WHITE PAPER

CHATSWORTH PRODUCTS

Powering the AI Revolution: Efficient Strategies for High-Density Data Center Environments

By Robb Jones, RCDD, DCDC
Vice President
Chatsworth Products

Published: June 2024

Powering the AI Revolution: Efficient Strategies for High-Density Data Center Environments

Artificial intelligence (AI) is here. From machine learning to new capabilities like generative AI (GenAI) powered by deep learning neural networks, this highly disruptive technology is poised to tackle complex challenges, streamline business, enhance customer experiences, and drive innovation and creativity. Experts say the potential is incredible – AI could add up to $25 trillion to the world's economy. Tech companies and enterprise businesses are jumping on board. AI adoption among organizations reached 72% in 2024, and the market is expected to explode at a staggering rate of 36.6% from 2024 to 2030. [2,3]

AI offers many advantages across a range of industries and use cases. In IT, AI can streamline complex operations, manage workloads, and enhance cybersecurity with cutting-edge anomaly detection. AI can provide real-time predictive analysis in data centers to optimize power, space, and cooling. The surge in AI adoption also creates demand for more powerful, affordable computer chips, fueling innovation and new business models in the semiconductor industry. The US government is allocating billions in funding under the 2022 CHIPS and Science Act to develop greenfield chip manufacturing facilities that will help keep up with the demand.[4]

One of the most significant impacts of AI is the need for more data center capacity. Research suggests the total capacity of hyperscale data centers will nearly triple in just six years. As enterprise businesses invest in purpose-built AI models for specific use cases, multi-tenant colocation and on-premises enterprise data centers will also require more capacity. This surge in capacity places unprecedented demands on data center infrastructure due to significantly higher rack power densities and higher-density fiber optic connections. At the same time, rising energy prices, regulatory pressures, and corporate sustainability initiatives require data centers worldwide to curtail their energy consumption and reduce their carbon footprints.

The good news is that with the right strategies and solutions, data centers can cost-effectively meet the needs of high-density AI environments while keeping things cool – for both the equipment and the planet.

Organizations that have adopted AI in at least 1 business function,[1] % of respondents

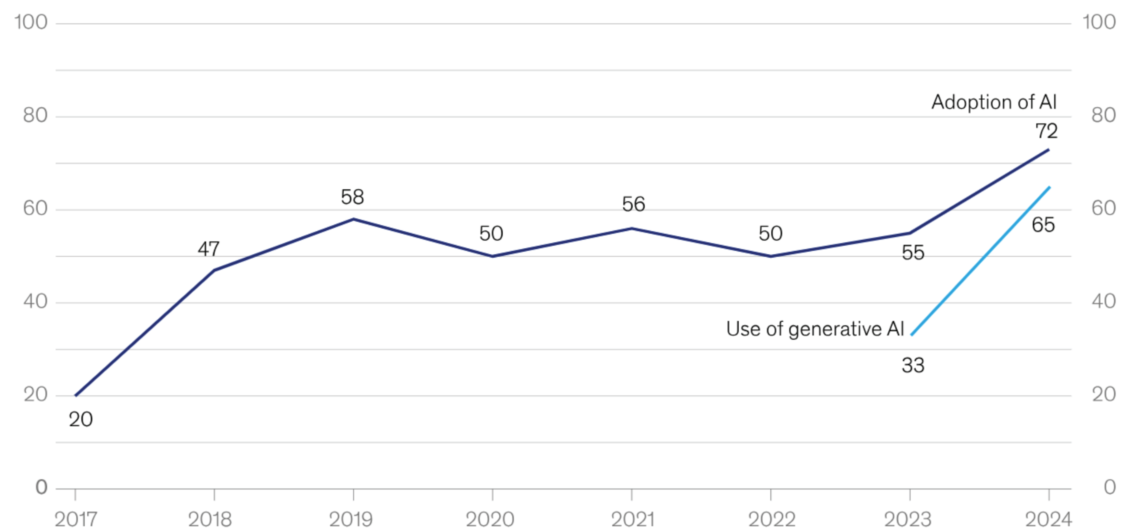

Powerful AI Processing Drives Up Rack Power Density and Heat

Traditional rule-based AI models excel at performing specific tasks and identifying patterns in historical data to make future predictions. Voice assistants like Siri or Alexa, recommendation engines on streaming services like Netflix, and Google's search algorithms are just a few examples of these traditional AI models. These models run smoothly on general-purpose high-performance servers with central processing units (CPUs). But things get more interesting when we step into the world of accelerated GenAI.

