BRIEFING REPORT

Resiliency considerations with direct liquid cooling

Jacqueline Davis, Research Analyst jdavis@uptimeinstitute.com
Daniel Bizo, Research Director dbizo@uptimeinstitute.com

Direct liquid cooling (DLC) promises attractive thermal performance and economics, and vendors are commercializing cold plate and immersion systems in anticipation of growing demand. Use of the technology is now set to expand beyond specialized technical computing applications, making its first significant inroads into mission-critical installations that tend to place strict demands on cooling resiliency.

Resiliency standards and best practices, however, have been shaped by the near-universal use of air cooling in data centers. DLC's design differences raise concerns for cooling-system resiliency because they introduce new failure scenarios and complicate the allocation of responsibilities between facilities and IT teams. Organizations deploying DLC face a shortage of applicable guidance and need to compensate with thorough planning and negotiation between staff teams. This report suggests guiding questions to define resiliency, beginning with the primary business objectives.

KEY POINTS

- DLC requires organizational changes: DLC techniques, such as cold plate and immersion, render the conventional relationship between data center facilities and IT teams inadequate — they will require much closer collaboration. This has implications for the entire data center life cycle, from design to operations.
- There is an absence of recognized DLC resiliency implementation standards: the resiliency profile of liquid-cooled IT hardware is both more complicated than, and inherently different from, that of air. Unlike with air cooling, there is currently no industry consensus on how to attain concurrent maintainability or fault tolerance in a DLC environment. Resiliency approaches to DLC implementations will likely fragment until guidance and practices mature.
- An integrated infrastructure resiliency posture may be required where DLC is used: the complexities and costs involved with achieving high resiliency standards with DLC invite an expanded view of resiliency, one that includes software resiliency measures and does not strictly prescribe cooling components or physical redundancy.

Introduction

Uptime Intelligence's research finds that the data center industry is widely anticipating that liquid-cooled IT hardware will become commonplace in data centers in the coming years, although the expectations around the speed and scale of this shift vary greatly. Infrastructure operators are expecting dense IT racks, which call for direct liquid cooling (DLC), to become more widespread in facilities.

Yet it is not ultra high-density racks that will make the technical case for mainstream liquid cooling, but server silicon power. With escalating processor thermal design power and corresponding temperature restrictions, the continued use of air cooling will prove ever more taxing in runaway fan power, high levels of silicon power leakage and untapped performance potential. Widespread enthusiasm for complex artificial intelligence (AI) is adding fuel to the drive for ever higher silicon power in pursuit of compute performance. This in turn has led many in the industry to contemplate the need for upgrades to facility cooling (and power).

Switching to liquids to cool IT electronics brings immediate gains in server energy performance and outright processing performance. As a bonus, it also opens up opportunities to boost infrastructure sustainability credentials via lowered compressor and water use. In selected cases, DLC is also more suitable for heat reuse schemes because of its ability to produce higher-temperature waste heat.

In light of these technical trends and sustainability considerations, data center industry stakeholders expect air-cooled IT to gradually give way to liquid-cooled IT in the coming years. A recent Uptime Intelligence survey on cooling, conducted earlier in 2023 (see *Direct liquid cooling shows promise, but challenges persist*), indicates an industry consensus that in about six years, liquid-cooled IT will become dominant in data centers that are larger than 1 megawatt in capacity. Furthermore, only one in five operators say they have no intention of deploying DLC in their facilities; a similar number say they are already using it, with the majority saying they will consider using the technology if necessary.

Placing DLC and air cooling in context

Among the data center operators with little to no DLC deployment today, those making long-term DLC plans often aspire to high resiliency standards that are comparable with air cooling. But many current DLC products, even those marketed for large-scale deployment, may not offer the concurrent maintainability and fault-tolerance features that many infrastructure operators require and have come to expect with air-cooling systems.

Decades ago, water cold plates cooled a large install base of mission-critical IT systems in the form of mainframes, which had a befitting redundancy design across the integrated DLC components. Then, as many mainframes were displaced by air-cooled volume servers, DLC found a growing application in high-density IT hardware, supporting technical computing workloads, such as scientific and engineering simulations, cryptocurrency mining and (more recently) training deep neural networks.

These workloads typically have several key differences from those of most businesses. Computing is often performed in batches rather than in response to user requests; an interruption from a thermal shutdown or equipment failure does not directly affect organizational revenue; and the loss of a single server (or node) may not affect the job's completion.

This stands in contrast to mission-critical workloads, which may be expected to handle end-user requests 24 hours a day with low latency — economic constraints may also prevent duplication or sharing of workloads to insulate business operations from a server failure.

In recent years, vendors have started marketing DLC systems for mainstream use, supporting all types of business applications in the data center. However, because the designs of many DLC products on the market were first developed for technical computing workloads, they are not geared to meet the demands of availability and serviceability required by business-critical applications. As such, vendors still tend to primarily compete on cooling performance and cost, and often promote extreme densification of compute alongside maximized free cooling for heat rejection.

This is partly because, for all its potential benefits, the switch to DLC raises many issues. One of these relates to the disruption of the current division of facilities and IT infrastructure functions. This is because adoption at scale of DLC (beyond its more established use in high-performance computing) calls for a resetting of the boundary of responsibilities and requires much closer cooperation between facilities and IT infrastructure functions.

