



VERTIV WHITE PAPER

Deploying Liquid Cooling in the Data Center

A Guide to High-Density Cooling




Table of Contents

Adopting Liquid Cooling at Existing Data Centers **3**

Getting Ready To Deploy Liquid Cooling **5**

Evaluating Liquid Cooling Options To Cool 1 MW IT Loads 8

Implementing RDHx 9

Best Practices for CDU Deployment and Management 9

Drilling Down on New Power Requirements for 1 MW
IT Workloads 11

Restructuring AC Power To Support a Liquid
Cooling Installation 11

Choosing the Right Racks for AI Computing Workloads 12

Leveraging Opportunities to Reduce Costs 13

Designing Mechanical Space 14

Designing Technical Space 15

Determining Piping Size and Placement 16

Selecting Manifolds and Couplings 17

Sourcing and Using Compatible Materials 18

Protecting IT and Liquid Cooling Systems Against Failure 18

Leveraging Project Services To Get Ready for New Capacity Turn-on **20**

Conducting Ongoing Maintenance

Moving Forward With Hybrid Cooling **21**

Appendix A: Creating the Optimal System Design **23**

Appendix B: Glossary **24**



Adopting Liquid Cooling at Existing Data Centers


The future of IT thermal management has arrived, and it's a hybrid of air and liquid cooling technologies. Enterprises are adopting high-performance computing (HPC) for artificial intelligence (AI) and machine learning (ML) model training and inference, causing a fast rise in chip, server, and rack densities, power consumption, and heat levels. Air cooling alone can't abate hot-running equipment effectively. As a result, many data center teams are strategizing how best to make their cooling strategy future-ready in support of evolving business requirements. Nearly one in five data centers (17%) already use liquid cooling, whereas another 61% of operations teams are considering it for their facilities.¹

While some new facilities will be specifically designed for AI workloads and liquid cooling, most deployments will occur in existing facilities. Multi-tenant data center (MTDC) owners and operators are driven to increase their competitiveness. They know that offering HPC capabilities for AI and other workloads will quickly become an industry standard. As a result, many are developing a business case to incrementally add liquid cooling across racks and rooms, evaluating different options, and creating a roadmap to phase in the adoption of new solutions.

IT, facility, and power teams will need to work together closely to deploy liquid cooling at existing facilities, as they will redesign buildings around rack, power, and cooling requirements. Readers of this technical guide are likely seeking insight into how to deploy liquid cooling to support rack densities up to, and in some cases exceeding 50 kilowatts (kW) per rack.

This guide discusses how to take a 1 MW IT load that is currently air cooled and add the incremental liquid cooling infrastructure to create a hybrid system (hereafter called hybrid cooling infrastructure). The guide covers evaluation of cooling, power, and rack requirements, strategies for cost reduction, designing the physical space, fluid network sizing, monitoring requirements, and services. In addition, it provides an appendix with recommended equipment and tips on which equipment should be purchased together. Teams can use this information to convert 1 MW IT loads to hybrid cooling systems and scale with business growth.

Thermal Management Capabilities At Various Rack Densities



Liquid-cooled systems are often used with air-cooling systems to cool racks at higher densities.

¹ Daniel Bizo and Lenny Simon, The coming era of direct liquid cooling: it's when, not if, Uptime Institute, July 2022, page 13, <https://uptimeinstitute.com/>

Top Three Takeaways on Deploying Liquid Cooling

Here's high-level guidance any team can use to navigate the process of developing a hybrid air-liquid cooling infrastructure.



Create a team to oversee the addition of liquid cooling to cool high-density applications

- Assemble a team of subject matter experts who will provide input on hybrid cooling infrastructure design, selection, installation, and maintenance.
- This group will include internal experts (IT, facilities, cooling, and power), consultants, manufacturers, and other vendors.



Get ready to deploy liquid cooling technology

- Gather application and workload requirements to determine future-state needs.
- Develop a target state infrastructure to support the new requirements by using design and budget guidance.



Use project services to deploy and maintain new systems

- Commission systems in line with their design specifications.
- Start up services, train teams on new systems, and ensure an effective handover to operations teams.
- Begin scheduled maintenance.

Getting Ready To Deploy Liquid Cooling

Deploying liquid cooling is a significant initiative that requires careful planning and consideration of the existing facility's footprint, current thermal management strategy, workloads, and budget, among other considerations. Here is a roadmap for getting started.

1 Determine current and future workload requirements

At MTDCs and enterprise data centers, rack densities are growing, spurred by the latest x86 and AI-capable graphics processing unit (GPU) chipsets, surpassing thermal design power (TDP) of 300 watts (W) and 800 W, respectively, and fast approaching 1,000 W or more. These chipsets are used for cloud and enterprise applications including deep learning, natural language processing, AI chat generation, training, and inferencing. With this trend, the newest AI servers are now approaching a TDP of 6 kW to 10 kW per server.

IT and facility teams must decide how much space to allocate to new AI/HPC workloads to support current demand and growth over the next one to two years. Some will convert a few racks at a time, while others will allocate entire rooms for these workloads and supporting the addition of liquid cooling systems.

2 Conduct a site audit

Before developing a business case, teams need to know if retrofitting a facility with liquid cooling systems is technically and economically feasible.

The IT and facility team should work with partners to conduct a thorough site audit. A partner should perform a computational fluid dynamics (CFD) study of existing airflows in the facility. This expert will analyze existing air-cooling equipment to see if it provides enough capacity to be leveraged in the new hybrid cooling infrastructure and if existing piping can be reused. They will also want to perform a flow network modeling (FNM) analysis to select the correct coolant distribution units (CDUs), size piping appropriately, choose manifolds, and evaluate the ability of the liquid cooling system to support server liquid cooling requirements.

Teams will want to execute water and power usage effectiveness (WUE and PUE) analyses to determine how efficiently they are using water and power resources to identify areas for improvement. A total cost of ownership (TCO) study is also recommended to optimize operations by replacing old or inefficient equipment to lower operational costs. Multiple online tools are available to help calculate TCO specifically for liquid cooling applications.


The power team should also analyze infrastructure to see if it can be adapted for use with more power-intensive workloads, such as AI. The larger group — IT, facility, and power — should review physical space to see if raised floors can support the combined weight of new power and hybrid cooling systems and determine access routes for piping. A check of the facility for potential required maintenance of existing infrastructure should also be conducted, as there could be contamination or degradation in the quality of existing pipes or equipment that could lead to inefficiencies or failure. The joint team should review the on-site water supply and determine if it is suitable for use in planned liquid cooling systems. Finally, any safety regulation compliance concerns should be addressed to ensure the new solution will be up to standards and safe for use.

All of this information will inform the business case, so that senior stakeholders can make critical decisions about which approach to use to develop a hybrid cooling infrastructure and provide the investment and resources needed to create it.

3 Model new infrastructure in desired space

With this data and partner support, IT and facility teams can model the desired hybrid cooling infrastructure in the data center and identify obstacles to overcome. These obstacles can include weight restrictions, a lack of on-site water, the need to install new piping, access route concerns, and other issues.

The team can model a greenfield build if space is available at the data center. If the team wishes to replace existing servers with higher-density alternatives, they must optimize technology placement, including busbars, connectivity, and racks. Once all issues have been addressed, it is a good idea to contract with a vendor to construct a digital twin replica of the new design to explore new systems and processes in 3D. The tool can also provide a methodology for modeling different scenarios to implement additional IT loads, all at once or in a phased model. With this information, the team can finalize their system design, making any needed enhancements.