GenAI leverages deep learning and neural networks that mimic the human brain to learn, solve problems, and generate innovative ideas and content. This is achieved through two fundamental processes: training and inference. An AI training model analyzes existing data sets to learn a new skill, while an inference model puts that learned skill to use.

AI training models pull massive amounts of data from various sources in parallel. These models execute thousands of complex computations simultaneously. While inference models can run on high-performance CPU servers that perform operations in sequence, CPU servers simply aren't powerful enough to handle the heavy workload of training. This is where graphical processing units (GPUs) come in. A single GPU server can match the processing capabilities of dozens of CPU servers. But with this processing capability comes dramatically increased power consumption. GPUs can consume up to ten times more power than CPUs. Consider an AI training cluster with hundreds or thousands of interconnected GPUs and high-speed switches. The challenge becomes clear—exponentially higher rack power densities.

AI is pushing average rack power densities significantly higher, with forecasts predicting an average rack density of 50 kilowatts (kW) by 2027.[7] This trend is already happening. In 2022, 25 percent of enterprise data centers reported having rack densities greater than 20kW, and some hyperscale data centers have reported reaching even higher densities of 80kW or more[8]. These higher rack power densities lead to higher heat generation, challenging data center operators to effectively cool equipment within the recommended operating temperature range. Exacerbating that challenge is ASHRAE's latest Thermal Guidelines for Data Processing Environment that includes a new class H1 for high-density systems, which narrows the recommended temperature range from 18° to 27°C (64 to 81°F) to 18° to 22°C (64 to 72°F).[9]

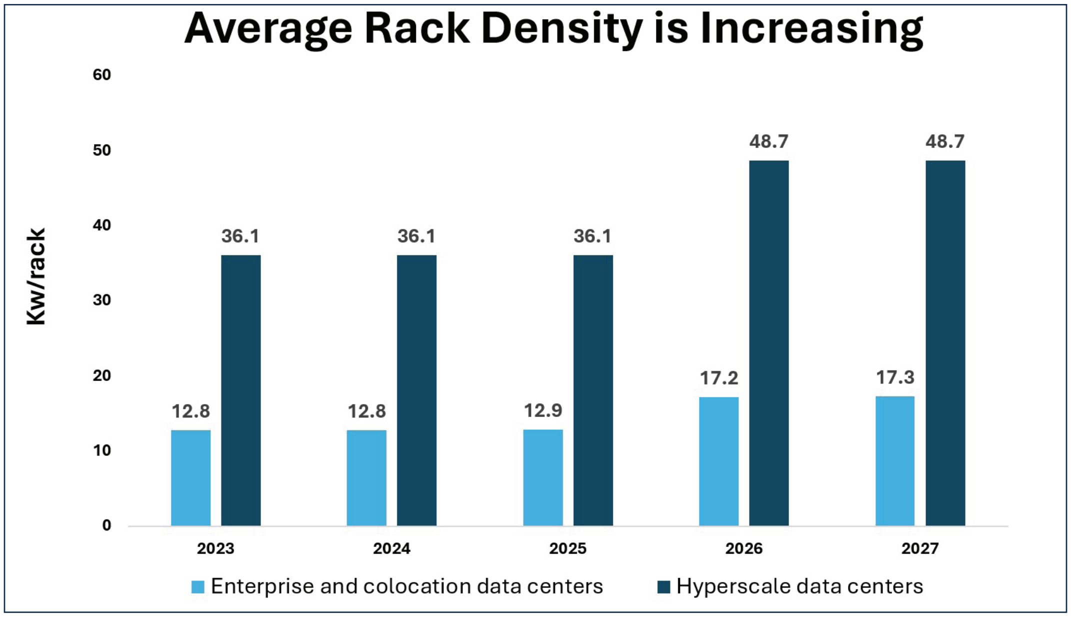

Average rack power densities are increasing, especially in hyperscale data centers. Source: JLL

Air Cooling is Effective for Emerging AI Systems

The tech industry is buzzing about liquid cooling being the next big thing for heat removal, claiming traditional air cooling can't handle high-powered AI systems. Several sources state that air cooling cannot effectively support rack power densities much above 30kW and properly cool equipment with a socket power (i.e., power per individual processor) above 300 Watts (W). However, leading supercomputing companies are selling powerful GPU platforms for GenAI with a socket power of 350W or higher and total power consumption of 10 to 15kW per server as air-cooled solutions. Furthermore, many of these platforms have an operating temperature range that stretches to 30°C (86°F), far exceeding ASHRAE recommendations. The truth is that a well-designed passive air cooling strategy can handle the higher power densities and socket power of emerging AI platforms.