Moving to DLC is not a trivial effort. Design differences fundamental to DLC systems complicate the conventional model for cooling system planning and implementation. Worse still, a lack of standards offers data center and IT operators no clear path to implement DLC systems to high, uniform levels of resiliency that are comparable with the definitions used for air-cooled infrastructures in mission-critical environments.

This report outlines the main points that data center infrastructure operators will want to consider when it comes to resiliency planning for DLC implementations. In preparation for this report, Uptime Institute Intelligence spoke to many industry experts, including vendors, data center operators and worked with Uptime Institute's consulting engineers. However, please note that this research report does not constitute a formal position of Uptime Institute.

Defining objectives and equipment

This report compares some of the resiliency issues across different types of DLC, as well as with air cooling, evaluating their business and technical characteristics. Some manufacturers and standards bodies may use differing terminology to describe operational objectives and cooling system components. The definitions below are used throughout this report to compare the various cooling systems and their characteristics with consistent language.

Operational objectives in cooling

Data center cooling systems are engineered to meet the thermal management needs of IT systems — this includes not only thermal performance but also resiliency considerations as part of securing IT service availability.

Organizations that use or are considering the adoption of DLC are typically responding to a need to support high rack power, which can be costly or technically impractical to cool with air (in general, racks above 20 kilowatts are regarded as high power). Historically, many of these racks support scientific or technical computing tasks that do not place strict demands on system availability.

However, as the use of DLC gradually expands beyond traditional high-performance computing, more and more critical IT systems will be liquid-cooled. Businesses will not, however, accept compromises to resiliency standards and realized application availability for a new mode of cooling, regardless of the technical or economic benefits.

Data center operators consider several design and operational objectives when characterizing the resiliency of their cooling infrastructure and IT hardware (see **Table 1**).

Table 1 Key cooling system resiliency considerations and design objectives

Topological resiliency	Specifically, **concurrent maintainability**, where any part of the system may be shut down for planned maintenance without interrupting the cooling or IT operation, and **fault tolerance**, in which a failure does not interrupt the operation of IT equipment (owing to redundant isolated capacity components and paths, as well as automated failovers when needed). Lower resiliency standards, such as selective redundancy for some critical cooling components, and non-redundant basic capacity, are more straightforward to achieve and are not discussed in this report.
Ride-through time	The calculated worst-case time between a partial or full loss of cooling capacity and the resulting interruption to IT operations. In some scenarios, there is a minimum requirement for ride-through time to allow data center infrastructure controls to react to the event. A typical example is a loss of utility power that shuts down some or all of the cooling units, which need to reinitialize once the generator sets can take the load. DLC systems usually involve much shorter ride-through times than air-cooled systems.
Maintenance and operations	Clear protocols for facilities management, especially the division of responsibilities between facilities teams and IT teams. In the fragmented area of DLC systems without technical standards to govern engineering choices and operational practices, this division and the implications for infrastructure resiliency requirements are unclear.
Failure domain	Also known as “blast radius”, this is the number of IT systems and applications rendered inoperable or degraded (i.e., reduced performance, features or service availability) from a failure in the cooling system. Liquid coolant delivery mechanisms tend to add new risks and expand failure domains.
Interoperability	Design features that enable a DLC system to share infrastructure, such as chilled water loops, with cooling systems of different designs — including air-cooling equipment or differing DLC systems — and to integrate with existing monitoring and control systems.
Improved IT hardware reliability	The reliability of IT hardware can affect IT service availability. The use of DLC should, theoretically at least, have a positive effect on the failure rates of IT hardware. This is because DLC reduces thermal stress due to better temperature stability; it lowers exposure to vibration; and immersion systems can protect IT hardware against the effects of humidity and corrosive gases. This upside, however, remains largely unproven due to limited data as well as other DLC reliability factors weighing on the balance.

UPTIME INSTITUTE 2023 uptime INTELLIGENCE

Types of direct liquid cooling

The use of liquid to absorb and remove heat from IT electronics, as opposed to air, is what makes a fundamental difference to the technical and business characteristics of DLC. As such, the scope of this report includes liquid-to-air (also called air-assisted liquid cooling) implementations where the liquid that cools the IT hardware exchanges heat with air as opposed to facility water. Such systems include, for example, air-cooled coolant distribution units (CDUs) and liquid-to-air heat exchangers mounted at the rear of the rack.

This report excludes in-row or in-rack cooling methods that deliver liquid to the rack but stop short of the IT hardware, which remains air-cooled. Even if they often target similar use cases in densified computing, their use of air to extract heat from IT components affords a continuity in technical and business conventions that DLC fundamentally disrupts.

Commercialized, widely available DLC technology generally falls into one of two categories: cold plates or immersion, with each comprising three sub-classes.

Cold plate systems use closed-loop heat sinks: a thermally conductive plate is mounted on the chip's casing to transfer its heat to a liquid coolant. As the cooling fluid flows through the plate, it draws away the heat dissipated by the electronics. Cold plates typically cool only the most powerful components of IT, such as processors and other high-performance logic and, in some cases, memory modules.

