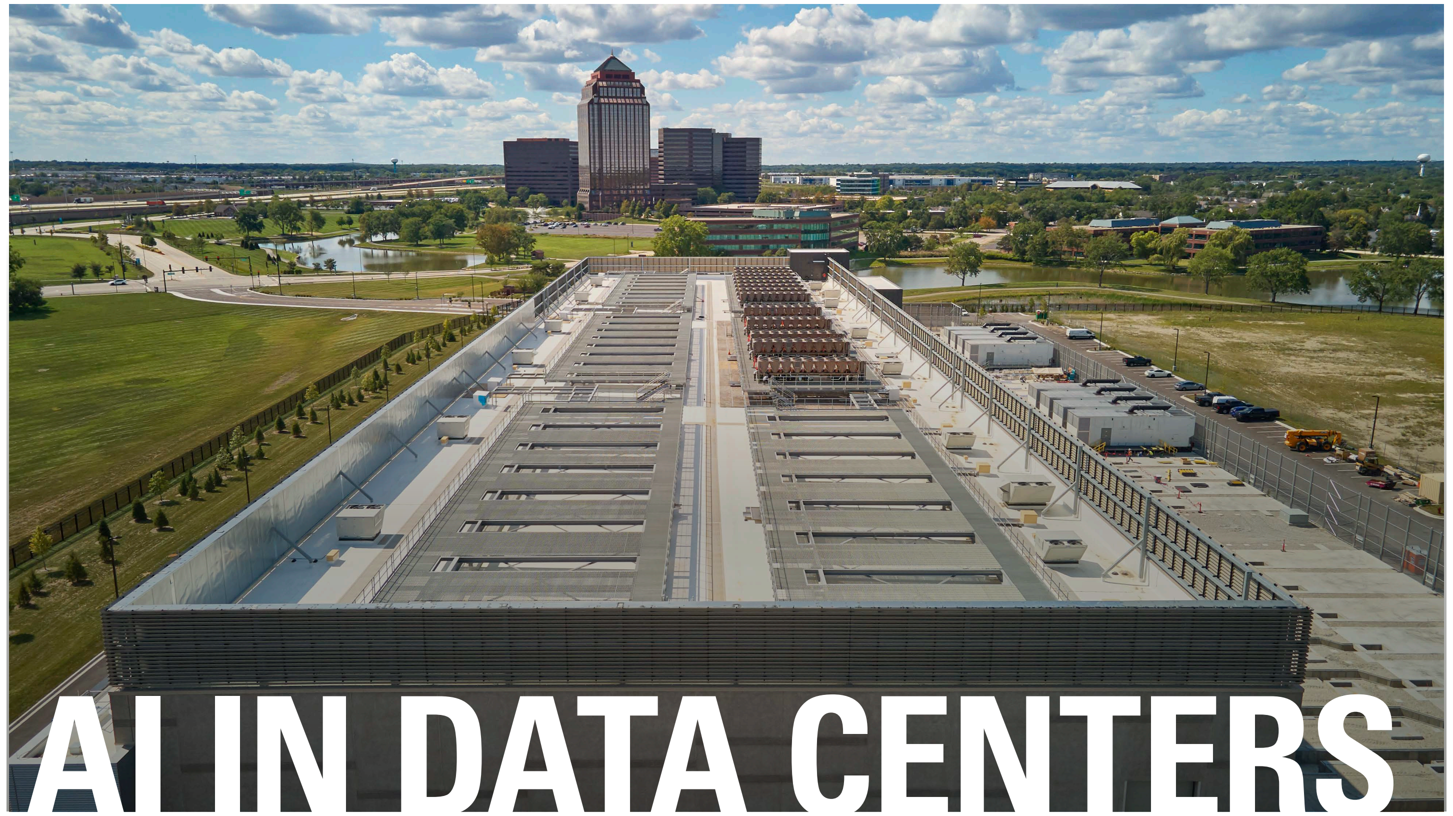

AI IN DATA CENTERS

Cooling Considerations for the Demands of Data Center Artificial Intelligence

Everyone is talking about Artificial Intelligence (AI). It is swiftly revolutionizing industries and applications, creating new efficiencies and value for organizations across the globe.

However, this rapid advancement comes with a potential consequence for data centers – increased heat generation. AI applications, especially those involving deep learning and complex computations, demand significant computational power and data processing capabilities. As a result, data centers are required to accommodate more powerful and energy-intensive hardware and utilize equipment for extended periods, leading to higher heat output.

