UNIVERSITÀ DEGLI STUDI DI PADOVA

UNIVERSITA' DEGLI STUDI DI PADOVA

Dipartimento di Ingegneria Industriale DII

Corso di Laurea Magistrale in Energy Engineering

CFD Analysis of a Data Center: Thermal and Airflow Optimization

Relatrice: Prof. Anna Stoppato

Correlatore: Massimiliano Maistro

Ali Charara, 2085781

Anno Accademico 2024/2025.

Acknowledgement

First and foremost, I would like to express my deepest gratitude to God, whose guidance and blessings have illuminated every step of my academic and personal journey.

I am profoundly thankful to my supervisor, Professor Anna Stoppato, for her invaluable guidance, encouragement, and unwavering support throughout my thesis work. Her expertise and dedication have been instrumental in shaping both this research and my development as an engineer.

My sincere thanks also go to my co-supervisor, Massimiliano Maistro, at Vertiv, and to the entire Application Engineering team. I am truly grateful for the warm welcome, continuous support, and inspiring collaboration I experienced during my internship. Their mentorship and readiness to assist at every stage made this journey both enriching and enjoyable.

A heartfelt thank you goes to my parents and siblings, whose love, sacrifices, and belief in me have been my strongest foundation. Your constant presence and encouragement have empowered me to reach this important milestone and to confidently embark on my career path.

Finally, to my fiancé and my partner in life, studies, and every adventure—thank you for being by my side through both the easiest and most challenging days. We have shared not only university lectures but also countless life moments, and without your support, patience, and unwavering faith in us, none of this would have been possible.

To all of you, I owe my deepest gratitude.

Abstract

The increasing computational demands and heat loads in modern data centers create significant challenges for maintaining efficient thermal management and preventing hotspots. This thesis presents a comprehensive Computational Fluid Dynamics (CFD) study using ANSYS Fluent to analyze and optimize airflow, temperature distribution, and cooling strategies in a high-density data center environment. A validated reference model based on the CQ-University data center was replicated and then substantially improved by implementing a cold aisle containment strategy, which reduced maximum rack temperatures by 33.79 % and successfully eliminated hotspot formation.

The research further investigates retrofitting for higher rack density, comparing several rack arrangements and identifying an alternating high/low-density configuration as the most effective for uniform cooling. Using real-world equipment sizing methods (Vertiv's Hi-rating and GRS 2.0 tools), the study then scales up to model a mega data center with a total IT load of 4.48 MW. Both normal and extreme operating conditions are simulated, with results confirming that fine-tuned cooling strategies and targeted airflow management can optimize Power Usage Effectiveness (PUE) from 1.06 to 1.021 and reduce annual energy costs by €374,760 (a 4.75% savings).

The outcomes of this thesis demonstrate the critical value of CFD-guided analysis and literature-backed best practices in designing, validating, and operating thermally efficient data centers. By systematically addressing thermal management challenges, this work offers a robust framework for sustainable, cost-effective operation of high-density digital infrastructure.

Table of Contents

List of Figures:

List of Tables

Abbreviations

- **AI**: Artificial Intelligence
- **ASHRAE**: American Society of Heating, Refrigerating and Air-Conditioning Engineers
- **CAD**: Computer-Aided Drawing
- **CACS**: Cold Aisle Containment System
- **CDU**: Cooling Distribution Unit
- **CFD**: Computational Fluid Dynamics
- **CPU**: Central Processing Unit
- **CRAH**: Computer Room Air Handler
- **CW**: Chilled Water
- **DOE**: Department of Energy
- **EED**: Energy Efficiency Directive
- **GPU**: Graphics Processing Unit
- **HACS**: Hot Aisle Containment System
- **IT**: Information Technology
- **PDU**: Power Distribution Unit
- **PUE**: Power Usage Effectiveness
- **RCI**: Rack Cooling Index
- **Re**: Reynolds Number
- **RHI**: Return Heat Index
- **SHI**: Supply Heat Index
- **UPS**: Uninterruptible Power Supply

- **CUE:** Carbon Usage Effectiveness

Symbols and Variables

- **ρ**: Density
- **t**: Time
- **d**: Diameter
- **μ**: Viscosity
- **x, y, z**: Coordinates
- **u**: Velocity Component
- $\vec{V}$: Volumetric Flow Rate
- **P**: Pressure
- G_k: Turbulence kinetic energy generated due to mean velocity gradients
- G_b: Turbulence kinetic energy generated due to buoyancy
- Y_M: Fluctuating dilatation in compressible turbulence to the total dissipation rate
- $C_2, C_{1\varepsilon}$: Constants
- σ_k: Turbulent Prandtl number for k
- σ_ε: Turbulent Prandtl number for ε
- S_k, S_ε: User-specified source terms

Introduction

The rapid growth of digital services and the rise of high-density, AI-driven computing have made thermal management a central challenge for modern data centers. As facilities increase in scale and power density, advanced cooling strategies are required to ensure hardware reliability, control costs, and limit environmental impact. While recent literature demonstrates substantial advances in airflow management and cooling technology, there remains a practical need for scenario-based CFD analyses that can inform real-world data center design.

This report is structured in two complementary parts. The first part delivers a comprehensive review of current data center cooling technologies and CFD simulation approaches, outlining key findings from recent research and highlighting the ongoing evolution of industry challenges.

The second part presents a systematic CFD-based investigation, progressing through a series of targeted simulation scenarios:

- Initial model validation is conducted by replicating a baseline data center configuration and implementing cold aisle containment to establish simulation accuracy.
- Next, a detailed study of alternative rack arrangements is performed, with the objective of identifying the optimal configuration for effective heat dissipation and maintaining maximum rack inlet temperatures below safe limits.
- Building on these findings, the validated best arrangement is then scaled up and applied to a full-scale mega data center scenario. This comprehensive model is subjected to two distinct operational conditions—normal load and extreme load—to evaluate system resilience and thermal performance.
- Finally, Power Usage Effectiveness (PUE) and economic analysis quantify the energy and cost savings achievable through the proposed cooling and arrangement optimizations.

Each simulation step is designed to build logically upon the previous, culminating in actionable insights for future high-density data center operation. This integrated approach not only grounds the study in current scientific understanding, but also demonstrates a clear, evidence-based pathway from literature review to practical engineering recommendations.