4 Consider budget and site impacts

The audit and modeling exercise provides the IT and facility team with insight into how extensive the liquid cooling deployment will be to develop a business case for executive consideration.

The IT and facility team will also want to consider how on-site construction will disrupt current operations and what impact adding extra heat loads on site will have on current workloads and service-level agreements (SLAs). For example, colocation and managed services firm teams will want to review proposed new heat loads and ensure that they can maintain temperature and humidity for existing customers.

5 Factor in efficiency and sustainability gains

Since liquid cooling removes heat at the source, it can be more efficient than air cooling alone and lowers facilities' PUE metrics. It also uses water or fluid to cool systems and allows teams to recapture and reuse heat, reducing WUE. After auditing systems and benchmarking data, teams can regularly capture metrics, demonstrating progress in reducing PUE and WUE. These gains can reduce indirect or energy-regulated emissions (Scope 2) for enterprises. As a result, liquid cooling can be an essential part of enterprises' sustainability programs.

6 Designing the new solution

With this information, the IT and facility team can work with the design consultant and partners to design a new solution customized for site requirements. The joint team can then use that information to create a bill of materials and services (BOMS), request quotes, and select the manufacturers to build and integrate the liquid cooling system.

Planning Liquid Cooling Deployments: Equipment and Timeframes

What type of equipment does it take to support a liquid cooling solution at an existing site? The product types below can vary based on the cooling approach, site requirements, and other factors, but they provide a starting point for making the first steps.

Category	Equipment Type
Liquid Cooling	Rack manifold, CDU, distribution manifold, RDHx
Heat Rejection	Freecooling or indoor split chiller; drycooler; DX condensing unit
Power Distribution	High-amperage busway, rack power distribution units

Want help navigating options, designing a custom system for your site, and creating a budget?

Contact Vertiv for more information.

Allow up to a year to deploy new systems

At least 2 months

To select vendors, suppliers, and contractors; perform CFD, PUE, WUE, and TCO analyses; conduct site planning; obtain a design quote; and create a BOMS.

At least 3 months


Manufacture and transport of liquid-cooled infrastructure, including pipework, rack manifolds, CDUs, heat rejection, power components where required, and racks.

2 to 3 months

Conduct site planning and deliver and integrate the new liquid cooling system.

2 to 3 months

Integrate IT equipment with the new cooling solution and check for any discrepancies. Conduct regular maintenance as needed.



Utilizing liquid-to-air CDUs could speed up deployment times.

Evaluating Liquid Cooling Options To Cool 1 MW IT Loads

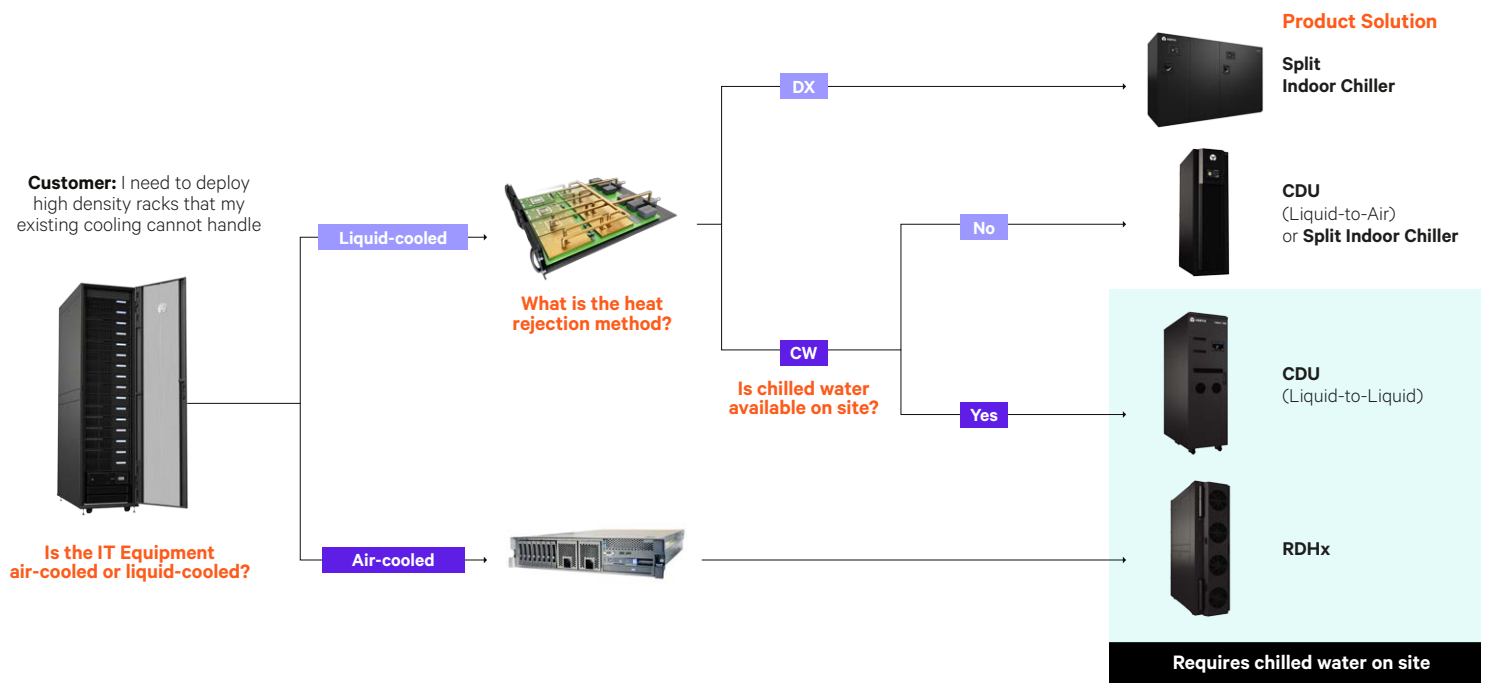
No single liquid cooling solution exists for 1 MW IT loads, as company requirements, site conditions, and budgets vary considerably. Some options teams can consider include the following:

- Leveraging existing chilled water:** If teams have access to chilled water on site, they can deploy a CDU that provides a liquid-to-liquid heat exchanger. To do this, the chilled water system must be a 24x7 design versus a traditional comfort cooling system which is often shut off after hours, on weekends, and during off-seasons. Available as in-rack, in-row, and perimeter designs, these CDUs isolate IT equipment from the facility's chilled water loop by providing a separate liquid loop to and from the racks. This solution enables teams to determine which fluid and flow rate to supply the racks, such as treated water from a water-glycol mix. Water quality may also be a concern when deploying CDUs and utilizing existing chilled water systems. Ensuring availability of filtration, or the capability to deploy filtration if needed, should be considered when leveraging the on-site chilled water.

- Leveraging existing DX condensers:** If teams have already deployed a direct expansion system, they can consider utilizing split indoor chillers as a solution. This system can reuse existing rooftop condensing units or additional units can be deployed to support incremental loads. Split systems offer the flexibility to be easily deployed as demand increases incrementally. When considering a split system, technologies such as pumped refrigerant economization should be leveraged to drive overall system efficiency. If the existing rooftop condensers are not enough in terms of quantity or capacity, or if there is not sufficient roof or ground space available for additional units, then a split indoor chiller may not be applicable.
- Overcoming a lack of access to reject heat from the liquid:** Teams can consider deploying a standalone CDU that provides liquid-to-air heat exchange, cooling IT equipment and blowing hot air into the data center, where computer room air conditioning (CRAC) units capture and remove it.

