WHITEPAPER

THE DATA CENTER ROADMAP TO NET ZERO CARBON

White Paper from Subzero Engineering

By Gordon Johnson

In an era where environmental consciousness is needed and expected, data centers are continually looking for ways to reduce energy usage and minimize their carbon footprint.

Innovative cooling solutions continue to assist data centers achieve energy-efficient cooling, with solutions such as cold or hot aisle containment resulting in significant energy cost savings which in turn is also great for the environment.

Many buildings including data centers are also seeing the benefits of receiving LEED (Leadership in Energy and Environmental Design) certification since this is widely recognized as a global symbol of achievement in sustainability. The LEED rating system is a green building certification program which provides a framework for the design, construction, and operation of high-performance buildings and communities. This certification often includes using recycled materials in the data center white space when available.

Others are taking energy usage and CO2 emission reduction more seriously with ESG (Environmental, Social, Governance) reporting either in place or soon required in their geographic location. Many are contemplating the move to a hybrid cooling infrastructure (simultaneously using air and liquid cooling to cool ITE) to handle expected rack kW increases resulting from AI (Artificial Intelligence) and HPC (high-performance computing) being deployed either presently or in the near future.

As we embrace changes while keeping the principles of sustainability in the forefront, various questions often arise. How large of a role should using recyclable, green, and sustainable materials be in our data centers? If we're making the transition from air to a hybrid cooling infrastructure, are there best practices to follow? How can we make sure we're shaping a greener, more sustainable future by protecting our planet and preserving its resources?

This paper will answer these and many more questions in the following topics:

1. **The Race to get Green**
2. **Start with Containment**
3. **CFD – The Untapped Green Initiative**
4. **Transitioning the White Space to Green Space with Recyclable Materials**
5. **Hybrid Cooling Infrastructure for AI and HPC**
6. **Best Practices for DTC**

The Race to Get Green

Data centers are the backbone of modern digital infrastructure. They store and process data and enable us to access a large collection of services from mobile apps to cloud services. However, they 're also responsible for consuming massive amounts of electricity while generating heat, resulting in carbon emissions and other environmental impacts. With more cases and applications of AI being developed, many data centers are either incorporating the infrastructure now or looking to do so in the next few years to support these high density systems. This means rising power consumption as server and rack kW continue to increase.[1]

This has resulted in data centers racing to get as green and sustainable as possible. This includes focusing their efforts on reducing energy consumption, using renewable energy sources, and optimizing operational efficiency. The industry goal is to reduce the carbon footprint of data centers and minimize their contribution to climate change while cutting organizational TCO (Total Cost of Ownership) by reducing long-term operational costs.[1] The carbon footprint is determined by the GWP (Global Warming Potential), a measure developed by the EPA to compare the global warming impacts of different greenhouse gases. Specifically, it quantifies how much energy the emissions of 1 ton of a gas will absorb over a given period, relative to the emissions of 1 ton of carbon dioxide (CO_2). The larger the GWP, the more a given gas warms the Earth over a certain period of time compared to CO_2 over that time period.

This race to get green is urgent, with data centers already stretching the limits of today's power grid, and the predicted unprecedented growth due to AI requiring even more power at a time when power availability is quickly become a source of local and global tensions. As new AI cases develop so will its impact on the industry, with 20% of data centers expected to have some kind of AI as early as 2026 and the AI market projected to reach US $407bn by 2027.

Start with Containment

It's often said that containment is the fundamental starting point for maximizing cooling capacity and energy efficiency in data centers, and that it's the smallest action with the greatest outcome because instead of consuming energy containment saves energy. Since the best energy saved is the energy not consumed in the first place, containment is an integral and one of the first steps taken in order to make any data center greener and a more environmentally conscious place. This is true if the data center is operating regular density servers, new high density HPC servers needed for AI, or a combination of both since containment makes other energy initiatives more efficient and economically feasible.