A more detailed examination of related works is presented in the following Literature Review section, which establishes the scientific foundation of this simulation-driven investigation.

Over the last two decades, purpose-built data centers have evolved into critical infrastructure for the global digital economy. They underpin cloud services, e-commerce, social-media platforms and the rapidly expanding universe of artificial-intelligence (AI) application sectors in which uninterrupted availability, low latency and data integrity are paramount (Shehabi et al., 2024). Each facility consolidates thousands of servers, network switches and storage array whose silicon devices convert nearly all electrical power into heat. Controlling this heat load is essential to preserve hardware reliability, curb operating expenditure and limit the sector's environmental footprint (ASHRAE TC 9.9, 2022).

Advancements in semiconductor technology have resulted in higher power densities within chips. For instance, the transition to three-dimensional (3D) integrated circuits, which involve stacking multiple layers of active components, has intensified thermal challenges. Studies indicate that 3D stacking can increase the maximum temperature by approximately 17 Kelvin compared to planar integrated circuits, underscoring the need for effective thermal management (Kingma & Welling, 2014).

Scaling effects at the chip level propagate upward to the facility's total energy demand. The 2024 United States Data Center Energy Usage Report estimates that U.S. data centers consumed 176 TWh of electricity in 2023—about 4.4 % of national demand—and projects a rise to 325–580 TWh (6.7–12.0 %) by 2028 under plausible growth scenarios for AI-accelerated servers (Shehabi et al., 2024). Similar upward trends are reported globally; the International Energy Agency forecasts that worldwide data-center electricity use could double between 2022 and 2030 if present trajectories continue (IEA, 2024).

Within the facility energy budget, cooling systems routinely account for up to 40 % of total consumption, second only to the IT load itself. Inefficient airflow—characterized by cold-air bypass and hot-air recirculation—creates *hotspots*, i.e., local regions where server-inlet temperatures exceed the safe operating envelope (ASHRAE TC 9.9, 2022). Hotspots not only jeopardize hardware reliability but also force operators to overprovision cooling capacity, increasing capital expense and power-usage effectiveness (PUE). Advanced computational-fluid-dynamics (CFD) modelling, coupled with dynamic control of computer-room air handlers

(CRAHs) and liquid-cooling schemes, is therefore emerging as a primary strategy for mitigating thermal risk while meeting aggressive sustainability targets (Bash et al., 2022).

Forward-looking projections suggest that data-center electricity demand after 2023 could diverge widely, largely because of variations in GPU-server deployment, utilization levels, and cooling-system efficiency (Shehabi et al., 2024). The U.S. Department of Energy likewise notes that improving cooling performance is critical for keeping future energy growth in check. Figure 1—adapted from the 2024 United States Data Center Energy Usage Report—plots the historical trend and the range of post-2023 scenarios (Shehabi et al., 2024).

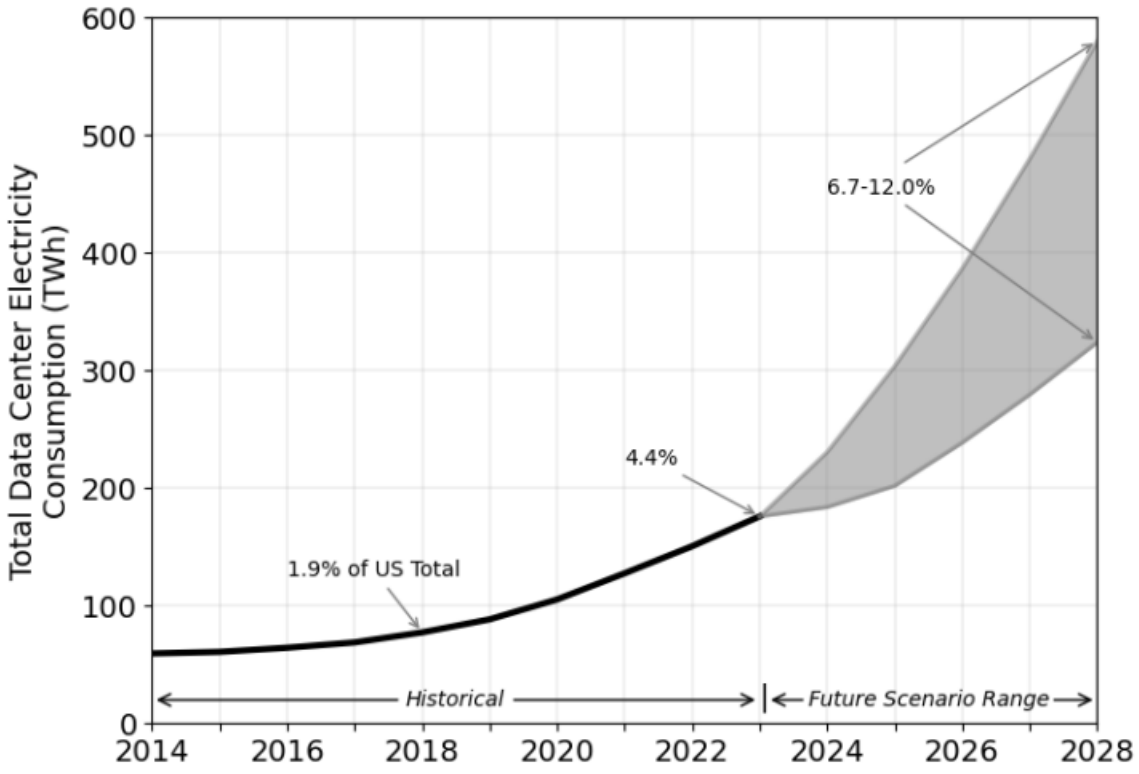

Figure 1. Historical and projected U.S. data-center electricity use (Shehabi et al., 2024)

Figure 1 shows that U.S. data-center electricity use could reach 325–580 TWh by 2028, equal to roughly 6.7–12 % of the nation's total power demand (Shehabi et al., 2024). To explain this projected surge, it is essential to examine how the sector's largest, most energy-intensive companies consume and manage power.