Determining Which Liquid Cooling Solution To Add To Air Cooled Facilities

Liquid Cooling Solution Tree



Use this decision tree to select the right liquid cooling solution.

Implementing RDHx

Many organizations use RDHx as a first step on their path to incrementally adding liquid cooling to develop a hybrid cooling system, as they can be used with existing air-cooling technology and don't require significant structural changes to white space. Rear doors do extend the depth of the cabinet and require a minimum distance between the door and other objects to ensure unobstructed airflow.

RDHx use air to cool racks but capture heat via a liquid-to-air heat exchanger mounted to the racks. The air leaving the heat exchanger is then expelled to the room at ambient temperature. Rear doors can be configured as either passive or active units. Passive heat exchangers use server fans to expel hot air through a liquid-filled coil, which absorbs the heat before the air is returned to the data center. Since these do not have fans or other moving parts requiring power, passive units have no annual usage cost to the doors themselves. These passive doors work well for loads from 5-25 kW. With no moving parts, these passive units will also reduce the amount of noise in the facility.

Active units assist the servers with fans attached to the rear of the door that can help pull hot air through the coil and reject room-neutral air to the IT space. The fans can modulate between 0-100% based on either the differential pressure across the coil or the measured temperature difference. Active rear doors are capable of cooling 50-70 kW based on application parameters.

Fluid flow control for passive or active doors can be adjusted based on the return water temperature using a modulating valve to make sure the door is operating efficiently. Additional monitoring and control capabilities are available with these units, providing remote access, alarms, and system status updates.

When properly sized, RDHx can eliminate the need for containment strategies, such as hot and cold aisles. However, it is still recommended that rack doors face each other in a hot aisle/cold aisle pattern.

Best Practices for CDU Deployment and Management

CDUs create an isolated secondary loop separate from the chilled water supply, enabling teams to contain and manage fluids used in liquid cooling systems and precisely control their pressure, temperature, flow rate, and even filtration quality.

CDUs are equipped with supply-side filters that are typically 50 microns in size. The filters keep the fluid supply free from contaminants and debris, protecting the integrity of the server cold plate and maintaining its performance. The cold plate fin pitch is the space between the raised fins on the surface of the cold plate. When teams deploy cold plates with a fin pitch of 100 microns, it is advised to use filters no greater than 50 microns to reduce the possibility of debris gathering within the fin pitch gap on the cold plate and negatively affecting system performance or life-span.

To determine CDU requirements, the application engineer will evaluate the number of racks to support, the heat-to-liquid ratio for liquid-cooled IT equipment in each rack, the flow rate required for each rack, and the facility layout. For example, CDUs that support systems with high required flow rates will run out of pumping capacity before the heat exceeds its thermal capacity. As a result, these applications require greater redundancy. After establishing requirements, the engineer will:

- **Connect CDUs to key equipment:** Liquid-to-liquid CDUs must be connected to the facility's chilled water supply or other heat rejection source to provide coolant needed to remove heat from liquids. Additionally, the CDU should be connected via inbound and outbound pipes to the liquid that needs cooling. The CDU contains a heat exchanger which transfers heat into the chilled water supply, completing the liquid-to-liquid cooling process.

Liquid-to-air CDUs do not require a chilled water supply to provide liquid cooling to the rack, but rather provide an independent secondary fluid loop to the rack and reject heat to the data center. The secondary fluid loop is an isolated system, delivering chilled fluid to liquid cooled servers in an otherwise air cooled environment. By rejecting heat directly into the data center, the liquid-to-air CDUs can rely on existing air cooling systems to capture and reject the heat outdoors.

- **Determine CDU placement:** CDUs can be placed in multiple locations. With in-rack or in-row placement, they feed individual racks or rows, respectively. They can also be placed at the end of a row, feeding the racks on that row.

Finally, they can be grouped at the end of a data center, feeding a common fluid circuit that supplies all racks in the rooms. Teams will make this determination based on the level of redundancy they want to achieve, how efficient they want the liquid cooling to be, and the capital expenditure (CapEx) funds they have for new technology.

- **Decide on the level of redundancy desired:** To achieve N+1 redundancy at the row level requires two CDUs: one to feed the racks and the second to back it up. This takes up valuable floor space and increases costs. However, MTDCs may choose this layout to provide redundancy for individual customer technology deployments.

Enterprise IT and facility teams will likely group all CDUs centrally to provide adequate fluid flow, achieve redundancy with fewer units, and scale more quickly.

- **Determine liquid cooling requirements:** The amount required to support 1 MW of IT load varies based on workload temperatures and flow rates.

The IT and facility teams determine rack power and hydraulic requirements. The table below shows an example of a direct-to-chip liquid cooling system that cools 20 50 kW racks. The heat-to-liquid specification is for the cold

plate liquid cooling system, where cold plates are placed on IT equipment components that generate the most heat. Other components rely on air systems for cooling.

The liquid flow rate is approximately 1.5 liters per minute/kW, which depends on cold plate heat transfer effectiveness and the heat generated by IT components cooled with liquid. The facility specifying engineer will consult with the IT equipment vendor to determine upper and lower flow rate requirements to set the desired flow rate. The facility team should review flow rates regularly and increase allowable ranges when IT equipment is refreshed.

The IT and facility teams and contractors should also inspect the liquid loop to ensure that the piping can support desired flow rate requirements. This group will also want to review the pipework, secondary inlet, and return header dimensions, ensuring they match connections to the CDU. The system will use pipework reducers to connect hydraulic loops to the CDU.

Representative Hybrid Cooling Deployment for a 1 MW IT Load

Quantity of liquid-cooled racks	20		
Total power per rack	50 kW	Total cluster power	1 MW
Heat to liquid	37.5 kW	Total cluster heat to liquid	750 kW
Heat to air	12.5 kW	Total cluster heat to air	250 kW
Liquid flow rate per rack	56.25 lpm	Total cluster liquid flow rate	1125 lpm
Airflow rate per rack	1875 CFM	Total cluster air flow rate	37500 CFM

For illustrative purposes only. Actual site conditions may vary.

Drilling Down on New Power Requirements for 1 MW IT Workloads

How do power requirements change with liquid cooling?

The IT power load may remain the same as data center teams enable new AI and other 1 MW IT workloads. However, rack densities will increase, consolidating power consumption into fewer racks.

In other scenarios, data center teams may want to increase overall IT power to expand computing operations.

Either way, the IT and facility team needs to ensure that power distribution provides the higher amperage that AI computing hardware requires. Teams can accomplish this goal by:

- **Installing a new busway:** A high-amperage busway can provide the overhead feed that new computing hardware requires. Tap-off boxes that draw off power should be appropriately sized to protect current flow.
- **Considering redundancy:** Teams may want to deploy A+B power feeds, each with automatic transfer switches, uninterruptible power supply (UPS) units, and power distribution units (PDUs), to provide redundancy to duplicate critical components and ensure power continuity for workloads.

The greater the redundancy, the greater the protection against downtime. However, additional systems may increase the overall budget. Multiple redundant models can be deployed, including N, N+1, N+2, and 2N.