What are the results of no containment? It may be possible to maintain enough airflow at the ITE intake to cool up to 6 kW per rack, but this will require more fan energy and lower supply temperatures versus some level of containment. Experts agree that ITE above 6 kW cannot be sufficiently cooled without either cold aisle containment (CAC) or hot aisle containment (HAC).

It terms of how much containment is needed, full or complete containment is always preferred. This option offers optimum airflow concentration and allows for higher rack power levels, assuming sufficient cooling airflow is available. Either full CAC or HAC can support up to 25 kW per rack with a 2-tile wide cold aisle (4'), however by going to a 3-tile wide cold aisle (6') it's possible to support up to and even exceed 35 kW per cabinet.

> **FACT: Containment SAVES energy**
> ***Containment is the fundamental starting point for maximizing cooling capacity and energy efficiency***

Of course, one of the most impressive benefits of containment is the energy savings which in turn lowers the data center's carbon footprint. Since containment separates the cold supply air and hot exhaust air, airflow and cooling optimization can now take place where prior to containment they could not. This includes running cooling unit supply temperatures as high as possible while still keeping server inlet temperatures within ASHRAE's recommended thermal guidelines. In addition, cooling unit fan speeds can be reduced for additional savings.

While airflow optimization and management enhance ITE air intake temperatures, genuine energy efficiency flourishes with cooling optimization. The result? A green and sustainable data center with a lower TCO, additional cooling capacity, lower energy consumption, and a lower overall carbon footprint.

Data Center with Subzero Engineering Containment

CFD – The Untapped Green Initiative

While recent focus has been on building new green and sustainable data centers, one area that often gets overlooked is modernizing existing legacy data centers. One of the easiest and quickest ways to start this process is with CFD (Computational Fluid Dynamics) which creates a 3D model or digital twin of the physical attributes of the data center, including the performance characteristics of the cooling units, ITE, and other significant equipment.

CFD provides a detailed understanding of the complex interactions between airflow, heat removal from the ITE, and cooling systems, enabling data center operators to make informed decisions to reduce energy consumption and improve efficiency. Key benefits include identifying potential server hot spots, validating data center pre-construction, correctly placing ITE and cooling equipment, airflow and temperature optimization, and simulate via transient modeling the ITE ride out time due to power and/or cooling system failure.

CFD aids in the design phase of data centers by simulating a variety of layout configurations. It looks at placement of servers, racks, cooling units and airflow containment systems to ensure optimal thermal performance starting with the design phase down to the operational phase of the data center. This means that designs can be validated while avoiding costly mistakes. For example, CFD can simulate air cooling strategies deployed in conjunction with liquid cooling designs, such as those using air cooling to remove the supplemental ITE heat not removed via the circulating liquid in DTC (Direct-To-Chip) cold plate cooling.

Legacy data centers can and should take advantage of the many benefits that CFD provides, including simulating proposed changes to the existing room layout such as capacity redistribution or ITE expansion. In all cases, a CFD is always beneficial because it gives legacy data centers a scientific foundation to build upon when executing data center modernization strategies without relying on costly testing or inaccurate assumptions. Furthermore, it can extend the lifespan of an existing data center rather than investing heavily in new data center space or moving to colocation and/or cloud providers.

Additionally, CFD now has the ability to accurately reduce risk from power and/or cooling system failure with transient modeling, showing via time increments what will happen to the ITE when either cooling is lost or both cooling and fan power (airflow) fail. When both cooling and fan power fail throughout the data center, cooling is not lost immediately since the fans will take some time to come to a complete stop. CFD can simulate this time period, showing how hot the ITE server inlets will get and how soon this happens before power is restored to the data center.