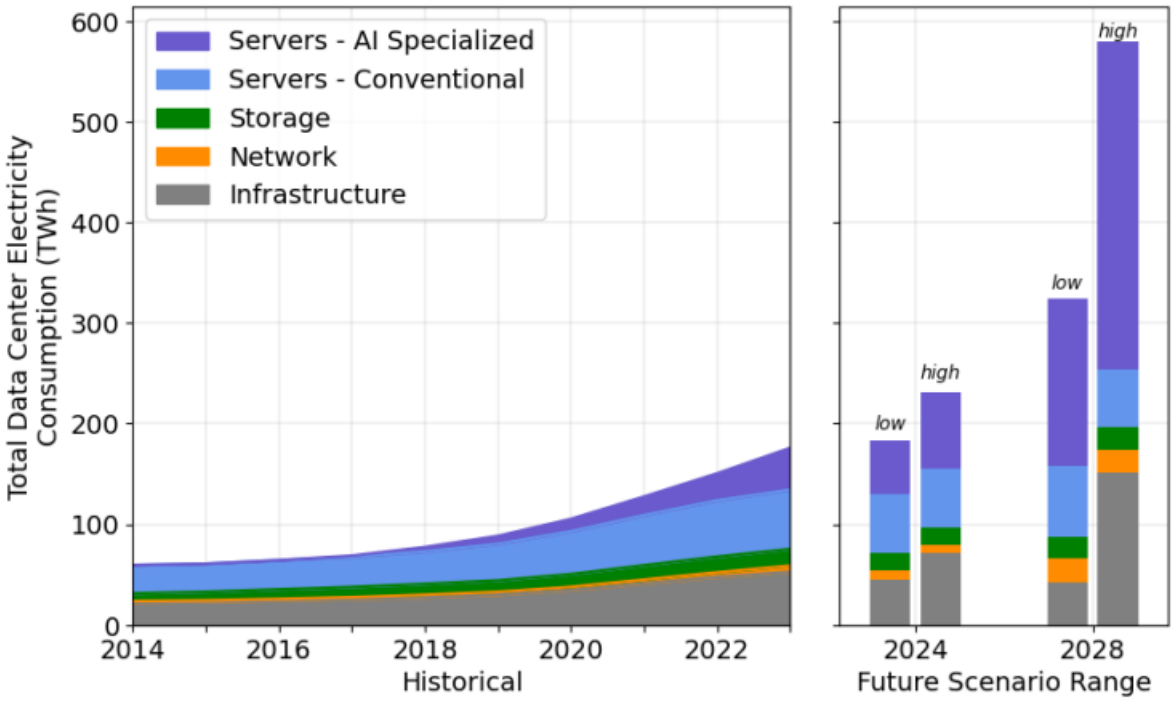

Figure 2. U.S. data-center electricity use by equipment type, (Shehabi et al., 2024)

Figure 2—adapted from the *2024 United States Data Center Energy Usage Report*—breaks total electricity use down by equipment class and shows that growth between 2014 and 2023 was fueled by both a rapid influx of AI-optimized servers and the steady expansion of conventional server fleets. Although storage devices continue to draw more power each year, their share of the total begins to shrink as GPU-accelerated servers dominate the load. When Figure 1 and Figure 2 are viewed together, it becomes clear that the trajectory out to 2028 is governed chiefly by how many AI servers are installed and how intensively they run (Shehabi et al., 2024).

Thermally, those servers concentrate heat in small footprints, creating hotspots that threaten reliability and efficiency. Most facilities still rely on air cooling, typically raised-floor or overhead-supply layouts—to manage this heat. To safeguard hardware, ASHRAE recommends maintaining server-inlet temperatures between 20 °C and 25 °C with relative humidity of 40–50 % for Class A1 equipment (ASHRAE TC 9.9, 2022).

Air and Liquid Cooling Basics:

Data-center racks can be cooled in two fundamental ways (ASHRAE, 2012):

- **Air-cooled systems.** Room air is delivered to the front of each rack, flows through the IT hardware, and exits as heated air at the rear. Inside the rack, individual components may still use heat sinks or liquid cold plates, but the rack as a whole rejects its heat load to the surrounding air.
- **Liquid-cooled systems.** A temperate liquid—typically water held above the room-air dew-point—is routed directly to the heat-generating parts of the electronics. The liquid removes heat at the source and transfers it to a heat exchanger (air-to-liquid or liquid-to-liquid) or to an external cooling unit located outside the rack.

Air cooling remains the dominant approach in most data halls. Supply air reaches the equipment inlets via raised-floor plenums, overhead ducts, or local (row-based) fans. Industry best practice arranges racks in a hot-aisle / cold-aisle pattern: server fronts face a shared cold aisle, draw in conditioned air, and discharge exhaust into the adjacent hot aisle. While this layout suits most deployments, certain specialized servers are not optimized for it, and under-floor systems can suffer from unwanted leakage at cable cut-outs unless those openings are sealed (ASHRAE, 2012).

1.1. Data Center Components and Cooling System Overview

A modern data center is a complex environment designed to ensure the reliable operation of IT equipment while maintaining optimal thermal and environmental conditions. The primary components can be grouped into IT infrastructure and supporting facility systems, with special emphasis on the cooling system, which is critical for efficient and safe operation.

Key Data Center Components

1. **IT Racks and Servers**
    - The core of the data center consists of rows of server racks that house computing, storage, and networking equipment. These racks generate significant heat loads that must be managed continuously.
2. **Raised Floor (Access Floor)**

- Many data centers use a raised floor system to distribute cooled air directly to the fronts of server racks. Perforated tiles allow precise delivery of conditioned air to the cold aisle, improving cooling efficiency (see Figure 1).

3. **Cooling Units (CRAH/CRAC)**

   There are two main air-cooling indoor devices used in data centers, which are:

   - **CRAH (Computer Room Air Handler) Units**: Utilize chilled water supplied from external chillers to cool air, which is then circulated within the data center.
   - **CRAC (Computer Room Air Conditioning) Units**: Use direct expansion (DX) cooling, typically with refrigerant, for standalone cooling needs.

   These units are usually positioned at the perimeter of the room or integrated into rows, controlling both temperature and humidity. They are essential for delivering cool air to the IT equipment inlets and for removing heat from space. In this thesis, the CRAH is the employed unit throughout the studies done.

