Quick Take

Insights into Liquid Cooling in Brownfield HPC and AI Data Centers

Jaclyn Ludema, Mark Nossokoff, and Earl Joseph
March 2025

HYPERION RESEARCH OPINION

Artificial intelligence (AI) and high-performance computing (HPC) workloads continue to drive higher power densities in data centers, and liquid cooling is becoming increasingly important for maintaining performance and preventing thermal throttling. These data center cooling requirements have reached a critical juncture, particularly for existing facilities, commonly referred to as brownfield facilities. As scientists and researchers increasingly integrate AI into their research and operational workloads, data centers will seek to adapt their cooling infrastructure to remain relevant. However, brownfield sites face unique challenges in implementing liquid cooling solutions due to architectural and infrastructure limitations.

Unlike new build data centers, so-called greenfield projects, that can implement optimal liquid-to-liquid (L2L) cooling systems from the ground up, brownfield sites require a more measured approach. L2L systems circulate liquid coolant directly to computing components and then to external heat exchangers, offering maximum efficiency but requiring comprehensive infrastructure planning and build outs. Recent industry research demonstrates that liquid-to-air (L2A) cooling systems can serve as an effective bridge technology for brownfield sites, offering significant improvements in cooling capacity and energy efficiency without requiring extensive facility modifications. L2A systems use liquid to cool components but reject heat into the existing air-cooling infrastructure, making them more suitable for incremental implementations.

SITUATION ANALYSIS

A recent Hyperion Research study reveals compelling trends in liquid cooling adoption across the HPC/AI sector. Current data shows that 67.0% of HPC/AI sites have implemented some form of liquid cooling, with adoption expected to reach 79.6% within 12-18 months. More notably, sites using liquid cooling for 50% or more of their cooling requirements are projected to increase from 38.8% to 51.5% during this period.

The cooling implementation strategies of these HPC/AI sites vary significantly, with multiple approaches often employed simultaneously at individual facilities. Based on a survey of 67 HPC sites worldwide, shown in Table 1, rack-level cooling leads adoption at 43.7%, followed closely by direct-to-chip solutions at 40.8% and board-level cooling at 39.8%. Aisle-level cooling, most associated with L2A cooling, is adopted less frequently, at 26.2%. The lower adoption rates for infrastructure-intensive solutions like immersion cooling (20.4%) reflect the practical constraints faced by brownfield sites. These constraints include extensive retrofitting to accommodate tanks, limited floor space, and challenges integrating the system with existing power and cooling infrastructure.

TABLE 1

Please allocate your liquid cooling implementations TODAY between these approaches (select all that apply)

Liquid Cooling Approach	Percentage of Sites Utilizing at Least Some Amount of this LC Approach	Average Usage Level
Rack-level	43.7%	24.1%
Chip-level	40.8%	35.5%
Board-level	39.8%	28.6%
Aisle-level	26.2%	11.8%
Immersion	20.4%	7.7%
Other	4.9%	3.5%

n=67

Source: Hyperion Research, 2025

The following box plot, Figure 1 represents the dataset of sites using liquid cooling, showing which implementation strategies are used at what percentages. Chip-level, board-level, and rack-level, typically implemented as L2L cooling solutions, are more frequently adopted as comprehensive cooling strategies among the sites represented in this study. This is seen in Figure 1, where the upper bounds of the IQR and dataset maximums highlight more sites going "all in" on one of these implementations.

- Notably, 28.4% of sites use one of the three cooling implementations for all their liquid cooling (13.4% chip-level, 6.0% board-level, and 9.0% rack-level).
- This trend aligns with the potential of greenfield data centers, where L2L cooling infrastructures are often integrated into the design and construction phase, enabling deployment of these systems.

