscalers, and other data center operators (enterprise, colocation) largely favor water cold plates, with dielectric-based cold plate and immersion systems being used less at present (see *Investments back two-phase cooling as water cold plate successor*). However, there is considerable variation among operators in how they prioritize the above coolant properties and other cooling system criteria. Correspondingly, not all operators prefer water cold plates and, importantly, not all water-based coolant mixtures are the same. They include additives to protect against corrosion and biological growth. Precise additive mixes vary between manufacturers, but they can be broadly grouped by their use of propylene glycol (PG), a common additive that has become nearly a de facto standard despite thermal performance drawbacks.

Active DLC deployments are increasingly converging on PG25 (75% water, 25% propylene glycol) for cold plate IT cooling. A few years ago, it was more common for manufacturers of cold plate systems — or their IT original equipment manufacturer (OEM) partners, where servers and cooling were sold together — to specify a coolant mix with no PG, based on deionized (DI) water instead. DI water outperforms PG25 in thermal conductivity, specific heat capacity and pumping efficiency. However, it is also more prone to absorb contaminants and metal ions from cooling system components and therefore requires treatment and monitoring.

Where operators are deploying liquid cooling at scale, they are almost universally seeking maximum cooling and compute performance. For this reason, it may seem counterintuitive that operators would favor PG25 over DI water, despite the associated thermal performance and pumping penalties. Yet, consistency across DLC equipment and IT vendors, well-understood material compatibility and familiarity in maintenance procedures (given PG25's widespread use in facility water loops) overrides DI's incremental performance advantage. Facilities staff are often well-versed in maintenance procedures for PG25 in chilled water systems and they can minimize disruption and errors if they apply many of the same principles to the IT coolant loops they are responsible for.

CDU scale, placement and ownership vary

Facilities staff are not always responsible for maintaining the CDU or the IT cooling loop. The line of demarcation between facilities and IT responsibilities in liquid cooling often requires deliberate, project-specific negotiation (see *Hold the line: liquid cooling's division of labor*). Uptime Intelligence briefings with colocation providers and other operators suggest the industry may be settling on a few preferred "responsibility templates" to reduce complexity.

CDUs vary considerably in physical size and cooling capacity, and organizations select them to match the thermal needs of their liquid-cooled servers as well as the scale of that deployment. In-rack CDUs provide coolant to IT in that rack, and larger CDUs serve one or more rows of racks. The CDU form factor is also closely tied to resiliency strategy and maintenance concerns. CDU placement and plumbing determine what redundancy approaches are possible and which staff (facilities or IT) are able to access and maintain the CDUs.

Figure 1 shows four potential boundary options for dividing team responsibilities using a simplified representation of a water cold plate system. Previous Uptime Intelligence reports used these four examples to illustrate the wide range of possible locations for the line of demarcation, but they do not indicate where it might be most effective or practical. Drawing this line still typically requires a negotiation for each individual IT tenant or project.

