INTELLIGENCE UPDATE

IT-agnostic DLC systems: key considerations

Daniel Bizo 14 Sep 2026

Although adoption of direct liquid cooling (DLC) has received an unexpected boost from dense AI computing, practices around how to design, build and operate DLC systems at scale are still maturing. In discussions between Uptime Institute and data center operators, a recurring topic is how best to support all liquid-cooled IT equipment (predominantly using water cold plates) with a single, shared DLC infrastructure.

A single approach for liquid cooling would be beneficial to operators: in its lifetime, a data center facility may see several overlapping IT hardware technology cycles from numerous different vendors. To enable liquid cooling for generic infrastructure use, facilities will need to maintain a single, standard approach within a facility and accommodate hardware from different IT vendors in the same fluid loop.

This report provides a concise, non-exhaustive overview of key considerations for IT-agnostic DLC using water-based coolants. It draws on briefings with data center operators, with additional input from subject-matter experts at vendors Vertiv and Schneider Electric.

Many of the challenges associated with DLC stem from the limited industry standardization and lack of common specifications. This has led to technical fragmentation and commercial lock-in, increasing operational complexity for data center operators.

This is largely because DLC has evolved as a hardware cooling solution that is typically integrated with IT equipment, and was first deployed in early mainframes and supercomputers several decades ago. As the scale of installations now often span multiple racks or even entire data halls, data center operators increasingly prefer an IT-agnostic approach as they assume responsibility for the maintenance and operation of liquid-cooling infrastructure outside the IT rack.

However, ensuring compatibility with various current and future IT hardware generations without major modifications to the DLC system will require operators to weigh several design considerations and the trade-offs between them.

Material compatibility issues and water coolants

The first barrier to mixing hardware in the same DLC loop is resistance from the IT vendors. Technically, differing specifications for the required coolant by different IT vendors should not stand in the way of developing a shared secondary fluid network (SFN). A high-quality water coolant from a reputable coolant vendor, whether it is a propylene-glycol mix (typically PG25) or deionized water, will be compatible technically with the wetted components of commercialized water DLC systems (including cold plates, connectors, hoses, pipes, heat exchangers, pumps) used by any of the major IT hardware vendors. An important caveat is that each coolant requires the correct additive packages to inhibit biological growth, corrosion and any other degradation of the fluid. Maintaining and regularly reviewing a wetted material list with all vendors involved is still essential, and the Open Compute Project (OCP) offers applicable guidelines.

Disciplined commissioning and change management practices for coolant quality

The second concern immediately after material compatibility is maintaining coolant quality in larger SFNs spanning several rows of racks, particularly with diverse hardware from multiple IT vendors. The first step is to protect the hygiene of the DLC fluid network against contaminants when any new equipment is connected to it, including load banks used for commissioning tests. Ideally, all equipment would be filled, flushed and drained by the facility staff while testing for turbidity (particulates) and acidity (pH value) against the baseline of a retained, factory-condition sample of the fluid used in the SFN. In live production environments, taking reference samples from the SFN before and after installing any new equipment will help maintain traceability and accountability.

Filtration standards

The common industry filtration standard for water DLC networks is 25 microns. Full-flow filters are usually integrated with the coolant distribution units (CDUs) to protect the fine internal channels of cold plates from clogging up. For larger, multi-megawatt fluid networks, additional filtration can help maintain coolant quality and protect the CDU filters. Operators may opt for side-stream filtration to improve serviceability.

Dense rack connection points

IT rack manifolds can differ in their positioning, and racks themselves can vary in width. To support as many configurations as possible and allow for changes between them, the row-level SFN headers (also referred to as in-row manifolds) need to have closely spaced connection points.

Flow modulation

In a heterogeneous IT environment, rack power densities, system utilization rates, flow rate requirements and pressure drops can all vary significantly. Adjustable or automated valves help to meet cooling requirements of differing racks and appropriately balance coolant flow between them.