Interpreting the Box Plot

For Figure 1 below, box plots are used to identify dataset tendency, variability, and potential outliers. The data is represented with two averages; the line inside the box represents the median, while the "x" marks the mean of the data. The box represents the interquartile range (IQR), which contains the middle 50% of the data. The "whiskers" extend from the box to the minimum and maximum values within a designated range considered representative of the dataset. The size of the IQR and length of the whiskers indicate the variability within the dataset. Values outside this designated range are considered outliers and are represented by individual dots.

FIGURE 1

Liquid Cooling Implementations

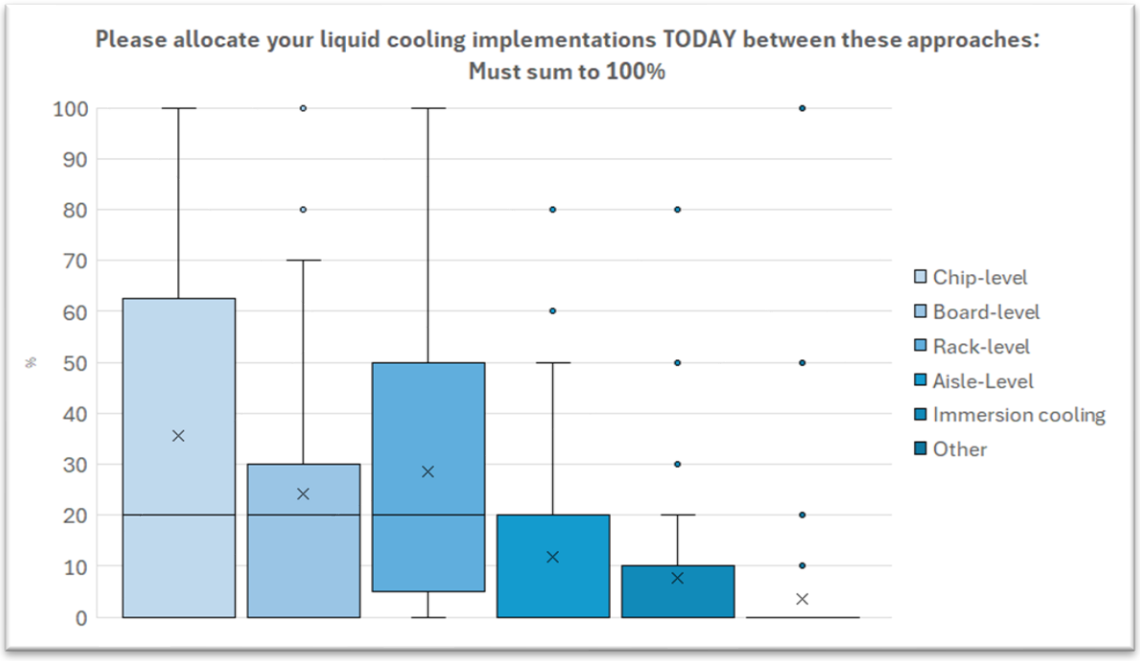

n=67

Source: Hyperion Research, 2025

Brownfield sites, on the other hand, have a different narrative. Since these sites are retrofitting their facilities for liquid cooling, they appear more likely to adopt incremental solutions like L2A cooling, which avoids the need for a complete overhaul of the existing infrastructure. Aisle-level cooling is a good example of this and is a common L2A solution that integrates selectively with the data center infrastructure already in place. Taking a closer look at implementation of aisle-level cooling:

- Aisle-level cooling averages 11.8% usage and tops out at 50%.
- Sites using aisle-level cooling solutions are also using a variety of other cooling solutions in tandem (both liquid and air).
- None of the sites using aisle-level liquid cooling are going "all in" on any given liquid cooling solution.

This data suggests that these aisle-level implementations, mostly at brownfield sites, are taking a practical, step-by-step approach to improving cooling performance gradually while keeping costs and structural changes manageable. This hybrid approach allows these brownfield sites to make the most of their existing infrastructure while steadily upgrading their cooling capabilities to meet modern needs.

