

Navigating Liquid Cooling for AI Driven Data Centers

Key insights for
technology leaders

Many AI servers with accelerators (e.g., GPUs) used for training LLMs (large language models) and inference workloads, generate enough heat to necessitate liquid cooling.

These servers are equipped with input and output piping and require an ecosystem of manifolds, CDUs (coolant distribution unit) and outdoor heat rejection. There are six common heat rejection architectures for liquid cooling where we provide guidance on selecting the best one for your AI servers or cluster.

Let's address why liquid cooling is important and its benefits to your data center.

Life Is On

Schneider
Electric

Why liquid cooling matters



Increased heat from AI Workloads

AI workloads generate over ten times more heat than traditional servers, challenging conventional cooling methods.



Limitations of air cooling

Air cooling limits the number of GPUs in a AI rack, while AI prefers more GPUs for parallel processing. Liquid cooling helps to increase GPU density.



Need for advanced thermal management

Liquid cooling solutions are essential to maintain performance and reduce downtime in AI-intensive data centers.



Strategic infrastructure upgrade

Adopting liquid cooling enables a long-term strategy to meet growing AI application demands sustainably.

Business benefits of liquid cooling

But liquid cooling goes further. It improves performance, boosts energy efficiency, and supports sustainability in data centers.



Cost reduction and efficiency

Liquid cooling reduces operational expenses through enhanced energy efficiency in data centers.



Increased accelerator reliability

Liquid cooling can improve chip and hard drive reliability by providing a lower stable operating temperature.



Sustainability advancement

Liquid cooling contributes to sustainability by lowering water and energy usage in data centers.



Scalability for future demands

Scalable liquid cooling solutions meet next-generation workloads and evolving infrastructure needs.

Life Is On

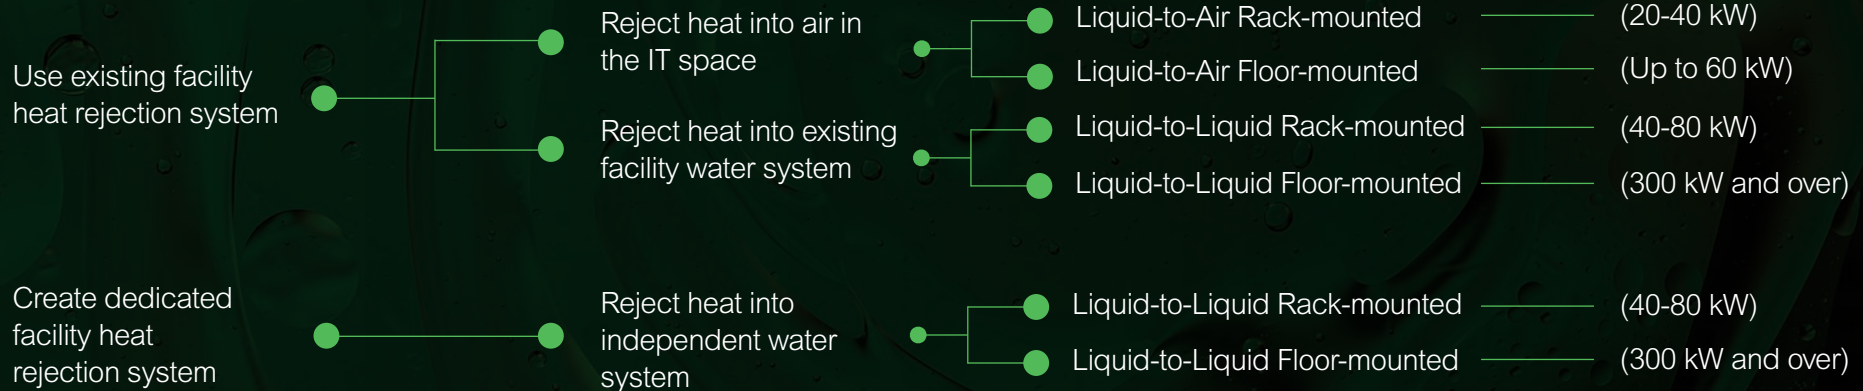
Schneider
Electric

Six liquid cooling architectures

Heat rejection method

CDU type

CDU capacity



Definitions

Liquid to air (L-A)

Technology cooling system (TCS) liquid loop heat is pumped to a coil (i.e., radiator) where the heat is rejected directly into the data center air

Liquid to liquid (L-L)

TCS liquid loop heat is transferred to a facility water system

Rack-mounted

CDU mounted within a rack provides a TCS loop for a single rack and can be pre-integrated with servers.

Floor-mounted

CDU provides a TCS loop for several racks. Form factor may be similar to an IT rack or larger as capacities increase

Author



Victor Avelar

Chief Research Analyst, Data Center
Research and Strategy Group at
Schneider Electric

[LinkedIn](#)

If you want to learn more about what is the best liquid cooling architecture for your organization and best practices, explore the following content:



Executive Report - Optimizing AI Infrastructure: The Critical Role of Liquid Cooling



White Paper - Five Reasons to Adopt Liquid Cooling



White Paper - Navigating Liquid Cooling Architectures for Data Centers with AI Workloads

Schneider Electric

For more information, please contact us at
dcsc@se.com

[se.com](https://www.se.com)

THIS DOCUMENT IS TO BE CONSIDERED AS AN OPINION PAPER PRESENTING GENERAL AND NON-BINDING INFORMATION ON A PARTICULAR SUBJECT. THE ANALYSIS, HYPOTHESIS AND CONCLUSIONS PRESENTED THEREIN ARE PROVIDED AS IS WITH ALL FAULTS AND WITHOUT ANY REPRESENTATION OR WARRANTY OF ANY KIND OR NATURE, EITHER EXPRESS, IMPLIED OR OTHERWISE.

©2025 Schneider Electric. All Rights Reserved. Schneider Electric | Life Is On is a trademark and the property of Schneider Electric SE, its subsidiaries, and affiliated companies.

Life Is On

Schneider
Electric