Cold Aisle Containment

Passive air cooling has been around for decades, starting with simple hot aisle/cold aisle configurations in the early days to limit cold supply air from mixing with hot exhaust air and eventually evolving to more effective air isolation via the following methods:

Airflow management: Solutions such as cabinet blanking panels, air dams, sealing, and grommets block airflow around the sides and top of equipment and through open rack units and cable openings to prevent exhaust air from recirculating around equipment. Proper cable management within cabinets also helps prevent cables from restricting airflow.

Hot Aisle Containment

Cold aisle containment (CAC): For use with raised-floor environments or in-row cooling and where there is no overhead hot air return space or duct system, CAC solutions trap cold air into the contained intake space between two cabinet rows. With CAC, the area outside the contained space becomes the hot aisle, elevating overall room temperature.

Hot aisle containment (HAC): HAC solutions direct hot exhaust air out of the contained hot aisle to overhead return air systems. With HAC, the area outside the contained space becomes the cold aisle, resulting in a colder, more comfortable room temperature.

Vertical Exhaust Duct

Vertical exhaust duct (VED): VED solutions, called chimneys, function like HAC but rather than directing hot exhaust air into a hot aisle, they direct hot exhaust air from the cabinets directly to the overhead return air system. This eliminates hot aisles, keeping hot air out of any working space, and resulting in a more comfortable room temperature. VEDs can be used on a single cabinet or shared across multiple cabinets in a row.

Unfortunately, data centers have often deployed these solutions in a way that can only be defined as "partial containment." For instance, some use curtains at aisle ends instead of sealed doors, while others use only doors without overhead panels. Some data centers also neglect to seal empty rack units or cabinet openings. While these partial measures provide some benefit, they don't completely isolate hot and cold air, leading to unused or wasted cooling capacity. Curtain barriers made of flexible plastic strips are leaky and ineffective. End-of-aisle doors alone only prevent airflow from wrapping around the row ends and do little to stop air from recirculating over the tops of the equipment cabinets.

With the higher rack power densities of GenAI, many data centers with partial containment are struggling and forced to lower air supply temperatures to maintain proper equipment operating temperature. However, over-provisioning of cooling is a costly and slippery slope considering rising energy costs and the need to reduce power consumption in line with new regulations. As a result, some data center operators are considering liquid cooling solutions. However, achieving true full containment is effective for emerging AI systems and can be a cost-effective alternative for data centers not ready to deploy liquid cooling.

Achieving Full Containment, the Secret to Effective Air Cooling

Full containment establishes a precisely controlled airflow path through the active equipment with zero mixing of hot and cold air. This data center discipline requires due diligence with a comprehensive airflow management strategy that involves implementing HAC, CAC, or VED solutions and thermal management accessories to create a meticulously sealed airflow pathway.

The first step to achieving full containment is identifying airflow patterns and any unsealed gaps that compromise the isolation of hot and cold air. Computational Fluid Dynamics (CFD) modeling software is a valuable tool for this purpose. It can pinpoint where conditioned air is wasted, identify mismatches between airflow demand and supply, and model the potential impact of implementing solutions. Armed with this information, data center operators can address issues with the following solutions and best practices:

- Leverage Build To Spec (BTS) aisle containment solutions that accommodate varying site conditions and cabinet sizes to completely separate hot and cold air.
- Implement sealed doors rather than curtains and consider automatic doors to ensure they remain closed.
- Ensure a strong seal between ceiling panels (CAC) and duct components (HAC/VED).
- Minimize floor leaks by sealing openings in raised floor tiles with brush-sealed or rubber-gasketed grommets.
- Seal off unused rack-mount spaces in cabinets with filler blanking panels and block airflow under cabinets with bottom panels.
- Optimize cabinet airflow using air dam kits to seal the space between equipment mounting rails and the top, bottom, and side panels of cabinets.