4. **Chilled Water System (Cold-water System)**

   - The backbone of large-scale data center cooling. This system includes:
     - **Indoor Units:** CRAH units that use chilled water to condition air within the data center (highlighted in orange in Figure 2).
     - **Outdoor Chillers:** Large chiller plants located outside the building supply cold water to the CRAH units via insulated piping (highlighted in blue in Figure 2).
     - **Auxiliary Components:** Hydraulic connections, valves, buffer tanks, and pumps maintain water flow, pressure, and temperature throughout the cooling loop.

5. **Airflow Management Features**

   - **Hot Aisle/Cold Aisle Containment:** Physical barriers or enclosures are used to separate cold supply air (delivered to the front of racks) from hot exhaust air

(leaving the rear), significantly increasing cooling efficiency and allowing precise environmental control (see Figure 1).

- **Perforated Tiles and Floor Grilles:** Used in conjunction with the raised floor to direct cool air exactly where it is needed.

1.2. Figures and Placement

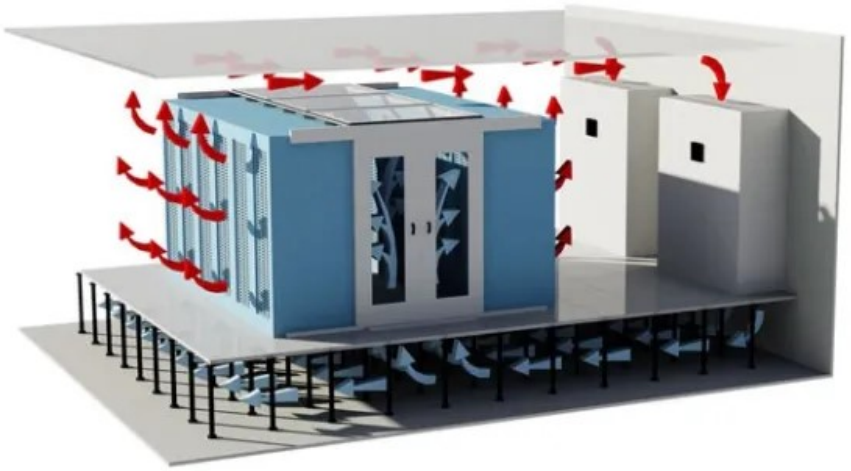

Figure 3. Illustrates the airflow paths in a data center with cold aisle containment, raised floor supply, and overhead return air (ENERGY STAR, 2024).

The blue arrows show cool supply air delivered underfloor and to the rack inlets, while red arrows depict the removal of hot exhaust air back to the cooling units.

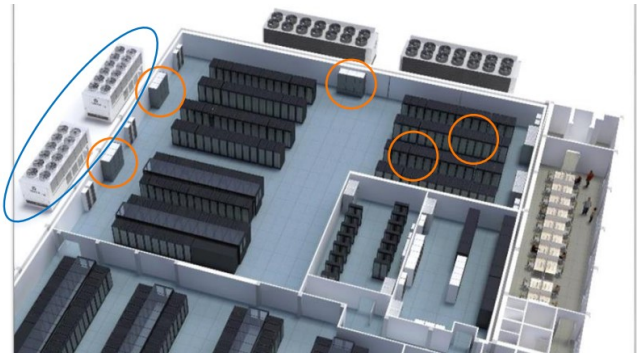

Figure 4. Shows the overall layout of a data center's chilled water system.

Indoor cooling units (CRAHs) are distributed among the racks (orange circles), while outdoor chillers (blue ovals) provide chilled water to the facility. This schematic highlights the integration of both indoor and outdoor components essential for effective data center cooling.

Components in Figure 2

1. **IT Server Racks** – Rows of black cabinets housing computing, storage, and network equipment.
2. **Indoor CRAH Units** – Large white cabinets placed inside the data hall, responsible for cooling the circulating air using chilled water.
3. **Outdoor Chiller Units** – Large white/grey units outside the main building envelope, supplying chilled water to the cooling loop.
4. **Auxiliary Chilled Water System Components** – Include buffer tanks, hydraulic valves, pumps, and piping (usually located near the chillers or in a service area).
5. **Raised Floor** – The blue-tinted floor with visible grid, supporting underfloor air distribution and cable management.
6. **Control Room/Equipment Area** – The small separate room shown at the side of the data hall, typically housing monitoring and electrical equipment.

1.3. Data-center air-cooling configurations

Modern facilities usually distribute conditioned air to the IT racks in one of two ways (Schmidt & Iyengar, 2018):

Table 1. Comparing Data Center Air Supply methods - Raised Floor and Overhead Supply table.

Configuration	Principal airflow path	Notes
Raised floor (under-floor) supply	CRAH units discharge into a pressurized plenum beneath a 600 mm access floor; cold air rises through perforated tiles placed in the cold aisles.	Favored in legacy sites; tile layout strongly influences bypass and recirculation.

Overhead (ceiling) supply	Conditioned air is ducted or released from ceiling diffusers directly above the cold aisles; return air is extracted through ceiling plenums or rear grilles.	Simplifies cable management and allows deeper raised floors to carry power cabling instead.

Both approaches rely on a hot-aisle / cold-aisle rack layout to keep supply and return streams separate. Server fronts (cold-aisles) face the air-supply tiles or diffusers, while the rear exhausts from opposing rows form the hot aisle. Proper containment—using doors, baffles or curtains—reduces mixing and can cut fan energy by 15–20 % (ASHRAE TC 9.9, 2022).