Armed with data from the CFD analysis, data center managers can easily understand the impact from data center design changes which will help optimize and reduce energy usage while safeguarding the data center.[2]

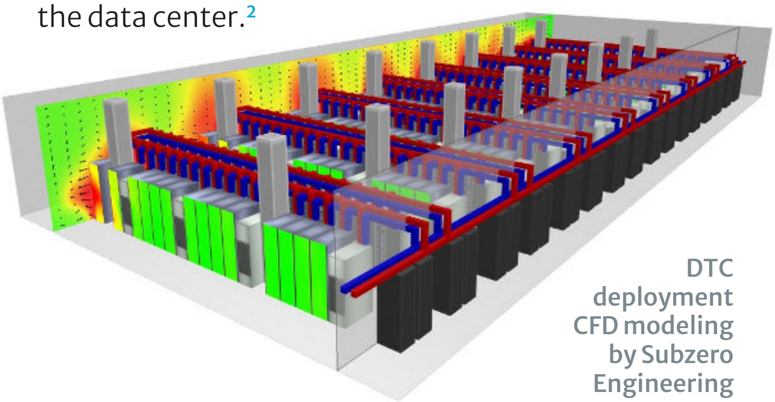

DTC deployment CFD modeling by Subzero Engineering

Transitioning the Whitespace to Green Space with Recycled Materials

Increasing environmental concerns have many looking at additional opportunities to lower the data center GWP by using recyclable materials, in a sense turning the white space into a green space. Optimizing the resource lifecycle of materials is an important aspect of any sustainable data center which includes incorporating recycled materials and designing data centers with a focus on resource efficiency, thereby minimizing their environmental impact.[3]

We've already discussed the importance of containment for lowering energy consumption and the overall carbon footprint of a data center. Let's now dive deeper into the benefits of using recyclable materials for our containment systems.

For example, primary aluminum production involves the Hall–Héroult process which requires high energy and CO2 emissions. Recycled aluminum reduces CO2 emissions by up to 93% and requires approximately 95% less energy compared with primary aluminum production. To put those numbers into perspective, a rough estimate of just 100 containment panels (8' x 2') of 3" recycled aluminum versus primary aluminum can reduce CO2 emissions by 6,819 kg (7.5 Tons) and energy by 24,393 kWh, enough to annually power 2.1 US homes. And that's just 100 containment panels! Significant additional GWP reductions can also be achieved using recyclable material for panels inserts.

By adopting a more sustainable approach to data center design, we can mitigate the environmental impact and create a more efficient and eco-friendly infrastructure.[3] Using green, recyclable materials also reflects positively on gaining credits towards LEED certification, because a green data center always focuses on energy-efficient technologies, including using recyclable materials in the white space, thus contributing to a more sustainable future.

Hybrid Cooling Infrastructure for AI and HPC

With the emergence of AI and HPC comes the challenge to cool these high density workloads. The power density requirements for AI and HPC can be 5-10 times higher than other data center use cases, with power densities approaching and even exceeding 100 kW/rack and predicted to reach 150 kW/rack over the next couple of years. Traditional workload densities can still be air cooled, however most AI and HPC workflows will require some form of specialized cooling technology like liquid cooling.[4]

Data center liquid cooling is an advanced method used to manage the heat generated by high density workloads within data centers. Unlike traditional air cooling systems which rely on fans and ambient air to dissipate heat, liquid cooling uses water or other fluids to transfer heat away from hot components more efficiently, based on the principle that water is a better conductor of heat than air (water has a heat carrying capacity nearly 3,500 times that of air). This approach offers several advantages in terms of efficiency, performance, and sustainability.

> **The Path to Sustainability**
> ***Our data center designs need to be future ready while maintaining flexibility***