As a rule of thumb, these components account for 70% to 80% of a compute server's heat output. This is not a theoretical limit of cold plates, but a practical economic balance: supplemental air cooling can handle the remaining low-density heat. There are exceptions in supercomputing where the cold plate system is used to cool the full heat load of the server due to the close spacing of dense compute hardware.

Uptime Intelligence recognizes three main types of commercially available cold plates, differentiated by the type of liquid used.

- **Water cold plates** use a form of treated water. Water is attractive as a coolant because it is abundant and low-cost — and has high performance in terms of thermal conductivity and energy storage capacity. However, possible water leakage carries a substantial risk to the IT equipment, which remains a key concern for general-purpose scale deployments in mission-critical environments. Operators should assess different water chemistries for thermal performance, material compatibility needs, long-term component reliability and related equipment maintenance requirements, but these considerations do not affect the analysis of this report. Because the operational risk profiles of the various types of water chemistries are similar, this report does not differentiate between them.
- **Single-phase dielectric cold plates** use an engineered dielectric fluid instead of water. This eliminates water leakage concerns and enables the liquid chemistry to be shared with immersion techniques in a mixed-design environment. However, such fluids have thermal performance that is inferior to that of water, as measured by thermal conductivity and energy storage capacity (specific heat). To achieve the same cooling capacity, these systems require a higher flow rate. Importantly, the thermal properties mean reduced ride-through times during some cooling failures, such as a pump failure in the DLC loop or the loss of a facility water loop.
- **Two-phase dielectric cold plates** use engineered dielectric fluids with a suitably low boiling point to carry away heat from server chips as latent heat, which is thermal energy absorbed by vaporization (a phase change from liquid to gaseous form). The vapor pressure that develops at the chip also drives circulation — vaporized coolant naturally flows towards the heat exchanger where it condenses, resulting in a drop in pressure. Depending on the implementation, an evaporator cold plate with a small pool of coolant may offer longer ride-through times because the fluid vaporization continues passively in the event of a cooling system failure, including major leaks or disconnection on the return side of the cold plate.

In an immersion system, the IT hardware is flooded with a dielectric liquid for cooling. It is sometimes referred to as total liquid cooling because it absorbs virtually all of the heat load of the cooled IT system. In all cases, an engineered fluid comes into contact with IT hardware components — potentially introducing new risks to the health of IT systems from fluid contamination or chemical leaching. This calls for a host of installation and maintenance protocols around the chemical and particle "hygiene" of immersion systems.

- **Single-phase chassis immersion**, sometimes referred to as partial immersion, is a technique in which a single server is flooded with a small pool of dielectric fluid inside a sealed chassis with a conventional form factor. Major benefits include integration with standard vertical air-cooled racks, using only as much dielectric liquid as needed and the ability to match coolant flow rates to individual compute requirements. Some maintenance tasks are also simpler to carry out because the sealed chassis is easy to remove and handle with a reduced risk of dripping. However, the thermal ride-through time is shorter compared with tank immersion because of the smaller volume of liquid that surrounds the server.
- **Single-phase tank immersion** involves submerging the IT hardware in a tank filled with a non-conductive liquid coolant that remains in a liquid phase. Fluid is circulated by pumps or by passive convection. The larger volume of coolant (compared with cold plate and chassis systems) can offer prolonged ride-through times during cooling system failures.
- **Two-phase tank immersion** submerges the IT hardware in a dielectric coolant with a low boiling point that changes from liquid to vapor as it absorbs heat. Two-phase tank immersion systems gain a performance advantage by using the latent heat of vaporization to cool the entire server. Vapor rises to the top of the tank, condenses back to a liquid as it transfers heat energy to the heat exchangers, then falls back into the tank for reuse. There are several key operational challenges surrounding two-phase tank immersion.
    - Dielectric coolant can be lost as fugitive vapor.
    - Coolant formulas require extensive testing to avoid chemical incompatibility with IT and power electronics that can cause abnormally high IT failure rates.
    - Opening the tank to service individual server boards also opens the coolant loop and is likely to affect all of the IT hardware in the tank, requiring all servers to reduce performance — if not shut down completely.

Future developments in DLC techniques may enable new cooling system designs that have markedly different maintenance and resiliency characteristics. In any case, these will take years to reach maturity and gain significant commercial availability.

Conceptually, while these DLC systems have individual underlying risk profiles and maintenance requirements to consider, they share major topological distinctions when compared with air-cooling architectures. This report discusses DLC resiliency in general and makes a distinction between the various types of DLC only when appropriate for additional detail or clarity.

Operational and resiliency differences with DLC

Even though liquid cooling is not a novelty in the history of IT (the first water-cooled computers appeared in the 1960s), data centers today are (de facto) standardized around air-cooled commodity hardware. Operations and maintenance practices have matured, along with equipment designs, around the predominant use of air cooling. In data center operations, liquid-cooled systems are treated as exceptions rather than the rule.

Air acts as a standard “line” of demarcation where responsibility for components and topological infrastructure resiliency shifts from facilities teams to IT teams. Typically, the arrangement is intuitive. Data center operations supply conditioned air to the front of the racks within acceptable temperature and humidity ranges; then within the rack, IT system fans drive the airflow to cool the electronics and any mechanical drives or tapes. The relationship between air handlers and IT hardware remains indirect, with no mechanical integration. This also means that topological cooling infrastructure resiliency goes only as far as producing the supply of air.