The schematics below highlight the airflow differences:

- Figure 3 shows cold air (blue arrows) rising through perforated tiles, travelling through the servers, and returning as hot air (red arrows) to ceiling-mounted CRAH intakes.
- Figure 4 illustrates an overhead layout where supply diffusers feed the cold aisle from above, while exhaust air returns via a dedicated ceiling plenum**.**

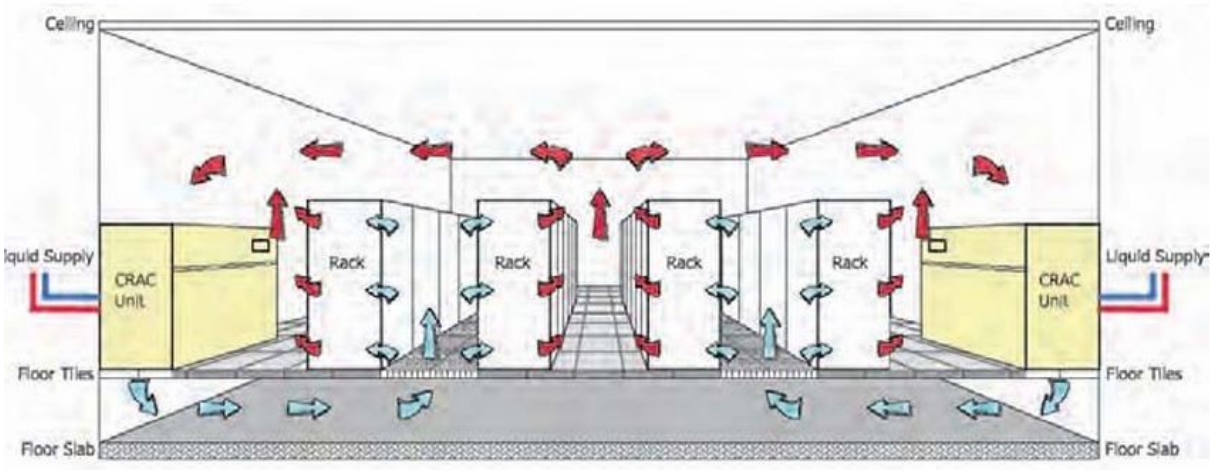

Figure 5. Under-floor air-supply configuration (ASHRAE, 2012).

Figure 5 – Under-Floor Supply

- The cooling unit units sit at the raised-floor level, discharging cold air (blue arrows) into the plenum below.
- Air travels horizontally under the floor, rises through perforated tiles in the cold aisle, moves front-to-back through the racks, and returns as hot air (red arrows) to the CRAH intakes at the top.

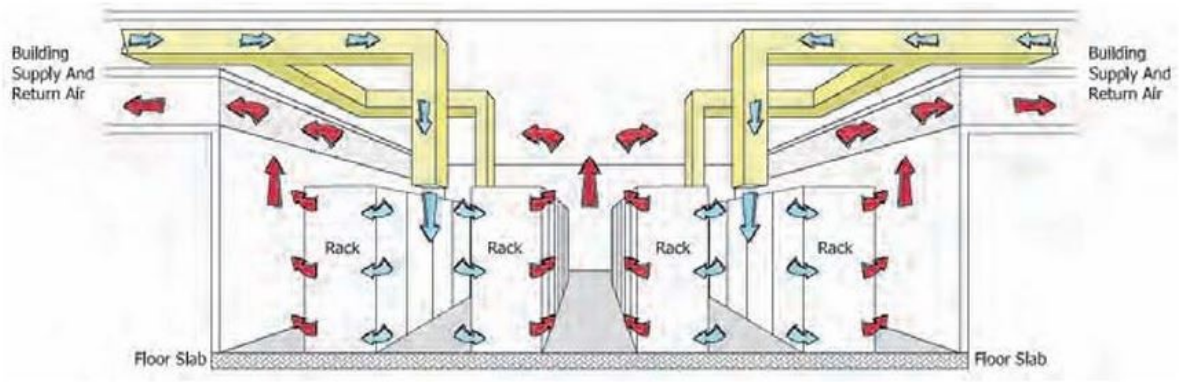

Figure 6. Overhead air-supply configuration (ASHRAE, 2012).

Figure 6 – Overhead Supply

- The cooling units feed conditioned air from above; diffusers drop cold air directly into the cold aisle.
- Exhaust air rises from the hot aisle and returns via a ceiling plenum.

1.4. Hot-Spot Formation and Mitigation in Air-Cooled Data Centers

Proper control of temperature and humidity is essential to data-center reliability, irrespective of the cooling technology in use. The TIA-942 standard recommends arranging racks in an alternating cold-aisle / hot-aisle pattern so that server inlets face one another across the cold aisle, while exhausts face inward across the hot aisle. Perforated tiles in the raised floor of each cold aisle admit conditioned air to the server fronts; the warmed exhaust air rises in the hot aisle

and returns to the cooling units. Under ideal conditions the two airstreams remain strictly separated, but in practice they often mix, creating localized regions of elevated temperature—hot-spots—that jeopardize equipment reliability (ASHRAE, 2012).

1.4.1. Typical causes of hot-spots

A survey of nineteen production computer rooms by the Uptime Institute showed that vertical hot-spots most often arise because lower-rack servers consume the entire supply of conditioned air, leaving upper servers to ingest recirculated exhaust (Sullivan et al., 2008). The same study found that roughly 60 % of the conditioned airflow bypassed the IT load, leaking back to the cooling units through cable cut-outs, floor voids or misplaced perforated tiles. As a result, operators were running between 3.2 and 14.7 times the cooling capacity actually required by the IT heat load, yet still experiencing hotspots. The problem worsens when high-density equipment is installed without re-balancing the airflow.

1.5. Analytical and CFD approaches

Computational Fluid Dynamics (CFD) remains the most powerful and widely adopted method for analyzing airflow and thermal behavior in data centers, especially when the goal is to identify, visualize, and mitigate thermal hotspots in three-dimensional detail. Numerous studies have demonstrated that CFD delivers unrivaled insights into complex airflow patterns, recirculation zones, and temperature distributions, enabling engineers to design advanced cooling strategies and validate layout configurations to comply with demanding standards such as ASHRAE TC 9.9 (Almoli et al., 2012).

While CFD does require careful geometric modeling, boundary condition calibration, and significant computational resources, these investments are justified by the level of accuracy and physical realism it provides. In contrast, empirical and analytical approaches—such as the spreadsheet-based model introduced by Wong, Chan, and Sze (2015)—offer facility operators a lightweight means to estimate hotspot risks based on measured rack surface temperatures and statistical fitting. These models enable quick “what-if” scenario testing and routine operational planning, often within common software platforms like Excel.