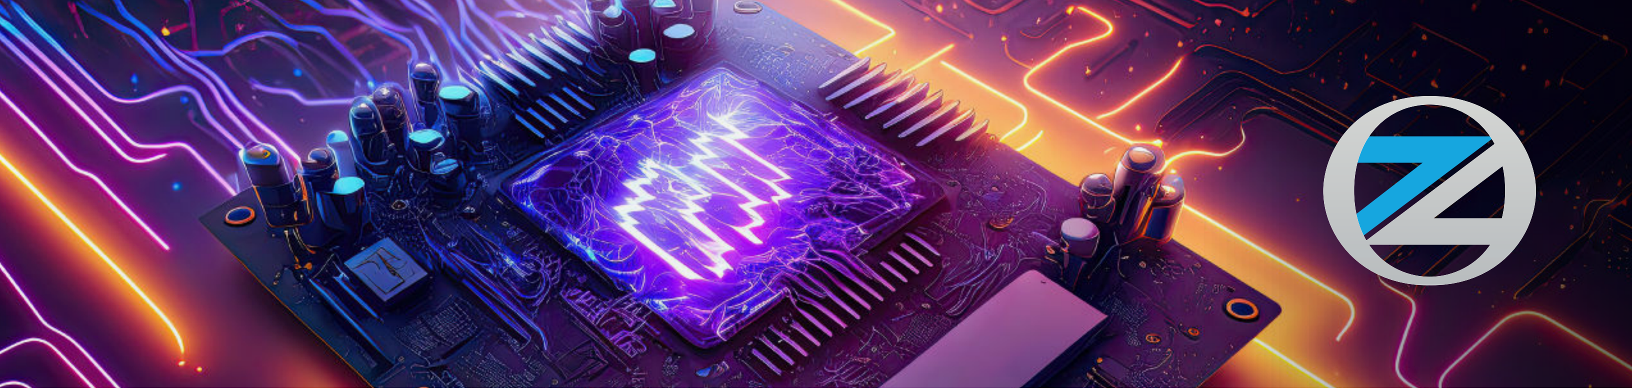

And while we can't predict the future, data center workloads continue to be denser and denser, meaning we'll need to think in terms of both power and thermal cooling when designing new or retrofitting legacy data centers. We'll want to optimize both our power and thermal cooling for higher loads so we can easily transition from lower to higher density systems in the future, and plan on deploying a hybrid cooling infrastructure with both air and liquid cooling into the foreseeable future. Our designs need to be future ready while maintaining flexibility.

While there are several forms of DLC (Direct Liquid Cooling), right now the industry is seeing an accelerated demand for single-phase DTC (cold plate cooling), and that demand continues to grow in part because it keeps the same data center form factor. Latest GPU (Graphic Processor Unit) technology requires liquid cooling, and if a data center is not at least thinking of or preparing for liquid cooling they'll be at a disadvantage moving forward. Single-phase DTC is the dominant form of liquid cooling being deployed right now, and that's the form of DLC that we'll focus on in this paper.

DTC deployments need air cooling since up to 25% residual heat still needs to be removed by air, meaning a 100 kW rack needs approximately 25 kW of heat removed via air cooling (see *DTC deployment with containment needed for air cooling portion of design*). Since we don't necessarily know how much air and liquid cooling we'll need in the future, this highlights the importance of designing holistically for the entire power and thermal cooling needs of our system. We'll also want to do it in an efficient way so that we can easily scale as the ITE density increases, thus future proofing our data centers.

Best Practices for DTC

DTC liquid cooling involves placing cold plates directly onto the processors or other high heat generating components within the ITE. A liquid coolant circulates through these plates, absorbing heat and carrying it away from the components. The shift towards DTC where needed expresses the industries desire to be more energy efficient and sustainable. What are some best practices to keep in mind when deploying DTC with a hybrid cooling infrastructure?

Before deployment and during the design phase, you'll want a trusted partner to perform a CFD study for both legacy and new greenfield designs. This will ensure that sufficient air cooling equipment is available to be leveraged in the new hybrid cooling infrastructure. You'll also want a FNM (Flow Networking Model) analysis to select the correct CDUs (Cooling Distribution Units), size piping for flow rate, choose manifolds, and evaluate the ability of the cold plates and liquid cooling system to support server liquid cooling requirements.[5]

As discussed earlier, containment is an integral part of any data center and DTC deployments are no different. While the majority of ITE generated heat is removed via DTC cold plate heat transfer, additional ITE heat (up to 25%) will still need to be removed by traditional air cooling, meaning either CAC or HAC should always be installed in either retrofitted or new DTC deployments. The chart (See *Hybrid Cooling Infrastructure for 2 MW IT Load*) provides an example of a DTC system designed to cool 2 MW of ITE load (20x 100 kW racks), estimating that the liquid flow rate is 1.5 liters per minute/kW (0.4 gallons per minute/kW) and 150 CFM of air is needed to remove 1 kW of generated heat.