- **Determining how to route electricity:** Teams can commission an electrician to hardwire electrical connections to the rack or combine the higher-amperage busway with higher-amperage rack PDUs. If teams choose the rack PDUs, they should decide whether to have an electrician hardwire them or use a traditional rack PDU with a plug.

While the traditional rack PDU is simpler to install, teams may be able to decrease the number of PDUs within each cabinet by hardwiring them. Fewer rack PDUs mean cabinets have more space for liquid cooling equipment, such as direct-to-chip liquid cooling pipes, manifolds, CDUs and RDHx.

Restructuring AC Power To Support a Liquid Cooling Installation

Restructuring AC power to support the addition of liquid cooling requires an evaluation of the full power train from server loads to utility feed. By so doing, the site electrical team can ensure sufficient critical power continuity for new HPC applications.

Currently, standard AC power systems support most high-density liquid cooling applications. However, some liquid cooling sites will require reconfiguration, redeployment, or full replacement of the power infrastructure.

To get ready for liquid cooling, IT and facility teams should work with a design consultant to design an integrated thermal-power solution, paying close attention to the physical layout of the space and potential mechanical interferences. This process allows teams to potentially upgrade critical power to support continued business growth or improve redundancy to ensure power continuity to liquid-cooled racks and rooms.

The AC power restructuring process includes:

- **Assessing the current state:** The electrical team will define the current state of power supporting IT workloads and infrastructure versus the target power distribution and infrastructure required to support HPC workloads.

The team will:

- Create a detailed power inventory, associated electrical one-lines, and target states to frame restructuring requirements.
- Capture available power backup, conditioning, and distribution capacity (such as rating and connectivity) to maximize the utilization of power restructuring and define any required supplemental support.
- Document power cabling sizing and layout, especially from the PDU, remote power panel (RPP), and busway to equipment where load concentration change is most acute.
- Evaluate overall rack and discrete internal rack loads for rack PDU deployment, considering redundancy strategies, phase balancing, and connection type and count.

- **Plan and implement restructuring:** The electrical team will determine the requirements and process to convert the power infrastructure to accommodate new liquid cooling systems and their buildout, evaluating potential improvements. The team should consider physical constraints associated with concentrated rack power loads (such as connections and wiring path) and how best to integrate seamlessly with the liquid cooling infrastructure.

- **Manage new systems:** The electrical team should consider working with power partners to help define and provide ongoing critical power monitoring, analysis, and service management to ensure critical power continuity. The facilities team should establish and monitor power performance and alarm parameters, as well as confirm preparedness to manage new power systems on an ongoing basis.

Choosing the Right Racks for AI Computing Workloads

To deploy new AI workloads, teams must change their approach to selecting racks. Racks will support dense computing environments up to 50 kW — and ultimately beyond. Thus, when selecting racks, bigger is better. Racks between 42U and 52U provide room for servers populated with accelerators, networking infrastructure, power distribution equipment, and cables. When selecting taller racks, ceiling height needs to be accounted for as some sites only have a 2 meter clearance height in the data center or through doorways. Teams should plan on having each rack support up to 3,000 pounds (1.36 kilograms) of static weight to handle increased gear density. As a result, it is highly unlikely teams will be able to reuse existing racks to support new AI/HPC computing equipment. They should also invest in cable management to reduce clutter and avoid the risk of unplugging essential components. Other guidance includes:

1 Select wide racks

The most common rack in an air-cooled data center is 600 millimeters (24 inches) wide, which aligns with raised floor tiles. When deploying liquid cooling, teams can benefit by selecting wider racks, with 800 mm (31.5 inches) being the most common size. This size provides an extra 100 mm of space (4 inches) on each side of the centrally located server mounting rails.

This extra space provides room for the coolant distribution piping and high-power cabling at the rear of a liquid-cooled rack. It also offers additional free, unobstructed space directly behind the IT gear that helps preserve proper airflow through the equipment and out the rear of the rack. Crowding this area by selecting a narrower rack can impede airflow, increase temperatures, reduce cooling efficiency, and degrade IT performance as temperatures rise.

2 Choose deeper racks

In addition to demanding greater width, liquid cooling systems benefit from deeper racks. New piping occupies space that previously wasn't required for air-cooled racks. In addition, high-power IT gear requires thicker power cables to provide increased power, while larger PDUs are needed to handle higher ampacity. While 1000 or 1100 mm (40 or 43.5 inches) deep racks have worked for some air-cooling cases, racks for liquid cooling should be 1200 mm (48 inches) deep. This is especially true when using a RDHx, as the extra depth allows for better airflow (static regain). Similarly, direct-to-chip cooling systems require space to mount rack distribution manifolds, cabling, and rack PDUs.

3 Consider where to place piping

Piping that's run under the floor to rack rows must also feed each rack individually. These pipes can be large in diameter and may change the underfloor airflow pattern that data center designers originally intended for power or data cabling. With a greenfield build, designers may use the raised floor for piping, running power and data cabling overhead. With a brownfield retrofit, designers may need to move existing racks and air cooling, which could interfere with pipe runs. This can be costly and disruptive to operational performance.

One alternative is to mimic greenfield designs, decommissioning underfloor space and using it solely for pipework, moving power and data cables overhead. Separating cabling from piping has multiple advantages, which includes alleviating the risk of a fluid leak affecting cabling and reducing "working interference" where workers accidentally push or nudge cables out of the way while working on a system, even if intentional and temporary. Working interference can negatively impact other systems.

Teams at data centers with slab floors will likely run power and piping overhead and need to consider issues such as system weight and available space. See “Designing Technical Space” for more information.

4 Plan for inter-row clearance changes

Inter-row clearance may be affected if liquid cooling equipment, such as a RDHx, adds depth to a rack. For example, consider two racks placed rear to rear, which commonly occurs in hot aisles. Each rack gains up to 330 mm (13 inches) in depth when adding a RDHx, while the aisle width is decreased by 400 mm (16 inches). This may result in an aisle that is too narrow for effective cooling and doesn’t enable functional work between racks.

Leveraging Opportunities to Reduce Costs

Regardless of the technology type or required updates needed to implement liquid cooling, deploying a hybrid solution is a major investment toward the future. While this can be a point of concern for some, data center IT and facility teams have many opportunities to reduce upfront CapEx costs to implement liquid cooling infrastructure. There are many ways to go about this, including:

- **Perform a thorough site assessment:** As mentioned previously, a site audit should be one of the first activities data center teams perform when looking to deploy a high density system. Review the existing infrastructure conditions to identify areas that may need modification, renovation, or new equipment to support this higher density IT workload. Work with internal and external experts to ensure the full scope is understood.
- **Reuse existing infrastructure:** Much of the installed infrastructure, such as cabinets, PDUs, UPS, and cooling, has the potential to be reused for high density applications. Rather than unnecessarily investing in equipment that renders existing product obsolete, look to use as much as possible.
- **Utilize existing cooling infrastructure:** A major benefit to hybrid cooling solutions is the ability to take advantage of existing cooling infrastructure. By deploying a split indoor chiller unit, customers can make use of existing outdoor condensers for heat rejection.

5 Prepare for different space use

High-density racks take up less white space, but the proportion of mechanical to white space in high-density, liquid-cooled data centers varies. For example, replacing 100 10 kW racks with 20 50 kW racks decreases the white space footprint. However, teams must deploy new equipment needed for liquid cooling, such as CDUs, indoor chiller units, piping, storage tanks, treatment devices, and leak detection and monitoring systems. Conversely, teams may be able to decommission some air-cooling equipment, potentially freeing up space. Regardless of an increase or decrease in available space, teams should also take into consideration the benefits that deploying an aisle containment system can bring to their existing white space.