DLC systems, whether cold plate or immersion, differ from this convention in that DLC equipment and components deliver the coolant directly to the IT electronics — IT hardware typically does not play an active role in the circulation of the liquid coolant. (Some servers use a self-contained loop with a mini pump and a liquid-to-air radiator, but because these systems do not receive liquid coolant from external equipment, they are not in the scope of this report.)

Since the cooling equipment and IT are tightly integrated, there are more possibilities for component failures and disruptions to distribution paths that would affect cooling.

Figure 1 illustrates this difference. Note that air-assisted DLC has been included because it is also an alternative form of heat exchange between facility equipment and the IT coolant. For simplicity, coolant return manifolds have been omitted.

Figure 1 Defining responsibilities in air cooling and DLC systems

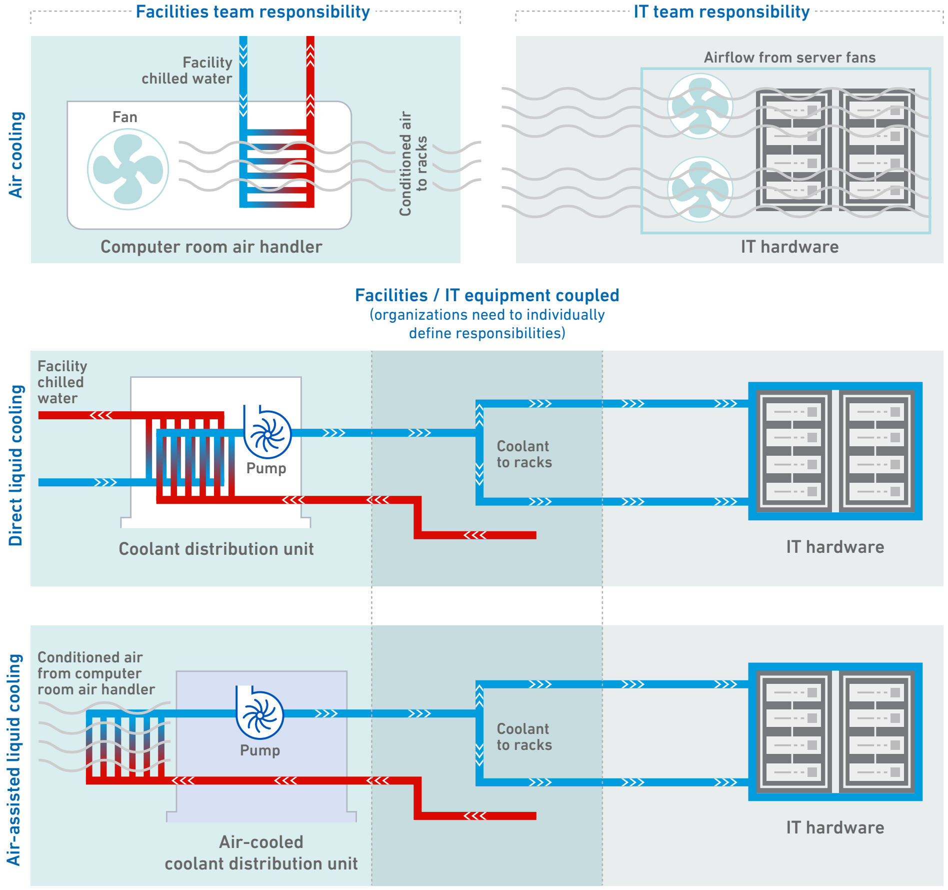

UPTIME INSTITUTE 2023

With air cooling, the hand-off point between facilities and IT teams is a useful convention that is effectively embedded in decision-making throughout the data center life cycle. It defines not just operation and maintenance responsibilities but also other functions such as commissioning and procurement.

Extensive use of DLC in mission-critical facilities calls for a renegotiation of this divide and a revisit of topological considerations as the integration between facilities and IT becomes tighter. Rather than producing a shared, large volume of cold air for uptake by IT system fans, DLC delivers a coolant (water or dielectric fluid) to IT electronics, crossing the boundary into the chassis — or, in the case of immersion, into the tank. Any disruption to the distribution of liquid coolant has the potential to immediately affect IT systems, and likely many of them.

However, where facility responsibility (and crucially, topology) ends and IT responsibility begins is far from clear-cut — and there is certainly no industry consensus on this. Variations in DLC system design further complicate the interface between the domains of facilities and IT. Even if the technical organization (and in colocation spaces, the customers) reach an agreement on a clear paradigm of responsibilities and resiliency, it will likely not be applicable across all the types of DLC systems, such as cold plates and immersion.

Even variations within one type of DLC can frustrate efforts to establish a stable consensus on the hand-off point, as well as topological resiliency standards. The category of water cold plate systems contains examples with markedly different approaches to procurement and resiliency. Some vendors produce centralized CDUs to cool multiple racks with redundancy, marketed primarily to data center operators. Other CDUs are sold integrated with a pre-racked IT system, often in the IT rack itself, with no redundant coolant distribution or ability to maintain without downtime. These integrated systems are more often purchased by IT teams with little input from facilities teams. **Figure 2** highlights four major boundary options between facilities and IT responsibilities.