Additional DTC best practices include ensuring power distribution provides for higher amperage required by HPC. Choosing the right rack dimensions is also important to support both the ITE and additional pipe spacing needed. This includes choosing rack heights between 42U and 52U, rack widths @ 800 mm (31.5") minimum, rack depths @ 1,200 mm (48") minimum, and racks that are able to support up to 1,360 kg (3,000 lbs.) of static weight.

Consideration also needs to be given to extra ceiling height due to increased rack height, and other infrastructure choices such as pipe sizing and pipe placement will need to be determined. For example, when piping is deployed under raised floors to avoid the potential for overhead leakage, the piping needs to adequately fit in the space under the floor. For slab floor data centers, the piping for the liquid cooling will need to be supported in some manner such as threaded rods hung from ceiling joists, assuming the ceiling can support this. Power and data cabling will also need to be properly supported, and a trusted containment partner will be able to provide a containment system that supports both power and cable trays.

DTC deployment with containment needed for air cooling portion of design

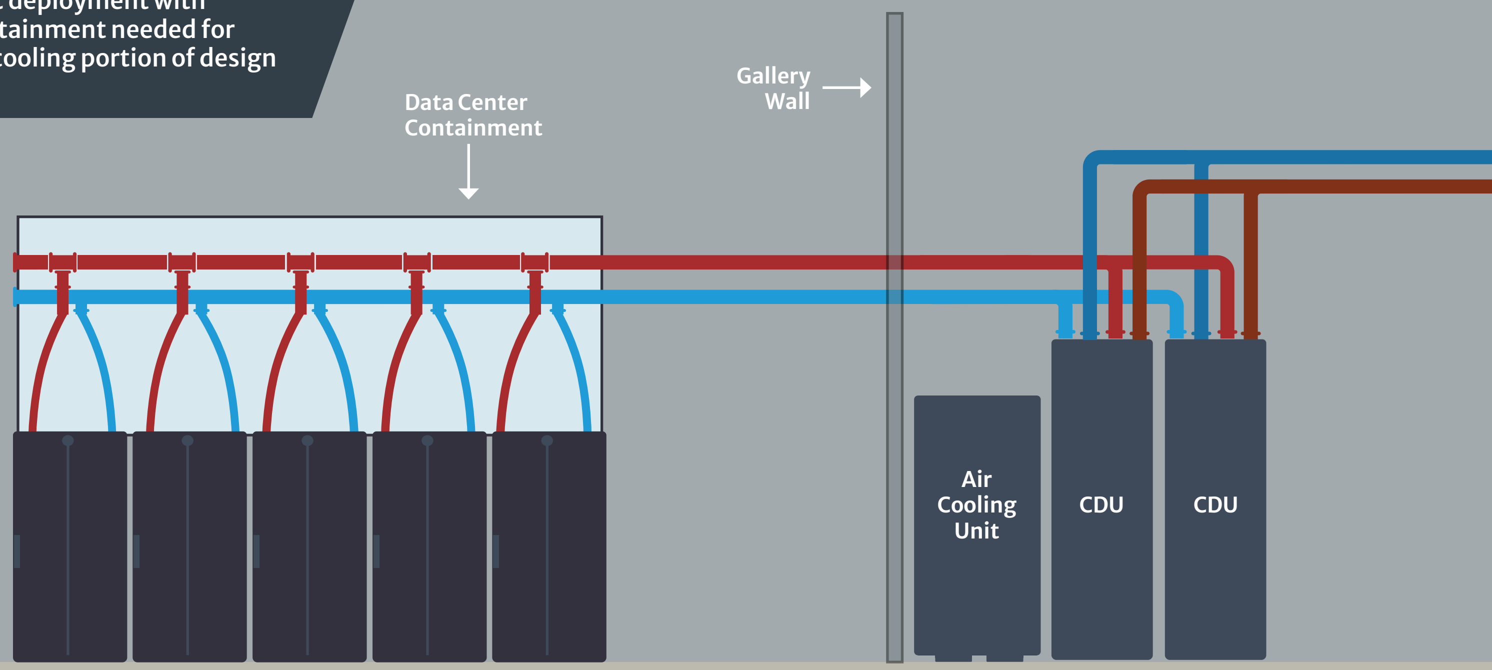

The best practices mentioned in this section are not inclusive but provide just a few examples for DTC deployments, however it's always recommended to work with trusted vendors from start to finish for all hybrid cooling infrastructure projects.

And finally, as data center energy consumption rises, the way we'll achieve ongoing green initiative goals as an industry is by developing and using more sustainable practices which makes net zero carbon targets possible and attainable.

Conclusion

Data center owners and operators have many decisions to make to meet business and customer demands for HPC workloads like analytics and AI. Data centers are producing more and more volumes of data at ever increasing speeds, with AI and HPC at the forefront. Energy consumption is going to continue to rise as the industry struggles to meet sustainability goals and requirements.

Liquid cooling is happening because high density chips and servers are happening, ushering in a new age of data center cooling technology where liquid will just be a portion of the total cooling system, but a lot of air cooling will still be needed as well.