Once full containment is achieved, ideally verified via CFD modeling, data centers can support the rack power densities associated with AI while significantly improving efficiency. Supply air temperatures can be safely raised to the upper limits of ASHRAE guidelines or equipment specifications. This eliminates over-provisioning and dramatically reduces energy consumption and costs.

The beauty of full containment lies in the flexibility of the cooling source. Cold supply air can originate from any source—traditional computer room air conditions (CRAC) units, chilled-water computer room air handler (CRAH) units, air-side economizers, water-side economizers, or a combination thereof. Economizers that leverage outside air or nearby natural water sources offer “free cooling” capabilities that reduce reliance on CRAC/CRAH units.

Scenarios and Solutions for Liquid Cooling in High-Performance Data Centers

Even with full containment, some scenarios could necessitate using liquid cooling. Technically, liquid cooling is only mandatory when the equipment manufacturer specifies it. Remember, even cutting-edge GPU platforms like NVIDIA's DGX H100 and DGX B200 are available as air-cooled solutions, and NVIDIA uses hot aisle containment for air cooled equipment in their turnkey DGX SuperPOD.

Other potential scenarios that could require liquid cooling include:

Data center limitations: If the data center cannot deliver the necessary volume of supply air or its structural design hinders having adequate return space for aisle containment, liquid cooling might be necessary for racks with high power densities.

Hot spots: If hot spots persist in high-compute cabinets, liquid cooling implemented at the cabinet or equipment level can prevent those hot spots from impacting equipment performance.

Rack space constraints: Liquid-cooled GPUs can be ideal for situations with limited rack space. Liquid-cooled GPU chassis are considerably smaller than air-cooled ones that necessitate extra space for heat sinks, typically offering a space saving of up to three times.

Maximizing compute performance: Liquid-cooled GPUs generally deliver the highest processing speeds, making them suitable for applications demanding maximum available compute power.

Data centers that require liquid cooling should understand the various options to choose the solution best suited to their specific environment, operations, and budget. Here's a breakdown of the choices:

Rear-door heat exchangers: These cabinet-level cooling solutions utilize liquid-filled coils (chilled water or dielectric fluid) installed in the rear door to cool hot exhaust air from equipment and return it to the room at ambient temperature. While considered a liquid-cooled device that requires a liquid supply system, plumbing to and from cabinets, and some type of heat rejection unit to cool the heated liquid before returning it to the coils, rear-door heat exchangers are not technically a true liquid cooling method since the equipment itself is still air-cooled and requires the use of fans. They can be considered more of an active containment solution.

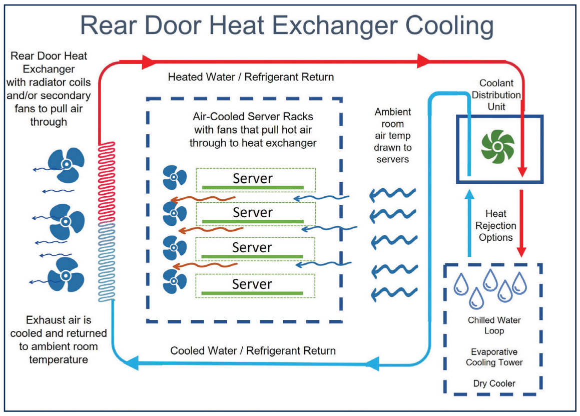

While rear-door heat exchangers support significantly higher densities than traditional open-air return, there are some key considerations. Rear-door heat exchangers place limitation on cabinets and are expensive and unreliable, resulting in higher CapEx and OpEx due to more frequent maintenance, replacement parts, condensation, pressure control, and coil and fan performance issues. They also have the potential for air leakage if not properly fitted and sealed and more frequent maintenance. It's important to remember that rear door heat exchangers require redundancy. While secondary cooling loop water supplies are one way to achieve redundancy, implementing a full containment strategy in conjunction with rear-door heat exchangers is critical to removing heat load and keeping it separated from the cold supply air if doors malfunction or are taken offline for maintenance.