If liquid-to-air CDUs are deployed for direct-to-chip applications, existing room cooling units are still required parts of the hybrid infrastructure.

- **Repurpose the power grid:** While some AI/HPC workloads increase power requirements, others will only consolidate power into denser rack deployments. In these instances, existing UPS units, busways, and PDUs can be rearranged to meet the new demands or be used to supplement new power requirements, limiting the initial investment.
- **Leverage existing IT racks:** Not all AI/HPC deployments will require additional U-space within the racks to support the increased workload. Whenever possible, teams should maximize usage of the existing racks to limit the up-front investment. Liquid-cooled servers take up less space compared to air-cooled heat sinks.
- **Carefully select fluids and coolants:** With the plethora of options available on the market at various price points, data center teams can be very selective about which fluid or coolant will be the best fit from a material compatibility, availability, and economic standpoint. Selecting one that is more cost-effective than others will help reduce upfront expenses.


- **Focus on energy-efficient solutions:** Energy efficiency is a top priority for investors and operators. Selecting equipment that has been designed to maximize energy efficiency will help reduce ongoing operational costs. There are many liquid cooling options that provide significant efficiency gains, which leads to overall long-term savings.
- **Implement the changes incrementally:** The new high density deployment and supporting hybrid cooling infrastructure doesn't need to be installed in one large system upgrade. Start with critical areas or equipment, and gradually expand to the entire data center to spread out the capital costs over time. Work with technology partners to decide which products need to be deployed first to support future expansion. Consider using modular systems that can be added incrementally to facilitate the addition of liquid in the data center.

Designing Mechanical Space

Where should teams place cooling systems? In some instances, equipment may be located in the white space. However, most MTDC or enterprise facility operations teams prefer to locate liquid cooling equipment in grey space, which may be shared with electrical equipment.

A typical scenario is to place equipment in a hallway or gallery surrounding the perimeter of white space. This gallery may have a raised floor that's level with the white space raised floor or have a sunken floor.

Another less common scenario is to run a deeper concrete trench around the gallery with the raised-floor level above the trench — level with the white space raised floor. In this scenario, large pipes are typically located in the trench floor or its walls. The goal is to create adequate space for ongoing service and maintenance.



Air-cooling systems that serve IT systems in the white space may penetrate the gallery via ducting or wall cut-outs. Depending on their design, these areas, which enable free air pass-through, may require fire dampers.

When designing mechanical space, teams and contractors should ensure that it provides adequate room to install, remove, and replace any items located therein without disrupting other equipment. For this reason, aisle space may need to exceed required fire, electrical, and accessibility code requirements.

Designing Technical Space

Next, the IT, facility, and power team must work with the design consultant to design technical space, considering where racks, air cooling distribution, fluid piping, power distribution systems, and monitoring equipment will be placed. A key factor guiding decision-making will be whether the data center has a raised or non-raised floor design.

For a non-raised floor, the team should consider a technical design that places power and pipework conduits above the racks with threaded rods hung from ceiling joists, providing ceiling support.

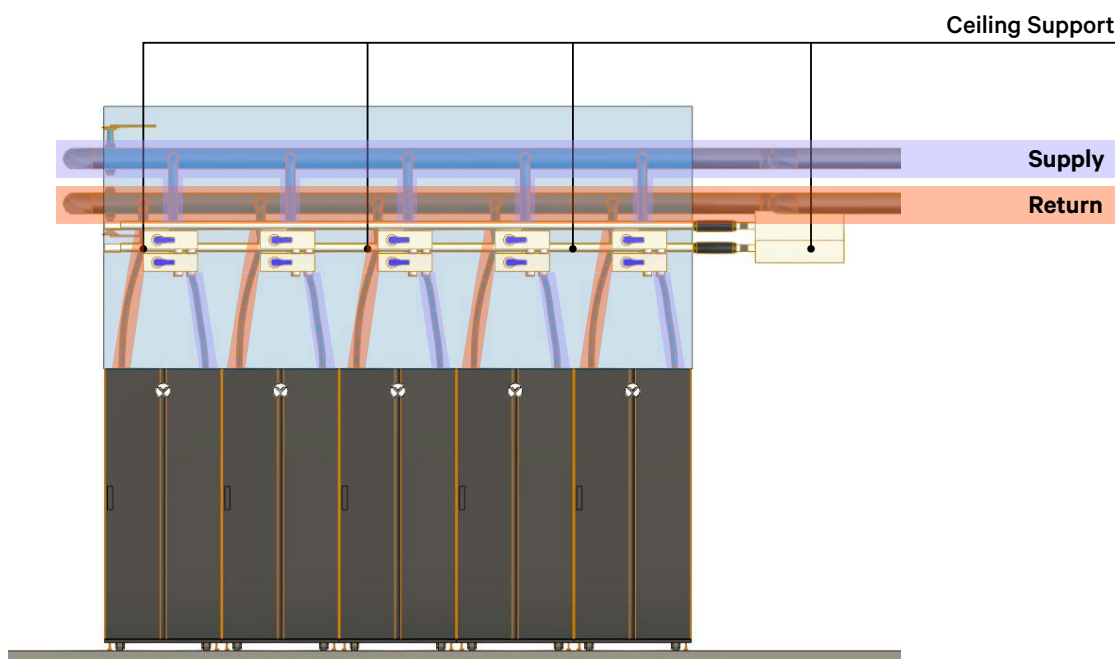
In addition to space considerations, an initial calculation is recommended to determine if the existing floor can support the load of the new equipment and if any bolt down brackets will be required.

Since pipework is run overhead, IT and facility teams should ensure valves and other items are easy to access for maintenance.

The design should consider the weight of supporting cable trays and pipework to ensure overhead joists and ceiling structures can support them.


The team should also verify sufficient space to accommodate pipework and equipment servicing. To continue the 1 MW IT load example, a team would need 20 50 kW rack positions, each with a 1.5 lpm (0.4 gallons per minute/kW) of low rate.

How Threaded Rods Support Overhead Infrastructure




This illustration depicts an overhead liquid supply and return manifold and power bus.

Running Liquid Cooling Equipment Overhead in Raised-Floor Facilities



This illustration shows how running a rear-door heat exchanger adds depth to racks.




This illustration depicts the clearance required to accommodate underground piping in a facility with a raised-floor design. Dimensions will vary from site to site.

Determining Piping Size and Placement

Pipe flows, or capacity, are limited by velocity and pressure. High-velocity fluid flows can cause erosion at pipe turns and create noise. Pressure loss is driven by flow rate, the length of the run, and the use of devices such as valves and quick-connects.

Allocating Extra Space Underfloor for Piping Requirements



Underfloor designs for piping should leave adequate space for piping, insulation, and pipe crossover. Dimension will vary from site to site.


Pipes should also provide a loop to isolate valves from racks, reducing the possibility of leaks.

When piping is deployed under raised floors, it must fit in the available space. The design consultant should consider piping size and the insulation wrapped around it. If the team adds 25 mm (or one inch) of insulation to a pipe, the overall diameter of the wrapped pipe is increased by 50 mm (or two inches). The design consultant should also consider that pipes may cross over each other and account for extra height requirements when finalizing the underfloor height and piping system.