However, it is important to emphasize that no empirical or analytical method can fully replace the depth, precision, and predictive capabilities of CFD. While such tools are helpful for monitoring and tactical adjustments, they lack the resolution and physical modeling required for robust design, validation, and optimization, particularly in high-density or retrofitted data centers.

For these reasons, this thesis employs CFD as the core analytical framework. CFD enables comprehensive simulation and optimization of airflow and cooling strategies that are simply unattainable with alternative methods. The resulting analysis not only validates and enhances existing reference models but also provides actionable insights for the development of next-generation, energy-efficient data centers.

1.6. Monitoring and mitigation workflow

A practical, measurement-driven procedure for eliminating hot-spots is as follows:

1. **Instrument the racks.** Install temperature and humidity sensors at several rack elevations, or use infrared thermography to map cabinet surfaces (Balaras & Argiriou, 2002).
2. **Model the thermal profile.** Import the data into a spreadsheet-based model to visualize temperature stratification and quantify hot-spot intensity.
3. **Target the fix.** Where hot-spots are localized, apply spot-cooling measures—blanking panels, brush grommets, airflow baffles or rear-door heat exchangers—instead of raising the overall cooling set-point.
4. **Verify and iterate.** Re-measure after each intervention; refine the model and repeat until inlet temperatures fall within the ASHRAE envelope.

This iterative workflow addresses the root cause of hot-spots and curbs energy use, rather than simply “throwing more cooling” at the entire room.

1.7. Computational Fluid Dynamics (CFD) and Its Impact on Data-Center Thermal Analysis:

Computational Fluid Dynamics (CFD) is a numerical technique that solves the governing conservation equations of mass, momentum and energy on a discretized, three-dimensional grid. Because it reveals detailed airflow paths, temperature fields and pressure distributions, CFD has transformed the way engineers design and operate mission-critical facilities such as data centers (Bash et al., 2015). Instead of relying on "build-and-test" prototypes, operators can now perform virtual experiments—altering rack layouts, tile patterns, fan speeds or containment schemes—and quantify their effect on inlet temperatures, energy use and hot-spot risk before any physical change is made (Schmidt & Iyengar, 2018). Recent CFD studies show, for example, that optimizing the hot-aisle/cold-aisle arrangement or switching supply plenum depth can cut the number of racks operating above 27 °C by more than 60 % (Nakao et al., 2019).

CFD's predictive capability is particularly valuable because most data halls lack windows, introduce little fresh air and are built primarily for equipment rather than occupants. Consequently, the environment must stay within the limits specified by ASHRAE Technical Committee 9.9. The committee defines four classes of data-processing environments, each with steady-state temperature and humidity envelopes designed to protect IT hardware (ASHRAE TC 9.9, 2021). Table 1 summarizes those requirements.

Table 2. Environmental envelopes for data-processing equipment (ASHRAE TC 9.9, 2021).

Class	Air Conditioning	Environmental Control
Class 1	Temperature: 18-24°C	Humidity: 45-50% RH
Class 2	Temperature: 18-27°C	Humidity: 40-60% RH
Class 3	Temperature: 20-27°C	Humidity: 45-60% RH
Class 4	Temperature: 21-27°C	Humidity: 50-60% RH

2. Literature review

The thermal management of data centers has gained significant attention due to the rapid increase of energy demands and the critical need for efficient operational reliability. Efficient airflow and cooling strategies are important to prevent thermal hotspots, which can degrade hardware performance and increase operational costs if unmonitored. This section reviews existing literature on data center cooling challenges, the application of Computational Fluid Dynamics (CFD) in thermal management, strategies for solving thermal hotspots, validation of CFD models, and relevant industry standards.

2.1. Data Center Cooling Challenges

Despite steady advances in facility design, three technical hurdles continue to dominate data-center thermal management: (1) ever-rising rack power densities that concentrate heat in smaller footprints, (2) airflow inefficiencies—cold-air bypass and hot-air recirculation—that create localized hot-spots, and (3) the disproportionate share of site electricity consumed by cooling systems, now averaging 30–45 % of total demand (ASHRAE TC 9.9, 2022). Any cooling strategy or CFD study discussed in this review must therefore be judged on its ability to lower inlet temperatures, cut hot-spot incidence, and reduce energy consumption against these benchmarks.

2.2. Heat Generation and Equipment Density

Average rack densities have climbed from about 5 kW in 2010 to 15 kW today, with high-performance and GPU-accelerated clusters often exceeding 20 kW per rack (Uptime Institute, 2023). At these loads, local heat fluxes can drive processor junction temperatures 10 K above specification unless airflow is carefully managed. The resulting thermal stress shortens component life and triggers protective throttling, eroding application performance (Bash et al., 2015).

2.3. Inefficient Airflow Management

Thermal performance depends mainly on airflow of the cooling unit. Mis-aligned racks, unsealed cable cut-outs and poorly placed perforated tiles allow conditioned air to bypass the IT load or mix with exhaust, cutting cooling effectiveness. Field measurements by Schmidt and Iyengar (2018) reported temperature differentials exceeding 5 °C between the bottoms and tops of

adjacent racks when only 60 % of supplied air reached server inlets. Such stratification forces operators to lower supply-air set-points or install additional computer-room air handlers (CRAHs), both of which inflate energy use. Mitigation options—blanking panels, aisle containment, brush grommets and CFD-guided tile layouts—are examined in Section 2.2.

These bypass and recirculation losses make aisle-layout optimization—especially the use of containment barriers—the decisive lever for reducing hot-spots and fan energy; recent evidence is summarized in the next subsection.

2.4. Containment-Based Aisle Layouts

Aisle containment—whether hot-aisle or cold-aisle—is increasingly recognized as a cost-effective strategy to reduce hot-spots and improve cooling efficiency in data centers. By physically separating cold supply air from hot exhaust air, containment minimizes airflow mixing, directing cooling precisely where it's needed. This enhances thermal uniformity and reduces wasted fan energy.

Recent studies demonstrate that even simple containment retrofits can deliver double-digit gains in both thermal performance and energy savings:

- ENERGY STAR (2019) documented a successful retrofit in a 650 m² data hall:
    - Floor-area hot-spots were reduced by 35%
    - CRAH (Computer Room Air Handler) fan energy consumption decreased by 18%
    - These improvements were achieved without the addition of new cooling infrastructure, highlighting the value of airflow management alone
- ASHRAE Technical Committee 9.9 (2018) reported results from controlled laboratory tests:
    - Installing either hot-aisle or cold-aisle containment resulted in an average ~20% increase in cooling efficiency.