Liquid immersion cooling: This approach submerges electronic components in a coolant, typically a dielectric fluid, within a sealed enclosure or tank. While immersion cooling effectively cools equipment and receives significant industry attention, it necessitates substantial modifications, specialized servers and IT equipment, increasing initial costs. Equipment is also vertically oriented and must be lowered in and out using an overhead lift, which can inhibit vertical scalability, lower power density per square foot, and necessitate more floor space. The increased weight of the fluid-filled tank can also require a higher floor load rating for structural integrity.

Another concern is the large volume of dielectric fluid. These fluids are expensive and evaporate whenever the tank is opened for maintenance. There have also been accounts of dirt, oils, and jacketing materials on cables and other components contaminating the fluid and damaging equipment. Many immersion providers have added filtration to help mitigate these known issues. Additionally, the chemicals used have raised concerns about long-term health effects and potential environmental contamination. Future sustainability regulations may limit their use.

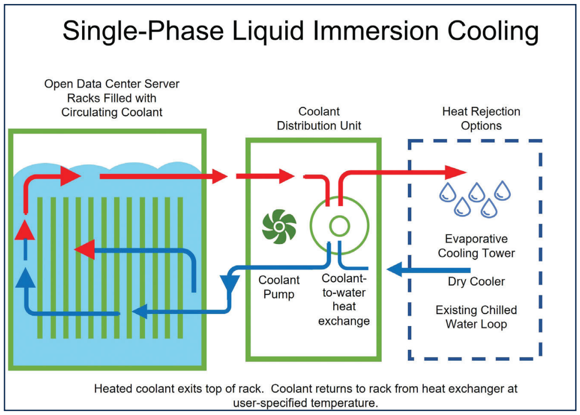

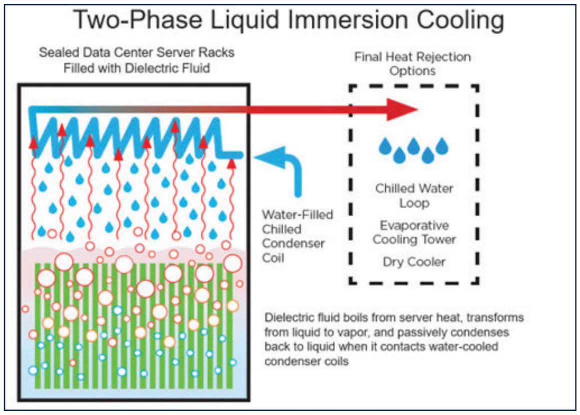

Direct-to-chip liquid cooling: This method involves circulating chilled water or dielectric fluid to cold plates that directly attach to the equipment's heat-generating components (i.e., CPUs and GPUs). Direct-to-chip is considered the most efficient liquid cooling solution, and it is much easier and less expensive to adopt than immersion cooling. GPUs with this technology are readily available and fit in standard racks and cabinets in a typical horizontal position with components that are small enough not to impede airflow. This allows them to reside in the same rack as air-cooled equipment in a full containment scenario. Two-phase direct-to-chip solutions use significantly less dielectric fluid than immersion cooling and fully contain the fluid.

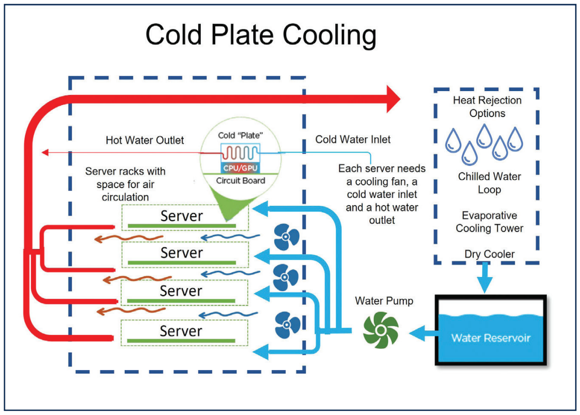

For data centers exploring liquid cooling for the first time, direct-to-chip liquid cooling is an ideal choice due to its simplicity and minimal impact on the data center's overall design—especially since very few data centers need 100% liquid cooling. Direct-to-chip cooling integrates well with existing full aisle containment systems and can be strategically deployed in specific locations like hot spots. The primary considerations are ensuring sufficient cabinet load capacity to accommodate heat rejection units, mounting space for vertical manifolds that distribute fluid to servers, and inlet and outlet connections to the cold plates, all while preserving ample space for PDUs and cable management.