Figure 2 Facility / IT boundary options for a DLC system

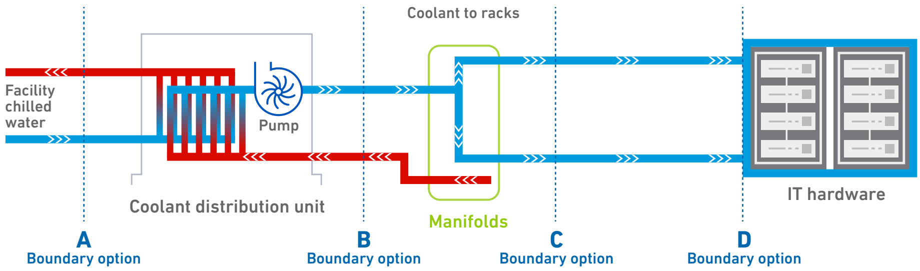

UPTIME INSTITUTE 2023

To help interpret these major boundary options, **Table 2** offers brief descriptions and further notes. Facilities and IT teams can manage expectations and minimize operational risks if they agree on a standard approach in the division of responsibilities. Note that regardless of the boundary, detailed maintenance and operations procedures, specifically for points of cooperation between facilities and IT, are necessary to avoid confusion and misunderstanding.

Table 2 Explanation of facility / IT boundary options for DLC systems

Boundary option	Description	Notes
A	The boundary is at the inlet / outlet connections between facility water pipes and the CDUs.	Facility responsibility stops with delivering water to a CDU (or air to an air-cooled DLC system). Downstream topological resiliency is beyond the scope of the data center physical infrastructure.
B	The CDU is part of the facility resiliency posture, but not downstream components to racks / tanks.	Higher resiliency standards call for either redundant CDUs connected through valves or CDU designs that can be maintained without downtime.
C	Coolant distribution paths and manifolds (often in rack) are included in physical infrastructure resiliency.	Tubing and manifolds are considered for topological redundancy to offer path diversity for coolant supply and return. Failover and isolation can require extensive use of valves.
D	Coolant inlet / outlet of IT chassis or tank internal space.	Correct tube connection / tank installation of IT equipment is facility responsibility.

UPTIME INSTITUTE 2023

Design differences, both between and within types of DLC, can have far-reaching implications for redundancy, commissioning, maintenance and risk profiles. If data center operators start deploying liquid cooling at scale (as some already have for specific services), a singular convention for the division of responsibilities and resiliency scope between facilities and IT will likely lack the flexibility to cover all deployment scenarios.

Conceptually, boundary options C or D (where liquid coolant distribution to the rack or the IT equipment is part of the facility domain) represent practices where facilities management ensures resiliency design and maintenance can be brought to high, uniform standards. Equally, this division may also be favored by IT teams because it does not require them to acquire new skills in mechanical and chemical engineering. This approach is analogous with common power delivery practices in which the data center provides two independent power distribution units in the rack for topological redundancy.

However, the reality of DLC deployment today is that IT systems typically come pre-integrated with specific flavors of DLC systems, along with their own CDUs (or other forms of heat exchange units), shipped and supported by the IT vendor. Even if technically compatible with an operator's preferred CDU (sometimes they are even the same make under a different brand), IT vendors may be reluctant to approve their hardware for use with third-party DLC systems due to commercial and legal liability reasons.

For cost reasons, most data centers use uninterruptible power supply systems to protect IT against power cuts, but not the cooling system.

For this reason, many will find boundary option A attractive even though it is fraught with issues. Chiefly, this means it is now IT's responsibility to look after the CDU and all downstream components, including valves and connections, which all affect the risk profile of the operation. And even if facilities assume the maintenance and operational responsibility for such integrated DLC systems, those deployments will likely fall short of the data center's resiliency standard.

Industry resiliency standards, including terminology, may also need to evolve to adapt to DLC — for example, ride-through time. For cost reasons, most data centers use uninterruptible power supply (UPS) systems to protect IT against power cuts, but not the cooling system. If there is a loss of power, cooling typically stops and reinitializes only when the generators have started, which can take up to a minute. In a typical air-cooled data hall, this will tend to raise the temperature, although not enough to impact IT availability. Only fault-tolerant data center resiliency standards require continuous cooling.

This convention may not be appropriate for all forms of DLC. Some implementations (often single-phase immersion tanks) naturally offer long ride-through times. They can keep IT operating without issue, often for tens of minutes after a loss of power to cooling — especially when starting at moderate supply temperatures or a partial server load. Others, particularly cold plate systems that operate near their cooling capacity, may offer only a few seconds of ride-through time.

This means that data center designs not requiring fault tolerance — but instead concurrent maintainability or simply redundancy — may need to protect cooling systems (or their components: CDUs, facility water pumps and computer room air handlers) with UPS power.