These findings emphasize that optimizing the alignment of racks, perforated floor tiles, and CRAH airflow direction especially when paired with containment, can significantly reduce hot-

spot incidence and cooling power requirements. Importantly, such strategies often require no mechanical upgrades, making them highly attractive for both new and existing data centers

2.5. Dynamic Heat Loads

The dynamic nature of IT workloads in data centers adds a further layer of complexity to thermal management. As computational demand shifts across servers—due to fluctuating user activity, virtualization, and especially machine learning tasks, the associated heat generation varies, creating transient hotspots that traditional static cooling systems struggle to address (Ahmed et al., 2021).

Static cooling systems are typically designed for steady-state conditions, assuming thermal loads are constant and evenly distributed. However, as real workloads change, these systems can cause overcooling in some areas and undercooling in others, leading to wasted energy and possible equipment damage if not properly controlled (Ahmed et al., 2021).

Recent studies highlight the benefits of adaptive cooling systems that respond in real time to varying heat loads. For example, Ahmed et al. (2021) demonstrated that adaptive cooling systems improved power usage effectiveness (PUE) by 20%, reducing it from 1.6 to 1.3 by dynamically adjusting cooling to actual server activity. This improvement underscores the value of responsive cooling for optimizing energy consumption and maintaining thermal stability.

Computational Fluid Dynamics (CFD) plays a crucial role in the development and analysis of these advanced cooling strategies. CFD simulations provide detailed insights into airflow and temperature distribution under variable workloads, helping data center operators predict where hotspots may arise and allowing them to optimize cooling mechanisms accordingly (Liu et al., 2019). For instance, CFD can model how shifting server loads lead to localized temperature spikes and inform more effective cooling solutions.

Looking forward, integrating CFD with machine learning or artificial intelligence (AI) offers promising new directions for thermal management. This approach could enable cooling systems to not only react to present conditions but also anticipate and proactively prevent potential hotspots by learning from historical and predictive data—thus creating more intelligent and efficient data center environments (Ahmed et al., 2021).

.

2.6. Energy Efficiency and Environmental Impact

The environmental footprint of data centers has become a pressing concern as the demand for cloud computing, AI, and digital services continues to surge. Recent estimates indicate that global data centers now account for over 1-2 % of total global electricity consumption, with a significant portion of this energy used by cooling systems (Jones et al., 2022). As these facilities expand, the pursuit of sustainable and energy-efficient cooling solutions has taken on increased urgency.

Several strategies are emerging to reduce both energy use and carbon emissions. Free cooling, which leverages cool ambient air, and liquid cooling, which enhances heat removal from high-density equipment, are gaining traction in modern designs. The integration of renewable energy—such as solar or wind—into cooling infrastructure further mitigates the sector's environmental impact (Geng et al., 2022).

2.7. Case Studies

Case Study 1: Airflow Optimization

A CFD-based investigation by Schmidt et al. (2020) analyzed a 1 MW data center and found that implementing hot-aisle containment reduced thermal hotspots by 35 % and improved energy efficiency by 15 %.

Case Study 2: Dynamic Cooling Systems

Kapadia et al. (2021) deployed adaptive cooling controls in a hyperscale data center, using real-time thermal sensor feedback to dynamically adjust airflow and cooling power in response to shifting workloads. This system improved power usage effectiveness (PUE) from 1.6 to 1.3—a 20 % boost in energy efficiency.

Case Study 3: Environmental Impact

Geng et al. (2022) conducted a life-cycle assessment of data centers integrating evaporative cooling and solar-powered chillers, reporting a 30 % reduction in carbon emissions compared to traditional systems.

Key Challenges Identified

Based on the reviewed literature and case studies, several persistent challenges remain:

1. The need for advanced cooling systems to address high-density heat loads effectively.
2. Inefficiencies in airflow management contribute to the formation of thermal hotspots.
3. The limitations of static cooling systems in adapting to dynamic and unpredictable workloads.
4. Ongoing environmental concerns and the need for more sustainable cooling practices.

This research builds on these findings by using CFD simulations to optimize airflow management, explore dynamic cooling strategies, and identify sustainable solutions tailored to modern data center requirements.

2.8. Computational Fluid Dynamics in Data Centers

Computational Fluid Dynamics (CFD) has become an essential tool in understanding and optimizing thermal management in data centers. By simulating the movement of air, heat generation, and the interaction between various components in the environment, CFD provides valuable insights into the underlying dynamics of cooling systems. The ability to visualize and quantify airflow, temperature gradients, and thermal hotspots helps engineers and operators make informed decisions to optimize design, reduce energy consumption, and improve the reliability of data centers. The most important aspect of the CFD analysis is the ability to simulate the system before performing a real-life test, which also would eliminate safety risks and costs.

CFD has revolutionized the analysis of thermal and airflow behavior by offering a more refined, predictive approach to understanding how air circulates within the space. Traditional methods, such as simple theoretical models or physical testing, often fail to account for the intricate complexities of airflow patterns. Nevertheless, CFD simulations allow for real-time, accurate analysis that can capture variables such as velocity, pressure, temperature, and turbulence, providing a more accurate picture of a data center's cooling performance with several options of modelling.

A major advantage of CFD is the flexibility to explore various turbulence and flow models depending on the level of detail and computational resources available. In ANSYS Fluent and similar platforms, the most commonly used models for data center airflow analysis include:

- **k-ε (k-epsilon) model:** Strikes a balance between computational cost and accuracy, making it suitable for simulating turbulent airflow around racks and equipment in large spaces.
- **Laminar flow model:** Used for scenarios where air movement is slow and smoother in high-density data centers, but sometimes appropriate for micro-scale studies.
- **Reynolds-Averaged Navier–Stokes (RANS) models:** Provide robust, averaged solutions for turbulent flows and are widely adopted in industry CFD simulations.
- **Large Eddy Simulation (LES):** Captures more detailed turbulence effects, but at significantly higher computational cost—typically reserved for research or particularly complex geometries.