Taming Higher Fiber Densities Within AI Clusters

GPUs for real-time, simultaneous AI training computation connect at 100 to 400 Gigabit speeds. In-cabinet GPU connections in a top-of-rack (ToR) configuration typically connect via short high-speed SFP and QSFP copper direct attach cables (DACs) that offer low power consumption and low latency. Larger end-of-row (EoR) and middle-of-row (MoR) AI configurations rely on parallel optic transmission using multi-fiber MPO connectivity to transmit and receive data over multiple fibers. For example, 100, 200, and 400 Gigabit applications in the data center often use four fibers transmitting and four fibers receiving at 25, 50, or 100 Gigabits per second (Gb/s). GPU servers also frequently connect using breakout configurations to maximize switch port utilization. Breakout configurations allow a single high-speed switch port to support multiple lower-speed connections. In an AI training cluster, it's common to see an 800 Gigabit switch port connecting eight 100 Gigabit GPUs or four 200 Gigabit GPUs. It's easy to see how cabling in an AI cluster is becoming increasingly dense and complex, emphasizing the importance of proper cable management.

Cable Management Bend Radius Category 6A cables

Managing high-density cables within AI clusters isn't as challenging as it might seem. Existing server cabinets and cable management systems were designed initially for Category 6A cables, which have a similar diameter to DACs and about four time greater diameter than fiber. However, there are still some essential fundamental aspects to consider.

Cable management should be able to maintain the specified minimum bend radius of the cabling at critical entry and exit points to avoid damaging the cables. The minimum bend radius is 10 times the cable diameter for DACs and fiber optic cables.

Cable management should have the strength and capacity to handle large numbers of cables routed between GPUs, switches, and patch panels while keeping everything neat and organized.

Proper cable strain relief should be maintained to protect cables and connectors from any strain during cable routing or from hanging cable weight in vertical managers.

Cable management should offer ample space to properly manage cables without compromising airflow or hindering access to components.

Cable management should be flexible and scalable, with the ability to add support, bend radius protection, and strain relief wherever needed in the cabinet.

Cable Management Bend Radius Fiber Optic cables

Balancing Efficiency, Sustainability, and Availability

With cooling systems accounting for up to 40% of total data center energy consumption, effective strategies like full containment are vital to reducing cost, especially with energy prices increasing year over year. Even with effective cooling, data centers must further reduce overall power consumption and carbon footprints to comply with stricter regulations and sustainability initiatives while ensuring high availability and minimal downtime.

The new European Union's (EU) Energy Efficiency Directive (EED) is a prime example. It requires data centers with a demand of 100 kilowatts (kW) or more to report their energy performance. Similar initiatives are underway in the US. The US Securities and Exchange Commission (SEC) is expected to require data centers to disclose greenhouse gas emissions and climate risks. Even individual states are getting involved. Virginia, a major data center hub, is considering strict carbon reduction requirements. At the same time, Oregon has a proposed bill that would force data centers to slash carbon emissions by 60% within a few years.

Following are several key strategies for balancing efficiency, sustainability, and availability in AI data center environments:

Free Cooling and Sustainable Power: Leverage air- or water-side economizers that reduce reliance on cooling systems and chilled water. Procure or deploy renewable energy sources like solar, wind, and hydro coupled with battery storage to lessen grid dependence and lower carbon emissions.

Heat Reuse and Circular Economy: Capture equipment-generated heat and reuse it to heat other spaces or nearby buildings. Embrace circular economy principles by prioritizing equipment reuse and recycling to minimize waste.

High-Efficiency IT Equipment: Implement energy-efficient IT equipment and utilize power management features to automatically reduce consumption during low-use periods.

Granular Power Monitoring and Optimization: Deploy intelligent PDUs that monitor power usage to prevent overloads and ensure optimal power distribution, avoiding stranded capacity. Leverage PDUs with outlet-level power monitoring to identify lightly used equipment and "ghost servers" that can be consolidated or removed to eliminate wasteful power and cooling over-provisioning.

Centralized Management Software: Deploy intelligent PDUs backed by centralized management software or integrated with DCIM (data center infrastructure management) tools to calculate PUE, track historical power trends, and support regulatory compliance, efficiency measurements, and future capacity planning.

Environmental and Security Monitoring: Continuously monitor environmental factors like temperature, humidity, and airflow to identify and prevent hot spots, ensuring the effectiveness of cooling systems. Implement sensors to detect leaks, smoke, and other potential hazards, preventing equipment or cooling system failures that could cause downtime.