Moreover, existing definitions for continuous cooling tend to specify measures (i.e., chilled water storage, plus power-protected pumps and fans) for air-cooling methods. Some of these measures may help DLC systems that use the same heat rejection infrastructure, but effective guidance on maintaining the cooling capacity through a power cut will need explicit inclusion of CDUs (and, if present, the dedicated heat-rejection system) in the list of equipment required to be connected to UPS systems.

Operational procedures and DLC infrastructure design will require careful coordination to address not just familiar issues, such as loss of utility power, but some failure scenarios that are specific to DLC as well. **Table 3** highlights some example failure scenarios that could affect DLC systems, including the expected effects on IT and design measures that can mitigate risk. The list is not exhaustive. "High-heat components" refers to those IT components (CPU, GPU and some memory modules) that are responsible for the majority of IT heat load. In cold plate systems, the high-heat components are typically those that are attached to the cold plates in the IT coolant loop.

Table 3 Example DLC failure scenarios

Scenario 1: Total power loss to IT coolant pumps	
Likelihood	Likely — usually results from utility power outage
Cooling types	Water cold plates*
	Single-phase dielectric cold plates*
	Two-phase dielectric cold plates*
	Single-phase tank immersion
	**Except systems where pumps are inside IT chassis and protected by UPS systems*
Estimated time before IT affected	Single-phase cold plate systems (water or dielectric): seconds
	Two-phase dielectric cold plates: about 30 to 60 seconds
	Single-phase tank immersion: minutes
Effects on IT	High-heat components can rapidly increase in temperature, leading to throttling, protective shutdown or chip damage
Mitigation or prevention	Supply IT coolant pumps with UPS backup power to continue coolant circulation

Scenario 2: Failure of IT coolant pump	
Likelihood	Moderate — depends on pump design and maintenance
Cooling types	Water cold plates
	Single-phase dielectric cold plates
	Two-phase dielectric cold plates
	Single-phase chassis immersion
	Single-phase tank immersion
Estimated time before IT affected	Single-phase cold plate systems (water or dielectric): seconds
	Single-phase chassis immersion: seconds
	Two-phase dielectric cold plates: about 30 to 60 seconds
	Single-phase tank immersion: minutes
Effects on IT	High-heat components can rapidly increase in temperature, leading to throttling, protective shutdown or chip damage
Mitigation or prevention	Consider a DLC system that can be implemented with redundant CDU or redundant coolant pumps
	Implement software resiliency to shift critical workloads away from affected systems in failure / maintenance

Table 3 Example DLC failure scenarios (cont.)

Scenario 3: Water leak inside IT chassis	
Likelihood	Low to moderate — increases with number of tubing connections, frequency of maintenance
Cooling types	Water cold plates
Estimated time before IT affected	Effectively zero
Effects on IT	Leaked water contacts IT electronics inside chassis, rapidly causing short-circuit damage that typically requires complete replacement of IT hardware
	Water likely to spill into other IT hardware in the rack
	Isolation and repair likely to require shutdown of all IT on the affected coolant loop
Mitigation or prevention	Implement software resiliency to shift critical workloads away from affected systems in failure / maintenance
	Consider a cold plate system with negative pressure IT loop to eliminate short-circuit risk and allow time for operator intervention
	Consider a dielectric cold plate system to eliminate short-circuit risk
	Consider a DLC system with redundant tubing and manifolds to facilitate isolation and repair (not always available from manufacturers)

Scenario 4: Dielectric fluid leak from cold plate system	
Likelihood	Low to moderate — increases with number of tubing connections, frequency of maintenance
Cooling types	Single-phase dielectric cold plates
	Two-phase dielectric cold plates
Estimated time before IT affected	Single-phase dielectric cold plates, large leak: seconds
	Two-phase dielectric cold plates, large leak: about 1 minute
	All dielectric cold plates, small leaks: hours to days
Effects on IT	Isolation and repair likely to require shutdown of all IT on the affected coolant loop
	If coolant loss is sufficient to stop supply to IT: high-heat components can rapidly increase in temperature, leading to throttling, protective shutdown or chip damage
Mitigation or prevention	Implement software resiliency to shift critical workloads away from affected systems in failure / maintenance
	Consider a DLC system with redundant tubing and manifolds (not always available from manufacturers)

Scenario 5: Loss of tank pressure / seal	
Likelihood	Low — immersion tanks likely to have robust construction and quality control
Cooling types	Single-phase immersion
	Two-phase immersion
Estimated time before IT affected	Small leak: likely hours
Effects on IT	Isolation and repair likely to require shutdown of all IT on the affected coolant loop
	A small leak in the tank or seal will not immediately interrupt cooling but will allow coolant to escape as liquid or vapor
	Two-phase immersion coolant loss as vapor can be significant, during tank / seal leak or opening tank for maintenance
	If coolant loss is sufficient to stop supply to IT: high-heat components can rapidly increase in temperature, leading to throttling, protective shutdown or chip damage
Mitigation or prevention	Implement software resiliency to shift critical workloads away from affected systems in failure / maintenance

DLC resiliency engineering: a work in progress

As best practices evolve to accommodate more DLC systems, the fragmentation in definitions and team responsibilities may be resolved. But this is a long-term expectation and will likely materialize over several years as standardization efforts come to fruition and product types coalesce around fewer design approaches.