Consequently, cooling systems must work harder to maintain an optimal operating temperature, resulting in increased energy consumption and potentially reduced cooling efficiency. As the demand for AI evolves, our data centers must evolve also. At HED, we're exploring each project's demands, expected development, and potential future needs when considering cooling systems and carefully weighing the pros and cons of cooling technologies, such as:

DIRECT EXPANSION (DX) COOLING

Pros | Refrigerant based cooling systems are a tried-and-true technology for data centers, offering scalability and grouped control.

Cons | AI is pushing the boundaries of what can be effectively cooled by direct air. For data centers with only DX, investment in supplemental liquid cooling infrastructures may be required to meet the demands of high-density AI deployments. For spot deployments, In-Row close-coupled DX cooling solutions with containment may meet demands.

CHILLED WATER COOLING

Pros | A closed-loop chilled water system can be operated at high efficiency for denser loads and provides the infrastructure for adding direct to chip cooling within the data center. This provides the flexibility to convert to liquid, being the primary cooling agent for liquid cooled servers, with directed air to make up the difference as applicable.

Cons | Chilled water may not be as practical for lower density colocation environments with uncertain load demands, cabinet/cage arrangements, and phased buildouts. Chilled Water may require a significant initial investment in cooling and structural infrastructure.

HOT AISLE / COLD AISLE CONTAINMENT

Pros | Containment solutions are a must in the absence of traditional raised floor deployments to maximize the efficiencies of room cooling and/or In-Row cooling solutions. With raised floor deployments they can greatly increase the control and efficiency of room-based cooling solutions.

Cons | Containment may not always be practical or uniform in colocation and other environments and adds additional cost to infrastructure. There are additional considerations for fire protection, equipment temp ratings, and cable ratings for use in plenum environments.

REAR DOOR HEAT EXCHANGERS

Pros | Rear door heat exchangers can be deployed as needed to high density cabinets without disrupting the functionality of existing perimeter-based cooling systems.

Cons | While heat removal directly downstream of the heat source is a highly efficient methodology, at scale it is costly and may be less practical than considering supplemental liquid cooling alternatives. While a rear door exchanger is an efficient method to remove heat and prevent recirculation, it does not guarantee adequate cool air delivery for high and ultra-high density AI compute environments.

IN-ROW COOLING

Pros | Close-coupled In-Row cooling is an efficient means of removing heat with lower fan energy use and is typically deployed with a hot or cold-aisle row or rack-based containment system. High density environments can be hosted in a smaller footprint than a typical perimeter-based approach, and can be deployed in environments not originally designed to host IT.

Cons | Impractical to deploy in existing server rows. It is typically situated between or above racks, making it a viable containment solution package for unused floor space, but adds significant cost in addition to typically already installed perimeter cooling. Requires refrigerant or chilled water pipe to be routed in the white space, adding infrastructure cost in piping and to manage or contain water leakage.

DIRECT-TO-CHIP LIQUID COOLING

Pros | Liquid is a far better medium for heat removal than air and can efficiently manage the heat removal of much higher rack densities than is practical for air alone. Liquid to Chip Cooling puts the cooling medium precisely where it needs to be, making for highly efficient operation. For 100% liquid cooled deployments, raised floor and traditional containment is not necessary and there are no cooling fan losses to factor in.

Cons | There is no margin for error with liquid to the chip deployments, as a failure in the system can lead to immediate overheating and lengthy downtimes. Water quality is paramount and conditioning requirements may exceed that of a typical chilled water system, adding infrastructure and maintenance costs. The industry is still in early stages of development so upfront costs are still high and there is a very limited install base and commodity products to choose from.

Anticipating the impact of AI on heat generation is tricky. It's safe to assume the demands of all facilities will increase – but how much is unclear. To be prepared, data centers should be designed with a plan for adequate cooling solutions to tackle an increase in heat generation from todays' benchmarks. Innovative cooling technologies, such as liquid cooling and advanced airflow management may become more prevalent to dissipate the higher heat loads efficiently or incorporate AI-driven cooling systems that intelligently adapt to dynamic heat fluctuations, optimizing cooling resources in real-time. In situations where future demand is difficult to predict, a combination of systems or planning for expansion of those systems can be the best method for future-sizing.

By Scott Hays, PE | Mission Critical Market Sector Leader

Explore more market and design insights at www.hed.design/insights