Figure 1 Facility and IT boundary responsibility options for a DLC system

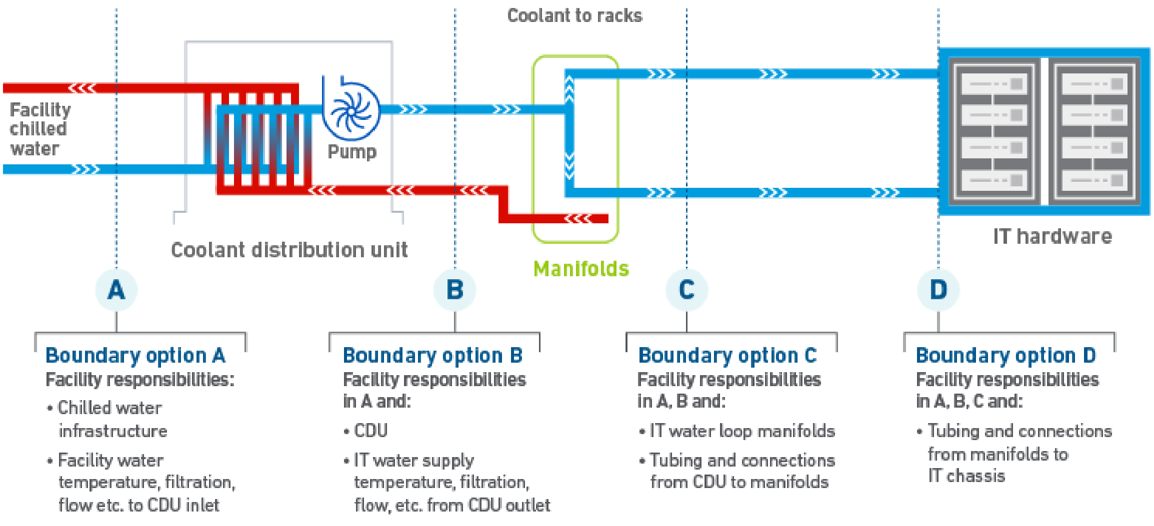

Facilities and IT teams may begin these discussions with incompatible preferences, if each seeks to minimize the amount of DLC equipment that they are responsible for. Management, IT and facilities teams need to clearly define their resiliency objectives and agree on a division of responsibilities that supports them, while also accounting for worker skillsets and training. Negotiating early in the capacity planning process to assign unambiguous ownership of cooling equipment can minimize delays, miscommunication in operations and unnecessary risk.

In deployments of water cold plates, facilities and IT teams tend to reach agreement on boundary options A or B. Both arrangements place rack manifolds and cold plate tubing under the purview of IT, requiring their workers to check and maintain tubing connections. Crucially, the difference in maintenance burden between options A and B can be significant in practice, because this dictates which team owns the CDU. This also tends to shape the resiliency outcomes that organizations expect from the CDU. Each boundary option has its own implications for cooling system maintenance.

Boundary option A:

- Facility teams often prefer this option initially because they can meet their obligations without needing access inside the rack (and sometimes without access into the whitespace).
- The facility is responsible for chilled water supply only, including temperature, filtration and flow to the CDU inlet.
- The IT team procures and maintains the CDU. In practice, this is most often a small CDU mounted in the IT rack and packaged together with the servers. These systems are validated and warrantied by the IT OEM and are often pre-filled with coolant.
- IT staff owns and maintains all IT coolant plumbing and maintains IT coolant chemistry.
- Typically, the IT team alone determines redundancy of CDUs, manifolds and IT coolant plumbing. Redundant CDUs are possible but uncommon. CDUs may have some internal redundancy of pumps, valves, controls, etc.

Boundary option B:

- Large DLC deployments (usually belonging to hyperscalers and often in wholesale colocation) commonly use this option, and it can act as a compromise position between the initial preferences of facilities and IT teams. Facilities staff do not need access inside the rack but they may need access into the whitespace if CDUs are placed at end of row.
- The facility owns and maintains the CDU and is responsible for IT coolant supply, including temperature, filtration, flow etc. from the CDU outlet. In practice, the CDUs are typically high capacity (hundreds of kilowatts to megawatts) and are often installed outside the whitespace itself.
- The CDU is typically specified and procured as part of the data center project, and not as part of the IT procurement. The IT team may have input into CDU specifications and selection. Warranty coverage may be underwritten by a party other than the IT OEM.
- The IT team owns and maintains IT coolant plumbing from the rack manifolds to the IT itself.
- Redundancy of CDUs and plumbing to the manifolds is determined by both facilities and IT; each team can define their own minimum acceptable redundancy. CDUs are commonly N+1 or N+2 redundant. CDU internal components are less likely to be redundant in this arrangement. The IT team alone determines redundancy of manifolds and plumbing to the IT itself.