Another key finding from this research suggests the strong correlation between AI adoption and liquid cooling implementation. Sites that reported increased AI workloads in the past 12 months tend to show

higher rates of liquid cooling adoption and more ambitious future implementation plans. This trend is driven by the high-power requirements and heat output of the newer generations of GPUs.

Currently, 67.0% of surveyed sites report using some amount of liquid cooling

The data presented in Table 2 illustrates this trend. Currently, 67.0% of surveyed sites report using some amount of liquid cooling for their HPC systems. Among sites that have increased the scale of AI workloads in the past 12 months, this figure rises to 71.6%. A similar pattern appears when looking at near-future plans: 79.6% of all sites anticipate implementing some form of liquid cooling within the next 12–18 months, while 86.5% of sites that have expanded their AI workloads plan to do so.

A similar trend is observed for sites deploying liquid cooling for 50% or more of their cooling requirements. Currently, 38.8% of all sites meet this threshold, compared to 44.6% of sites with expanded AI workloads. Looking ahead, 51.5% of all sites plan to have liquid cooling meet at least half of their cooling needs within 12–18 months, while this figure rises to 55.4% for sites with increased AI workloads.

These numbers highlight the role that liquid cooling appears to play in addressing the thermal and energy challenges of growing scale of AI workloads. Sites experiencing a surge in AI activity are more frequently adopting liquid cooling today and seem to be investing in its expansion at a faster rate. This suggests that as AI workloads continue to scale, liquid cooling will become an increasingly essential component of HPC infrastructure.

TABLE 2

Increases to AI Workloads and Liquid Cooling

Please estimate the liquid cooling deployed for your HPC systems TODAY and in 12-18 months

Have your AI plans changed in the past 12 months? If so, how?

	Liquid Cooling Now	Increased the scale of AI workloads in the past 12 months/Liquid Cooling Now	Liquid Cooling in 12-18 months	Increased the scale of AI workloads in the past 12 months /Liquid Cooling in 12-18 months
Any amount of liquid cooling	67.0%	71.6%	79.6%	86.5%
Liquid cooling for 50% or more of cooling requirements	38.8%	44.6%	51.5%	55.4%
N value	103	74	103	74

Source: Hyperion Research, 2025

FUTURE OUTLOOK

Brownfield data centers are increasingly adopting a staged approach to liquid cooling implementation, favoring solutions that can be deployed incrementally while maintaining operational continuity. This evolution is creating a dual-track market:

- Greenfield sites: Implementation of comprehensive L2L solutions from the ground up
- Brownfield facilities: Adoption of hybrid approaches combining enhanced air cooling with targeted liquid cooling solutions

The success of brownfield sites may depend on their ability to:

- Implement adaptable cooling solutions incrementally
- Maintain operational stability during transitions
- Balance improved cooling capabilities with practical infrastructure constraints
- Support increasing power densities from next-generation HPC and AI systems

This evolutionary approach to cooling infrastructure represents a practical and sustainable path forward for legacy facilities as AI and HPC computing continue to expand. By focusing on pragmatic solutions like L2A cooling, brownfield data centers can effectively support modern workloads while maximizing their return on existing infrastructure investments.

About Hyperion Research, LLC

Hyperion Research provides data-driven research, analysis and recommendations for technologies, applications, and markets in high performance computing and emerging technology areas to help organizations worldwide make effective decisions and seize growth opportunities. Research includes market sizing and forecasting, share tracking, segmentation, technology, and related trend analysis, and both user & vendor analysis for multi-user technical server technology used for HPC and HPDA (high performance data analysis). Hyperion Research provides thought leadership and practical guidance for users, vendors, and other members of the HPC community by focusing on key market and technology trends across government, industry, commerce, and academia.

Headquarters

365 Summit Avenue
St. Paul, MN 55102
USA
612.812.5798
www.HyperionResearch.com and www.hpcuserforum.com