Electronic Access Control: Enforce proper user access with electronic cabinet access control systems that log access attempts to prevent human error and physical security breaches while ensuring compliance with security regulations.

CPI is Ready to Help You Take on the AI Revolution

CPI's ZetaFrame® Cabinet System uses a modular, holistic approach ideal for effectively supporting critical, high-density AI environments. This innovative ecosystem integrates thermal management, cable management, power management, environmental monitoring, and access control—all working together as a single-vendor turnkey platform that supports increasing power and fiber densities while reducing costs and optimizing efficiency, sustainability, and availability for AI environments in hyperscale, multi-tenant, enterprise, and edge data centers.

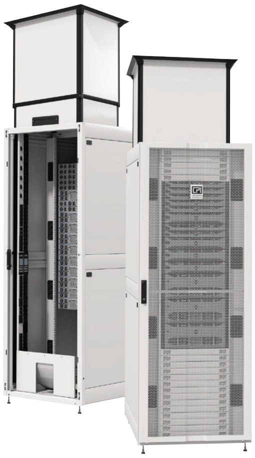

ZetaFrame® Cabinet: The cornerstone of the ZetaFrame System is the highly engineered, configurable ZetaFrame cabinet that offers the industry's highest load capacity, superior scalability, and enhanced high-density cable management. ZetaFrame cabinets are fully customized and can be purpose-built to integrate with advanced direct-to-chip liquid cooling solutions while fully accommodating PDUs and cable management.

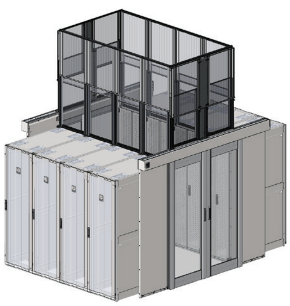

Aisle Containment and Thermal Management: ZetaFrame cabinets seamlessly integrate with CPI's industry leading hot aisle containment solutions, cold aisle containment solutions, and thermal management accessories that are highly efficient and built to spec to effectively achieve full containment.

eConnect® Intelligent PDUs: Designed to support higher rack power densities with up to 57kW on a single PDU, eConnect PDUs seamlessly integrate with the ZetaFrame cabinet. eConnect PDUs provide real-time power monitoring and control at the rack and equipment level and can be outfitted with various sensors to monitor environmental conditions and hazards. eConnect PDUs also seamlessly integrate with access control door locks, enabling power, management, and control of cabinet-level security while eliminating the need for additional infrastructure to connect and power locks.

Power IQ for eConnect: eConnect PDUs are backed by Power IQ®, a centralized management platform that provides data center operators with real-time measurements and alerts and the ability to calculate and report on energy consumption, PUE, access control, and other critical factors for complying with regulations, improving efficiency, and maximizing availability.

And that's not all. The ZetaFrame ecosystem is backed by CPI's unparalleled support, extensive expertise, and superior R&D, manufacturing, and logistics. From airflow and thermal analysis to capacity assessments and rapid built-to-order customization, CPI is ready to help you tackle the challenges of high-density environments to take on the AI revolution.

References

1 AI's Economic Potential Could Reach $25.6 Trillion, McKinsey Predicts, Tech.co, June 2023

2 The State of AI in Early 2024, McKinsey & Company, May 2024.

3 Artificial Intelligence Market to Reach $1,811.75 Billion By 2030, Grand View Research, May 2024

4 Biden-Harris Administration Announces Preliminary Terms with TSMC, Expanded Investment from Company to Bring World's Most Advanced Leading-Edge Technology to the U.S., U.S. Department of Commerce, April 2024

5 Hyperscale Data Center Capacity to Almost Triple in Next Six Years, Driven by AI, Synergy Research Group, October 2023

6 Uptime Institute Global Data Center Survey 2023

7 Data Centers 2024 Global Outlook, JLL, 2024

8 Looking Back at Data Center World 2023: 3 Concerns That Stood Out, Data Center Knowledge, May 2023

9 New ASHRAE guidelines challenge efficiency drive, Uptime Institute, October 2021

10 Tougher Reporting Mandates Ahead for Data Centers, Data Center Knowledge, February 2023

chatsworth.com
techsupport@chatsworth.com
800-834-4969