Selecting Manifolds and Couplings

Like piping, an in-rack manifold is a relatively simple set-and-forget product. However, IT and facility teams still have some unique design requirements to consider, as manifolds are placed in racks and are the last touchpoint of the secondary fluid network supplying the server. When they select manifolds, teams should consider material compatibility, flow rate, pressure, flow distribution uniformity, and pressure drop.

Ensuring material compatibility is especially important as teams want to prevent galvanic corrosion, which could cause catastrophic failure of the manifold or server chassis. Galvanic corrosion occurs when two different metals are in contact with each other in the presence of an electrolyte. The metal that is less resistant to corrosion will oxidize and flake.



In addition, components used to seal each junction, hoses used for fluids, and couplings must be compatible with the selected wetted material. If they're incompatible, seals could begin degrading, with small pieces breaking off, entering the flow stream, and damaging the downstream heat sinks.

Direct-to-chip cooling uses coolant to remove heat from cold plates and direct it to the CDU for removal from the data center. For this process to work, there needs to be an adequate flow rate. As a result, teams must ensure that the manifolds and couplings are sized appropriately and have enough pressure to supply an adequate flow rate through the cooling system.

Each system component must also withstand pressure drop and withstand high pressure, including manifolds and couplings. The manifold system must be able to isolate racks and individual servers to protect IT equipment from leaks and enable both planned and unplanned maintenance.

To prevent fluid from escaping the liquid cooling system when couplings are connected or disconnected, server and manifold couplings should have a non-spill quick disconnection with an internal valve and seal. The fluid should not flow until a complete connection has been made, and the valve should immediately reseal itself after initiating a disconnection.

In-rack manifolds that are being plumbed from below also require air bleeders, which control how much the valve opens to allow air release, including fluid pressure. In addition, teams may need to deploy circuit setters to relieve pressure from the secondary manifold as it connects to the in-rack manifold.

Finally, the secondary fluid network piping system connection to the in-rack manifold must also have an intermediary isolation valve. Typically, manufacturers use an isolation valve (such as a ball valve) placed between the secondary fluid network piping and the in-rack manifold's inlet to enable isolation of the secondary system. While it's possible to add controls, they significantly increase solution cost.

Sourcing and Using Compatible Materials

Key factors that drive the selection of fluids for liquid cooling include the presence or absence of chilled water on site and the need to ensure compatibility with all systems and components in the secondary liquid circuit.

Direct-to-chip manufacturers can guide what fluids are best for use with cold plate loops. Vendors are increasingly using propylene glycol to enhance and extend fluid quality throughout its lifecycle. However, propylene glycol provides lower thermal performance when compared to pure water. It is thus often treated with biocides and inhibitors to maintain thermal performance, increase reliability, and strengthen corrosion protections. Best practices include:

- **Ensuring material compatibility:** The selected fluid should be compatible with all systems that encounter it, including the CDU pump, heat exchanger, valves, couplings, sensor fittings, and pipework. Pipework includes hoses, couplings, seals, gaskets, and brazing flux. The manifold in the liquid-cooled rack includes couplings, hoses, and connectors for the cold plate loop. The cold-plate array includes couplings, connectors, hoses, and the cold plate itself. Since most cold plates are manufactured from copper, wet materials should not include aluminum, which would cause galvanic corrosion of the cold plate.

For more information on material compatibility requirements, please reference Liquid Cooling Guidelines for Datacom Equipment Centers, Section 6.1.5 – Wetted Material Requirements from the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE).

Teams should also consider the impact of fluid velocity as well. Degrading wet surfaces due to galvanic corrosion or erosion will contaminate the flow stream and cause catastrophic failure downstream at delicate and critical components such as the cold plate.

- **Conducting regular testing:** Teams should contract with a partner for regular chemistry testing to ensure fluids maintain their desired properties, typically recommended to be done twice a year. The same fluid composition should be used throughout the system lifecycle to avoid unwanted chemical interactions caused by a different material.

Fluid testing provides two main benefits: maintaining fluid health by managing contamination and the impact of contamination on fluid life and ensuring that the fluid meets the desired lifecycle targets without requiring flushing or additives.

Protecting IT and Liquid Cooling Systems Against Failure

IT and facility teams can use multiple strategies to reduce the risks of damage to IT from fluid leaks and proactively mitigate leaks when they occur. These strategies include:

- **Isolating IT infrastructure:** By keeping fluid piping separate from IT and power equipment or cabling, teams can reduce the risks of leaks. Each rack should also be able to be isolated from the system through the use of valves, and each server should be able to be isolated from the in-rack manifold to enable routine and emergency maintenance. By designing systems that allow users to shut off fluid flow to individual racks and individual servers, teams can continue to perform maintenance on specific sections of the data center without causing disruptions to critical operations. Tubing from the manifold to the server should have interconnections with dripless quick disconnects to avoid any damage caused when servicing the servers.

- **System visibility through monitoring software:** Monitoring environmental conditions around liquid cooled systems is pivotal to ensuring protection of the IT equipment. Liquid cooling is inherently different than air cooling when it comes to rapid system response time when failure scenarios occur due to the higher heat densities that exist. To ensure real-time visibility to system performance and IT equipment safety, additional data should be collected and acted upon quickly. Integrating intelligent flow monitoring and advanced leak detection into the cooling system monitoring solution will enable more proactive visibility for identifying harmful conditions. These scenarios could include system fluid leaks, humidity levels allowing condensation, and elevated fluid temperatures being supplied directly to IT gear.

With liquid cooling, another important concern is one of coolant fluid quality. Identifying fluid parameters that cause corrosion, blockage or limitation of fluid heat transfer capacity should all be considered in selecting a sufficient monitoring software package.

Liquid-cooled and air-cooled systems should both provide important data points upstream to higher level monitoring solutions through standard industrial protocols for easier integration into common existing interfaces to provide a wholistic overview of system status. This enables MTDC and enterprise facility operations teams to easily report on this information to meet customer and regulatory requirements.

- **Leveraging protective system controls:** For liquid-cooled systems, robust controls should provide not only exemplary cooling unit component logic, but multi-device awareness and system-level reaction capabilities. As heat densities increase, reaction time of the system in responding to events of a larger data scope (outside of the single unit controller) become more important. Utilizing advanced embedded controllers enable flexible and programmatic interfaces to collect non-traditional data sets for better IT gear protection and full system energy efficiency. Cross-device awareness through these interfaces enables system-wide operation through setpoint sharing and parallel protection of system flow and temperatures. Automated failover of redundant components and units ensure this protection while specific liquid cooling control parameters allow easy implementation of energy efficiency initiatives.
- **Using extensive leak detection systems:** Leak detection systems should be placed across system components reporting into or off of a CDU and at critical locations across the piping system.

They provide proactive alarms when fluid pressure drops at node, rack, and data center levels, indicating a leak. These systems may also identify where a leak occurs, enabling teams to do controlled shutdowns of the CDU, the servers, and its cooling systems, streamlining equipment troubleshooting and repair processes. Deploying systems that can manage temperature below the dew point is vital to limit condensation formation and avoid potential leaks.

- **Deploying systems with integrated cybersecurity measures:** One of the top priorities of data center managers is cybersecurity, according to Uptime Institute reports. Based on survey responses, 43% of respondents have initiated or increased cybersecurity initiatives in response to regulation changes, and 64% of respondents cited data security concerns as a main reason for not hosting mission-critical workloads in cloud data centers.² This extends past the IT equipment and into the supporting infrastructure. CDUs, RDHx, supporting power equipment and more all require protection from cyber threats. Selecting equipment from vendors that integrate protection into their units removes strain from the system managers and speeds deployment times.