Maintainability

There is no guarantee that every single IT rack will be supplied with redundant supply and return manifolds, which would be a standard requirement to achieve full concurrent maintainability of every part of the SFN without interrupting cooling to any single rack. Alternative options, such as dual-fed manifold or installations with T-connectors, remain less common. To mitigate this, operators can consider dividing the header into more serviceable sections with isolation valves to minimize the number of racks affected during maintenance work. Selecting the right granularity of maintainability requires input from the IT team and potentially other areas of the business, and, in the case of colocation providers, tenant customers.

Rack connection sizing

Rack connection sizes are dictated largely by the expected maximum rack thermal power the coolant will need to be able to handle, as well as erosional fluid velocity. As a rule, for facilities expecting to support high-density AI computing or high-performance technical computing racks, DN65/2.5-inch connections are recommended as a minimum standard. For installations that do not expect to be able to support ultrahigh power densities, DN40/1.5-inch connections will suffice. The capacity of pipework is limited by erosional fluid velocity. With a finely filtered SFN, a maximum of 2 meters per second (approximately 6.6 feet per second) is recommended to stay safely below the erosional velocity and avoid debris accumulation.

Flow rates

OCP's standardization work has yielded useful results for flow-rate assumptions in engineering calculations. This helps both product developers and facility designers to size components. Calculations tend to assume a flow-rate requirement of 1.5 liters per minute for each kilowatt of heat load (L/kW·min) at a 10°C (18°F) rise across the cold plate (ΔT). Operationally, a higher rise is possible as long as coolant temperatures are well below the permitted maximum levels, including return temperature limitations. This allows operators to reduce flow rates and corresponding pumping energy.

Temperatures

There can be a vast difference between theoretical maximum liquid cooling temperatures and real-world operator preferences. There is now renewed attention on chiller-free facility designs using warm water cooling (and hot return coolant), particularly following Nvidia's public endorsement of and hardware support for 45°C water coolant supply.

However, the benefits of a hot water coolant greatly depend on the facility's surrounding climate. Elevated coolant temperatures have various trade-offs — requires stricter thermal stability, limits higher temperature rise and, importantly, many if not most other (non-Nvidia) IT equipment are restricted to lower temperatures. For facility water infrastructure shared with air cooling, CDU supply set points tend to converge around 25°C (77°F). Operators looking at dedicated cooling infrastructure for DLC loads may be considering CDU supply set points around 30-32°C (86-89.6°F), running the facility water loop in the 27-29°C (80.6-84.2°F) range.

These temperature ranges are enough to enable compressor-free operation in a wider range of locations using adiabatic or evaporative trimming (using water only during the hottest hours) and can help to reduce compressor hours and sizing in many other locations with less favorable climates— hot and humid or hot arid. Set points that are even higher are typically features of single-purpose facilities, such as supercomputing and AI computing data centers.

ABOUT THE AUTHOR

Daniel Bizo

15 Sep 2026

Over the past 15 years, Daniel has covered the business and technology of enterprise IT and infrastructure in various roles, including industry analyst and advisor. His research includes sustainability, operations, and energy efficiency within the data center, on topics like emerging battery technologies, thermal operation guidelines, and processor chip technology.

dbizo@uptimeinstitute.com

About Uptime Institute

Uptime Institute is the Global Digital Infrastructure Authority. For over 30 years, the company has established industry-leading benchmarks for data center performance, resilience, sustainability, and efficiency. These benchmarks assure customers that their digital infrastructure will perform consistently with their business needs under any operating conditions. With more than 4,300 awards issued globally and over 1,100 active projects in 120+ countries, Uptime has helped tens of thousands of companies optimize critical IT assets while managing costs, resources, and efficiency.

The organization's Tier Standard is the IT industry's most trusted and adopted global standard for data center design, construction, and operations. Core offerings include Tier Certifications, Management & Operations reviews, AI Infrastructure Advisory, Data Center Risk Assessments, Energy & Sustainability Assessments, Cybersecurity Assessments & Ratings, and IT Advisory Services.

Uptime Education programs have been completed by more than 100,000 data center professionals, establishing Uptime as the Learning & Development partner of choice worldwide. Uptime Institute is headquartered in New York, NY, with offices in London, Sao Paulo, Dubai, Riyadh, Singapore, and Madrid, and full-time Uptime professionals based in over thirty-four countries around the world.

For more information, visit www.uptimeinstitute.com