In the short term, fragmentation is initially more likely to intensify in a phase of rapid innovation. The large variety in IT hardware, data center specifications and operator expectations make it unlikely that a single leading DLC type will emerge in a dominant market position.

In the absence of consensus or guidance, data center operators will be required to carefully consider how IT and facilities interface throughout the life cycle when adopting DLC — from design and construction through commissioning and operations — to set expectations and successfully meet them.

Reliability engineering in current DLC systems tends to focus on component quality and durability. To address the requirements of most business applications for higher resiliency standards, DLC system designs will need substantial engineering changes to add redundancy and improve serviceability, such as concurrent maintainability or fault tolerance, extended ride-through times or limiting failure domains.

Some vendors recognize this design gap and have plans to address it. But this is a recent development, one that will take time to reach commercial products, and even more time to generate experience from production installations.

DLC system designs will need substantial engineering changes to add redundancy and improve serviceability, such as concurrent maintainability or fault tolerance, extended ride-through times or limiting failure domains.

Most operators of mission-critical facilities are likely to move slowly with DLC installations at scale, limiting them to select applications where high resiliency is not needed. This is either because service criticality is low or because a form of software resiliency can, to a degree, make up for what is not yet feasible to implement in the DLC system's physical infrastructure. These concerns, among others (such as a lack of standardization, maintenance requirements and component supply diversity), may slow down the displacement of air cooling with DLC in scale deployments for several generations of IT.

Just as virtualization and containers decoupled many software applications from server hardware to make resource allocation flexible, resiliency features for business-critical IT encourage the use of distributed resiliency. With this approach, once application availability can be insulated from underlying IT components, the physical infrastructure does not need to be quite so highly resilient in order to achieve high service standards; concurrent maintainability and fault tolerance are possible with close support from software.

Such a strategy is increasingly common in cloud-native software, albeit not without some problems in practice. These include the difficulty of eliminating single points of failure in software, complications and complexity with workload management and application design, and extra cost burdens for the end user.

A shift toward a software resiliency strategy will be neither quick nor uniform across applications, and data center operators rarely favor designing around the loss of large chunks of IT capacity as an alternative to the cost of building in infrastructure resiliency. For the next several years, many critical applications will remain dependent on resilient physical infrastructure.

Resiliency guidance will change with DLC

For air cooling, the engineering principles and operational practices to achieve high resiliency standards are well understood and tested. The relationship between facility equipment and IT hardware is clear-cut: the facility supplies sufficient volumes of conditioned air to the IT equipment.

This boundary also clearly defines the scope of the data center air-cooling infrastructure and, in turn, how far the standard of resiliency extends. As long as cold air reaches the IT racks, the facility's job is done, and that means the resiliency posture is to focus on guaranteeing the availability of that cold air. A facility event, such as planned maintenance, a loss of utility power or a single failure of cooling equipment, should not have a significant impact on cooling.

It falls on the IT (or, less often, rack systems) to take the cold air to cool IT system electronics before returning it as exhaust — this includes having the necessary redundancy and serviceability of fans to meet availability objectives. Most IT systems have the option to be configured with redundant system fans — similar to redundancy in power supply units — and many also have the option for hot replacement of fans without the server needing to power down.

The same decoupled and well-defined boundary between facility and IT does not, for the most part, carry over to the use of DLC. A new coolant loop adds technical and operational complexity inside the rack, and potentially outside too, when compared with the use of air to cool the IT load.

Infrastructure operators looking to implement DLC in a mission-critical facility environment will need to map out how their technology stack, from software to the physical infrastructure, can meet business requirements for resiliency

Liquid coolant loops and reservoirs are an example — they inherently pose a risk of leakage or a loss of cooling capacity. Maintenance requirements of some mechanical components (or the liquid) in the DLC loop, or servicing the IT equipment, may affect system availability. Crucially, the radius for maintenance or failures can be much larger, affecting multiple IT systems in a rack (or tank) or even multiple racks.

Uptime Intelligence's research indicates that DLC vendors have not prioritized high availability and serviceability with a view to meeting concurrent maintainability and fault-tolerant standards. DLC sales have largely gone to high-performance computing uses that target scientific and technical workloads rather than business and mission-critical applications.

Today, there is no consensus or generally recognized approach to implementing DLC to achieve concurrent maintainability or fault tolerance. Infrastructure operators looking to implement DLC in a mission-critical facility environment will need to map out how their technology stack, from software to the physical infrastructure, can meet business requirements for resiliency.

The following areas, discussed in **Table 4** with suggested guiding questions, should be part of this exploration process. The list is not exhaustive.