This highlights the importance of working closely with an experienced partner and technology vendor, such as Subzero Engineering, for current and future data center designs. Whether its containment, CFD modeling, embracing innovative solutions for using recyclable materials in the white space, or 20 years of data center design experience, Subzero Engineering can help future proof any data center while lowering the overall TCO.

About the Author

Gordon Johnson is the Senior CFD Engineer at Subzero Engineering and is responsible for planning and managing all CFD related jobs in the U.S. and worldwide. He has over 30 years of experience in the data center industry which includes data center energy efficiency assessments, data center design, CFD modeling, and disaster recovery. He is a certified U.S. Department of Energy Data Center Energy Practitioner (DCEP), a certified Data Centre Design Professional (CDCDP), and holds a Bachelor of Science in Electrical Engineering from New Jersey Institute of Technology. Gordon also brings his knowledge and ability to teach the fundamentals of data center energy efficiency to numerous public speaking events annually, white papers, and industry leading researched articles worldwide.

References

1 Green Data Centers: Adapting to ESG for a Sustainable Future, Posted on April 23, 2024 by Aerodoc

2 Data Center Modernization: The Untaped Green Initiative, Posted on January 18, 2023 by Mark Fenton

3 Green Building Materials: Constructing the Next-Gen Data Center, Posted on October 8, 2023 by Harry Freeman

4 Ahead on the Curve on Advanced Cooling for AI & HPC, Posted on June 4, 2024 by Chris Sharp

5 Deploying Liquid Cooling in Data Centers: Installing and Managing Coolant Distribution Units (CDUs), Posted on March 15, 2024 by Vertiv

SubZero Engineering can lead you down the road to creating a net zero carbon data center. Scan the QR code to learn more.

Hybrid Cooling Infrastructure for a 2 MW IT Load

Quantity of Liquid Cooling Racks 20

Total Power Per Rack	100 kW	**Total Cluster Power**	2MW
Heat to Liquid	75 kW	**Total Cluster Heat to Liquid**	1,500 kW
Heat to Air	25 kW	**Total Cluster Heat to Air**	500 kW
Liquid Flow Rate per Rack	112.5 lpm (29.7 gpm)	**Total Cluster Liquid Flow Rate**	2,250 lpm (594.4 gpm)
Airflow Rate per Rack	3,750 CFM	**Total Cluster Air Flow Rate**	75,000 CFM

* Estimated values, actual numbers and conditions may vary