Boundary options C or D:

- IT teams often prefer one of these options initially. This minimizes training needs for fluid maintenance and minimizes the range of equipment failures that IT staff need to address. Facility staff would require access inside the whitespace and IT racks. Boundary options C and D are uncommon in practice.
- Facility staff own and maintain the CDU and IT coolant parameters as in option B. Options C and D extend facility responsibility into the back of the rack to include the manifolds (option C), and possibly the fluid connections to the IT servers (option D). IT staff own and maintain only the IT.
- In options C and D, the CDU would likely be procured by the data center operator as in option B; this may also be true for rack manifolds and IT coolant plumbing.
- Redundancy of CDUs, manifolds, and IT coolant plumbing would likely be determined by both facilities and IT.

A large portion of present-day liquid cooling deployment uses option B to divide responsibilities, or a very similar arrangement. This approach is most analogous to the one conventionally used for perimeter air cooling equipment, essentially treating a CDU as analogous to a CRAC or a CRAH — usually with a similar resiliency model and preference for redundancy. Even as project-specific negotiation remains commonplace, this approach is increasingly prevalent in large-scale deployments.

True consensus remains elusive

Over the next few years, data center operators (led by the hyperscalers) will continue to gather experience with multiple types of liquid cooling, and standards bodies will continue to develop guidance. These combined efforts will likely move the data center industry closer to consensus on best practices in liquid cooling system maintenance and resiliency design.

However, the divide between hyperscalers and smaller operators in terms of workload types, resiliency expectations, and familiarity with liquid cooling threatens to limit the applicability of past and current liquid cooling experience to all facility types. As long as DLC remains overwhelmingly concentrated in large homogenous AI training superclusters, smaller operators with more stringent resiliency objectives will likely see limited benefit from future convergence in system design or maintenance practices. A more "democratized" liquid cooling future, in which smaller enterprises can operate liquid-cooled environments with lower barriers to entry and robust guidance, is not yet in sight.

> ### The Uptime Intelligence View
>
> DLC's evolution and transition from supercomputing (and mostly into generative AI) was initially marked by fragmentation in DLC designs, coolant chemistry, resiliency approaches and maintenance practices. In 2026, notable trends are emerging in how operators address these problems. This suggests the industry is just beginning to converge on best practices. However, consensus is not yet strong — even among hyperscalers. Guidance has limited applicability for smaller operators.

Other related reports published by Uptime Institute include:
Investments back two-phase cooling as water cold plate successor
Hold the line: liquid cooling's division of labor

ABOUT THE AUTHOR

Jacqueline Davis

5 Jun 2026

Jacqueline is a Research Analyst at Uptime Institute covering global trends and technologies that underpin critical digital infrastructure. Her background includes environmental monitoring and data interpretation in the environmental compliance and health and safety fields.

jdavis@uptimeinstitute.com

About Uptime Institute

Uptime Institute is the Global Digital Infrastructure Authority. With over 4,000 awards issued in over 122 countries around the globe, and over 1,100 currently active projects in 80+ countries, Uptime has helped tens of thousands of companies optimize critical IT assets while managing costs, resources, and efficiency. For over 30 years, the company has established industry-leading benchmarks for data center performance, resilience, sustainability, and efficiency, which provide customers assurance that their digital infrastructure can perform across a wide array of operating conditions at a level consistent with their individual business needs. Uptime's Tier Standard is the IT industry's most trusted and adopted global standard for the design, construction, and operation of data centers.

Offerings include the organization's Tier Standard and Certifications, Management & Operations reviews and assessments including SCIRA-FSI financial sector risk assessment, the Sustainability Assessment, and a broad range of additional risk management, performance, availability, and related offerings. Uptime Education training programs have been successfully completed by over 100,000 data center professionals, such as the much-valued ATD (Accredited Tier Designer) and AOS (Accredited Operations Specialist). The Uptime Education curriculum has been expanded by the acquisition of CNet Training Ltd. In 2023.

Uptime Institute is headquartered in New York, NY, with offices in London, Sao Paulo, Dubai, Riyadh, and Singapore, and full-time Uptime professionals based in over thirty-four countries around the world.

For more information, visit www.uptimeinstitute.com