Table 4 Key resiliency considerations for DLC implementations

Area to address	Key question(s)
Application criticality	If the DLC system is for a specific application, what is its criticality (expected availability)?
	Alternatively, what is the targeted general resiliency standard (select redundancy, concurrent maintainability, fault tolerance)?
Software resiliency	Can the application stack contribute to the infrastructure resiliency posture through software mechanisms (i.e., high-availability clustering, workload migration, scale-out architecture)?
Facility / IT boundary	Where (at what cooling capacity component) should the boundary of responsibilities lie between the facilities and IT management teams?
	Does IT operations have the necessary resources to meet topological infrastructure resiliency objectives with liquid-cooled IT systems?
DLC equipment assessment	Does the redundancy and serviceability of all the critical components of the DLC system meet the organization's resiliency objectives?
	Can all maintenance and repairs be carried out without resulting in IT downtime?
	If not, is it possible to add redundancy (multiple CDU and facility water connections, added manifolds, tubing and valves) to meet resiliency objectives?
Ride-through time	What is the minimum ride-through time expectation in failure mode on components where there is no redundancy (e.g., loss of facility water, loss of CDU power)?
	How does this affect coolant temperature considerations?
	Can ride-through time be extended by protecting key functions with UPS systems?
Failure domain	Is the worst-case failure domain of the DLC system acceptable?
	If not, what are the technically plausible and economically viable options (e.g., redistributing workloads, smaller DLC system sizes, added redundancy levels, enhanced maintenance plans)?

UPTIME INSTITUTE 2023 uptime INTELLIGENCE

Inevitably, DLC introduces new points of failure, which will create expected and unexpected risks. Some of these risks will be impractical to address through topological redundancy. Key examples are potential leaching, degradation or contamination in the coolant; degradation / damage to tube connections to and in the DLC system; and problems with the structural integrity of cold plates or a tank. Historically, some data center operators have experienced regular leakage issues with DLC, even with highly engineered mainframes that use water cooling, and then phased them out as a result.

Conclusions and outlook

DLC is likely to move toward mainstream adoption, albeit at a measured pace, as server thermal management becomes more challenging. Operators and suppliers will, therefore, need to address the many issues discussed in this report. In the absence of standardized solutions, it will often be necessary to take a pragmatic stance.

The fundamental shifts in cooling system design exert pressure on organizational procedures throughout the data center life cycle, which are well established and may be difficult to change. At the same, suppliers of DLC products will not only respond to the demands of IT equipment, the use of DLC itself will influence new generations of IT as well. These two classes of data center hardware will coevolve as server manufacturers optimize more of their products in response to early growth in DLC deployment.

In the nascent and fragmented DLC product market of 2023, no singular set of guidelines to divide the responsibilities of facilities and IT teams is available. In the foreseeable future, it is likely that the variety of IT equipment and workloads will provide opportunities for expanded installations of both cold plate and immersion systems.

It is not apparent that one clear line of demarcation for maintenance and operations will emerge in this time span. At best, consensus may form around a handful of new paradigms based on organizational preferences — but system design diversity is very likely to frustrate efforts to establish a singular convention that is as widespread as today's virtually unanimous hand-off point in air cooling.

The lack of industry consensus will place the burden on data center operators' technical organizations to consistently define their positions and set expectations across various departments — and with customers in colocation facilities. An organization with plans for significant DLC deployments can best prepare by establishing and codifying rules for communication and collaborative practices across disciplines, from design and engineering to procurement and commissioning, through maintenance and operations.

As (or ideally before) DLC is adopted, data center operators should examine their procedures closely for areas where one or more teams are relying on unwritten conventions to divide responsibilities and replace assumptions with negotiated, express terms for consistency and risk management.

Ultimately, the transition from air to liquid cooling will bring some trade-offs. The resiliency performance of DLC infrastructure will go through a phase of rapid evolution in the coming years as vendors adapt their designs to meet the requirements of mission-critical facilities. But cost and lead-time considerations mean some data center operators — and their customers — will have to make some compromises.

There are two key issues to highlight relating to DLC resiliency:

- First, the resiliency profile of DLC will differ from that of air cooling, which means some compromises that will arguably lower the infrastructure resiliency standard compared with air cooling in the topological sense. This would be largely true even if the system availability projections from fault-tree analysis of components (a probabilistic exercise using estimated mean times between hardware failures and dependency chains between components) might indicate otherwise. This will create an additional barrier to clear communication between facilities and IT, (specialists in each domain tend to speak different languages of availability).
- Second, DLC calls for a greater need to coordinate across facility and IT disciplines to meet service resiliency objectives. Topological resiliency and maintenance of critical components have historically been the sole competence of facilities management. With DLC, IT infrastructure teams will need to better understand these risks and engineering considerations, because they may need to share responsibility for them. They will also need to coordinate with the facility side to be able to select and configure the DLC implementation that is appropriate for the application.

Ultimately, those with the technical capability to shift a part of their resiliency posture from the cooling infrastructure into software-based mechanisms (such as high-availability clustering, share-nothing architectures and stateless scale-out applications) may find themselves in a better position to implement DLC without increasing the risk to IT service availability. But such an approach adds its own set of complexities and further dissipates responsibility for resiliency.

The Uptime Institute report *Direct liquid cooling shows promise, but challenges persist* is available to Uptime Institute members **here**.

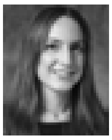

Jacqueline Davis is a Research Analyst at Uptime Institute covering global trends and technologies that underpin critical digital infrastructure. Her background includes environmental monitoring and data interpretation in the environmental compliance and health and safety fields. jdavis@uptimeinstitute.com

Daniel Bizo is Uptime Institute's Research Director. He has been covering the business and technology of enterprise IT and infrastructure in various roles, including more than a decade as an industry analyst and advisor. dbizo@uptimeinstitute.com