One caution: If the IT and facility team adopts non- or low-conductive fluids in the future, they will want to review and enhance existing systems to ensure they can continue detecting any changes in fluid conditions.

² Douglas Donnellan, et al., Uptime Institute Global Data Center Survey 2023, Uptime Institute, July 2023, pages 21-22, <https://uptimeinstitute.com/resources/research-and-reports/uptime-institute-global-data-center-survey-results-2023>

Leveraging Project Services To Get Ready for New Capacity Turn-on

The IT and facility team can leverage partner project services to prepare for on-site liquid cooling system installation, commissioning, and startup. As a major infrastructure project, this requires careful planning and coordination. Project phases and services include:

- **Navigating application design:** The IT and facility team will leverage available services and work with key partners to determine the scope of work that needs to be done on site by performing a thermal and electrical evaluation. This process includes providing written instructions for electrical contractors on where to install or upgrade busbars and PDU circuit breakers in remote power panels and distribution boards. For cooling, it might involve taking down walls, adding or redirecting piping, directing water through cold water loops, and deploying manifolds.
- **Installing the new system:** The external integration team will then implement the new hybrid cooling infrastructure, leveraging all guidance documents and preparation work. The IT and facility team should ensure that any external firms working with the liquid cooling project deploy only certified engineers to the site.

One caution is to keep the space and equipment clean to be ready for commissioning and startup. The integration team should install end caps in all open pipework, manifolds, and connectors.
- **Conducting off-site testing:** The system integrator will test all systems off site to ensure they work properly before installing them at the data center. The IT and facility team should pre-inspect equipment on arrival to ensure it functions properly.
- **Starting up systems:** To ensure that the new hybrid cooling infrastructure is ready for use, contractors will conduct a series of tests, including L3 (startup and leak detection testing), L4 (site acceptance testing, which may require specialist load bank testing), and L5 (integration testing). The integration testing will determine that cooling, power, room, and monitoring systems are all working effectively together to manage target heat loads that represent the daily operating conditions systems will be required to support. Contractors will also provide site data and closeout reports.
- **Providing service project management:** The IT and facility team should develop a document detailing who is responsible, accountable, consulted, and informed (RACI). This expanded group will include the IT, facility, health, and safety teams, the design consultant, equipment manufacturers, and mechanical, electrical, and IT consultants.

After creating the RACI document, the IT and facility team should work with on-site experts to create a health and safety plan that discusses tasks, methods, risks, and mitigation strategies. Then, the IT and facility team can create a site plan with the design consultant, including all hydraulic, mechanical, and electrical designs for the target space.

Finally, the IT and facility team will want to create operational guides that teams and partners will use for servicing, troubleshooting, and emergency maintenance. By creating defined processes, implementing controls, and maintaining strict governance, the group responsible for deploying the air-liquid hybrid cooling system can mitigate risks and ensure a safe and effective implementation and handover process.
- **Handing over the new system:** Equipment manufacturers can train the IT and facility team on new systems to understand how to operate and maintain systems effectively and what to do when they detect performance anomalies.

Conducting Ongoing Maintenance

Liquid cooling systems need to be constantly monitored and controlled. Teams will want to ensure that all fluids used in systems are high-quality and consistent across applications and that any leaks are proactively identified and mitigated. Equipment manufacturers can provide ongoing services to protect costly liquid cooling systems and ensure effective operations. These services include:

- **Conducting remote monitoring:** Equipment manufacturers can provide a remote web portal so IT and facility teams can manage liquid cooling systems worldwide. On-site teams will respond to local alarms from their building management system, while the partner will provide condition-based monitoring. By identifying anomalies, the partner can recommend and conduct proactive and preventive maintenance to prevent leaks and system breakdowns.
- **Performing preventive maintenance:** During routine on-site visits, a partner will perform mechanical, electric, and hydraulic visual and control checks, maintain pumps, replace filters, and check for any signs of corrosion. Fluid management can occur during these visits. Partners will manage systems to achieve predetermined SLAs and provide ongoing reports.
- **Managing fluids:** Any fluid that passes through manifolds needs to be exceptionally pure. A partner should sample fluids at least twice a year, send it to a lab for testing, and take any corrective actions to restore fluid quality.
- **Overseeing asset management:** Companies that rely on liquid cooled servers in their global data centers will likely seek to outsource system management. A partner can manage worldwide installations, meet SLAs, and provide reporting on when fluids and systems were last checked and what maintenance has been performed.

It's important to note that certain IT equipment vendors will have specific requirements related to the safe operation of their products, including the application of liquid cooling systems. It is recommended to verify any proposed solution with the IT equipment vendor to avoid potential voided warranty coverage.

Moving Forward With Hybrid Cooling

Data center owners and operators have many decisions to make at the rack and row level as power densities increase at their existing data centers to meet business and customer demands for HPC workloads like analytics and AI.

Navigating the critical infrastructure changes associated with higher rack power densities can be a daunting task. IT and facility teams can tap the expertise of design consultants to design hybrid cooling systems that meet their business and technical needs. They can also access experts at companies that manufacture critical infrastructure thermal management solutions to learn which solutions best meet their requirements and budgets.

Before recommending liquid cooling solutions, Vertiv works with customers to answer critical questions about their facilities, workloads, and more. Vertiv provides products that support the entire thermal management chain (air cooling, liquid cooling, and outdoor heat rejection), power distribution equipment, racks and containment, and remote monitoring software to provide end-to-end solutions and simplify procurement and other lifecycle processes. Finally, customers can leverage our remote monitoring, preventive maintenance, and troubleshooting services to keep liquid cooling solutions performing at peak levels.

Ready to buy?

Contact Vertiv for support in developing the optimal hybrid cooling infrastructure design to meet your business and site requirements.

Contact Vertiv for a complimentary discovery call, where we'll consult with you to learn about your requirements and work with you to develop a solution that will meet your current needs and position your facility for future growth.

Schedule your no-cost discovery session today.



Appendix

A Creating the Optimal System Design

B Glossary

Appendix A: Creating the Optimal System Design

The chart below provides guidance to teams getting ready to source and procure liquid cooling equipment. It lists product groups that should be purchased at the same time to streamline procurement and integration processes.

Product Type	Install Type	Description	Category	Budgeting
IT Rack	New	Racks support various equipment, including servers, storage, switches, routers, PDUs, UPS units, console port servers, and KVM switches.	Racks & Containment	The IT and facility teams should work together to evaluate these solutions to ensure they fit the thermal management solution
Cable and Airflow Management	New	Vertical cable manager accessories for the rack; provide a vertical row of cable management fingers.	Racks & Containment	
Rack Cable Trough	New	Cable trough for overhead power and network cabling.	Racks & Containment	
Blanking Panels	New	Closes unused U space within the rack, promoting proper airflow and minimizing bypass air.	Racks & Containment	
Coolant Distribution Unit	New and Retrofit	Designed to remove barriers to liquid cooling in an air-cooled environment. CDUs are available as liquid-to-liquid heat exchanger for chip and rear-door cooling applications that offer easy, cost-effective deployment in any data center.	Liquid Cooling	IT will provide power-density information, while the design consultant will provide input. These solutions should be bought early in the design selection.
CDU Accessories	New and Retrofit	Water detection cables provide integrated leak detection for the CDU.	Leak Detection	
CDU Hoses	New and Retrofit	One pair each for primary and secondary distribution.	Liquid Cooling	
Distribution Manifold	New and Retrofit	In-row manifolds enable support of liquid-cooled IT equipment by distributing fluid from CDUs to individual racks.	Liquid Cooling	
Rack Manifold Set	New and Retrofit	The in-rack manifold provides supply and return to connect liquid between the row/room facility cooling source and the server cold plates.	Liquid Cooling	It may be an option based on data center team needs and preferences. The condenser will be sized depending on ambient conditions and the thermal capacity required.
Indoor Chiller	Retrofit	Working in tandem with a RDHx, indoor chillers integrate pumped refrigerant economization to achieve higher efficiency. They quickly and efficiently cool pods of high-density racks in an air-cooled data center. Liquid-to-air heat exchangers are located in the rack rear door.	Heat Rejection	
High Density Condenser	Retrofit	A rooftop high-density condenser designed for use with an indoor chiller.	Heat Rejection	
Freecooling Chiller	New	An outdoor freecooling chiller that provides chilled water to CRAHs in an already existing solution. If correctly sized it can also be connected directly with the CDUs. A new larger freecooling chiller can be implemented to both supply cooling capacity to CRAHs and CDUs.	Heat Rejection	
Rack PDU	New and Retrofit	In-rack PDUs with switched outlet-level monitoring offer protection for IT equipment.	Power Distribution	The design consultant, the specifying engineer, and the facility and power teams should work together to align the liquid cooling solution to power specifications.
Busway	New and Retrofit	Distributes power from the facility to IT cabinets.	Power Distribution	
Busway Components	New and Retrofit	Tap-off box and breakers to support rack PDU and A+B feeds.	Power Distribution	
PDU and-or RPP	New and Retrofit	Floor-mounted PDU and RPP distribute power to the busway and direct rack connections.	Power Distribution	
Switchgear	New or Retrofit	Ensures proper flow of power, protects equipment, and isolates data center power systems.	Power Distribution	These solutions should be bought at the same time as cooling products.
UPS	New	Standalone floor-mounted UPS units provide optimized power per footprint, offering high efficiency, robust electrical protection, and intelligent paralleling to enhance performance at partial loads.	UPS	
UPS Batteries (Greenfield only)	New	Lithium-ion batteries for the UPS provide the best TCO.	UPS	

Appendix B: Glossary

Coolant distribution unit (CDU) – CDUs distribute fluids from the chiller to the data hall to the CDU, racks, and back via a secondary loop separate from the chilled water supply. CDUs can be used with RDHx, cold plates, and immersion cooling. They provide fluid filtration capabilities to protect equipment and offer flow rate and temperature monitoring, enabling teams to adjust both to optimize operations.

Computational fluid dynamics (CFD) – A numerical calculation performed to simulate both turbulent and laminar flow, providing pressure, density, and temperature values across a volume. For data centers, this means analyzing air-cooling and liquid-cooling equipment to determine current and future cooling capacity, air and fluid temperatures and flow rates, and to determine if existing piping is sufficient for the hybrid cooling infrastructure.

Direct-to-chip (DTC) liquid cooling – DTC cold plates sit atop the board's heat-generating components to draw off heat through single-phase cold plates or two-phase evaporation units. These cooling technologies can remove about 70-75% of the heat generated by the equipment in the rack, leaving 25-30% that must be removed by air-cooling systems.

Flow network modeling – An analysis process used to calculate full-system flow rates and temperatures. Data centers benefit massively from this type of study by validating current and proposed fluid network designs, piping sizes, pump sizes, etc. This helps to ensure the infrastructure can support the IT deployed in terms of cooling capacity.

Hybrid cooling infrastructure – A combination of air-cooled and liquid-cooled thermal management units designed to provide increased cooling capacity, efficiency and sustainability.

Liquid-to-air heat exchange – For a liquid-to-air system, the CDU transfers the heat to the ambient environment, utilizing a heat exchange coil design within the CDU to remove heat. The products in this category do not require water pipes to be connected to the building system for heat rejection. These units enable localized liquid cooling for high-output IT equipment but leverage the technologies from existing data center cooling systems to dissipate heat.

Liquid-to-liquid heat exchange – For a liquid-to-liquid system, the CDU transfers heat from one liquid to another for heat removal. Typically this is done by passing the two fluids by each other separated by a thin plate, typically in a plate heat exchanger. Liquid-to-liquid CDUs yield the best cooling performance, however, they require installing pipes and pumps to connect to the facility's water supply.

Manifold (rack and row) – Manifolds are a fluid distribution system that serves to bring coolant to the racks and the individual servers in a data center.

Row manifolds include the plumbing supply and return system to distribute coolant from the CDU to the rack. This is also known as the secondary fluid network. Rack manifolds include the supply and return system to distribute coolant from the secondary fluid network to the liquid cooled servers in the rack.

Power usage effectiveness (PUE) – A metric used to determine the overall energy efficiency of a data center. PUE is determined by dividing the total facility energy usage by the IT equipment energy usage. The lower a data center's PUE value, the more efficiently it is operating. The target value is 1.0 with the average value for data centers falling under 1.6.

Rear-door heat exchangers (RDHx) – Passive or active heat exchangers replace the rear door of the IT equipment rack with a liquid heat exchanger. These systems can be used in conjunction with air-cooling systems to cool environments with mixed rack densities.

Redundancy (N, N+1, N+2, 2N) – additional units deployed to avoid the risk associated with unit failure. N is defined as the minimum infrastructure (power, cooling, etc.) required to support the IT deployment in the data center and does not include any extra units. Systems with N+1 and N+2 redundancy have an extra one or two infrastructure units deployed, respectively. Systems with 2N redundancy have full 1:1 redundancy built in for each infrastructure component in the data center.

Scope 2 emissions – Indirect greenhouse emissions associated with the purchase of electricity, steam, heat, or cooling. Though the processes to produce these occur at the facility where they are generated, the organization is required to account for these as a result of the data center's use of the energy.

Thermal design power (TDP) – The maximum amount of heat generated by a central processing unit (CPU), GPU, or chipset, as measured in watts. The higher the thermal design power of a component, the more cooling capacity is needed to maintain efficient operation.

Total cost of ownership (TCO) – The amount required to purchase equipment plus the costs of operation and maintenance. By deploying more efficient equipment in the data center, teams can lower the cost of operation, and thus reduce TCO over the life of the product.

Water usage effectiveness (WUE) – The ratio between the amount of water used in the data center systems (water loops, adiabatic towers, humidification, etc.) and the energy consumption of the IT equipment. The lower a data center's WUE ratio, the more efficient its use of water resources. The average data center using evaporative cooling technologies has a WUE of 1.8 liters/kW.



Vertiv.com | **Vertiv Headquarters**, 505 N Cleveland Ave, Westerville, OH 43082, USA

© 2023 Vertiv Group Corp. All rights reserved. Vertiv™ and the Vertiv logo are trademarks or registered trademarks of Vertiv Group Corp. All other names and logos referred to are trade names, trademarks or registered trademarks of their respective owners. While every precaution has been taken to ensure accuracy and completeness here, Vertiv Group Corp. assumes no responsibility, and disclaims all liability, for damages resulting from use of this information or for any errors or omissions. Specifications, rebates and other promotional offers are subject to change at Vertiv's sole discretion upon notice.