By selecting and calibrating these models appropriately, CFD enables the rapid evaluation and optimization of diverse cooling strategies—from basic hot/cold aisle layouts to advanced containment and liquid-cooling systems (Liu et al., 2019).

2.8.1. Airflow Pattern Analysis

This section of the literature review is based primarily on two influential CFD studies:
(1) *Evaluation of Air Management System's Thermal Performance for Superior Cooling Efficiency in High-Density Data Centers* by Cho and Kim (2011), and
(2) *Investigation of Airflow Pattern of a Typical Data Center by CFD Simulation* by Liu et al. (2019).
Both works use Computational Fluid Dynamics (CFD) simulations to examine how different air management layouts impact airflow distribution, cooling effectiveness, and the formation of thermal hotspots in data centers.

Cho and Kim (2011) systematically compared six air distribution system (ADS) layouts—including various overhead and under-floor supply/return strategies—using detailed CFD modeling. Their study introduced two key performance indices:

$$\mathrm{SHI} = \left(\frac{\delta Q}{Q + \delta Q}\right)$$

$$\mathrm{RHI} = \left(\frac{Q}{Q + \delta Q}\right)$$

$$Q = \sum_j \sum_i m^r_{i,j} C_p((T^r_{out})_{i,j} - (T^r_{in})_{i,j})$$

$$\delta Q = \sum_j \sum_i m^r_{i,j} C_p((T^r_{in})_{i,j} - T_{ref})$$

$$\mathrm{SHI} + \mathrm{RHI} = 1$$

- Supply Heat Index (SHI): Quantifies heat gained by cold aisle air before entering racks, with high values signaling undesirable mixing or recirculation.
- Return Heat Index (RHI): Measures the efficiency of heat removal by the return air system, with higher RHI values indicating better exhaust extraction.

RHI and SHI values versus supply and return infrastructures for a data center feature.

Air distribution system (ADS)	ALT-1 (O-CS/CR)[a]	ALT-2 (O-CS/LR)[b]	ALT-3 (O-LS/CR)[c]	ALT-4 (O-LS/LR)[d]	ALT-5 (U-LS/CR)[e]	ALT-6 (U-LS/LR)[f]	Remark
		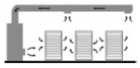	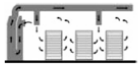	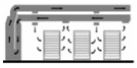	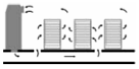	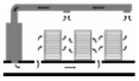	
Environment	Hard floor	Hard floor	Hard floor	Hard floor	Raised floor	Raised floor	Floor type
Supply air	CRAC flooded	CRAC flooded	Locally ducted (Overhead)	Locally ducted (Overhead)	Locally ducted (Under-floor)	Locally ducted (Under-floor)	
Return air	CRAC flooded	Locally ducted	CRAC flooded	Locally ducted	CRAC flooded	Locally ducted	
m	Supply air volume: 45,600 CMH	CRAC [12]					
T_{ref}	Supply air temperature: 13 °C	CRAC [12]					
T^r_{in}	20.7 °C	21.6 °C	21.7 °C	17.5 °C	19.1 °C	17.5 °C	Cold aisle [12]
T^r_{out}	26.0 °C	24.9 °C	27.1 °C	28.2 °C	29.6 °C	29.3 °C	Hot aisle [12]
SHI	0.59	0.72	0.62	0.29	0.37	0.28	Target: 0.0
RHI	0.41	0.28	0.38	0.71	0.63	0.72	Target: 1.0

Figure 7. Schematic of six air distribution layouts (ALT-1 through ALT-6) and summary of SHI and RHI performance indices. Adapted from Cho & Kim (2011).

Key Results:

From Figure 7 (adapted from Cho & Kim, 2011), a clear performance ranking of the six layouts emerges:

- **ALT-1** (CRAC flooded supply and return with hard floor) achieved SHI = 0.59 and RHI = 0.41. This reflects moderate performance, with some mixing.
- **ALT-2** and **ALT-4**, both featuring locally ducted return systems, show improved heat removal:
    - **ALT-2**: SHI = 0.72 (worse than ALT-1), but RHI drops to 0.28—indicating ineffective return air capture despite ducting.
    - **ALT-4**: SHI = 0.29 and RHI = 0.71—among the best results, demonstrating efficient separation and minimal mixing due to locally ducted returns and overhead supply.
- **ALT-3** (locally ducted supply with CRAC return) has a moderate SHI of 0.62 and RHI of 0.38, slightly better than ALT-1.
- **ALT-5** and **ALT-6**, both with raised floors and underfloor supply:
    - **ALT-5** (CRAC return): SHI = 0.37 and RHI = 0.63.
    - **ALT-6** (locally ducted return): SHI = 0.28 and RHI = 0.72—the best overall performance in terms of airflow management.

These results confirm that **ALT-6**, which combines underfloor supply with locally ducted return, offers the most effective configuration by minimizing cold air contamination (lowest SHI) and maximizing heat extraction (highest RHI). In contrast, layouts with CRAC-based return or non-ducted return pathways tend to suffer from greater air mixing, leading to higher SHI and lower RHI values.

Additionally, the temperature data support these trends:

- Inlet temperatures range from 17.5 °C (ALT-4 and ALT-6) to 21.7 °C (ALT-3), correlating inversely with SHI.

- Outlet temperatures rise in layouts with poor return separation, peaking at 29.6–29.3 °C for ALT-5 and ALT-6, as expected for higher thermal loads despite better airflow patterns.

Liu et al. (2019) further confirmed that CFD-driven adjustments to airflow paths—especially tile placement and flow balancing—can greatly improve performance even in existing layouts. Their findings reinforce the role of simulation in both diagnosing issues and testing retrofit strategies.

2.9. Design Optimization

CFD is not only useful for testing predefined strategies, but also for holistic design optimization in complex data center environments. By systematically simulating different configurations of racks, cooling units, and perforated tiles, CFD helps engineers identify layouts that maximize cooling efficiency, reduce energy consumption, and prevent the formation of hotspots (Yamada et al., 2021).

In a large-scale study, Yamada et al. (2021) used CFD simulations to optimize both the placement of IT racks and cooling units in a hyperscale data center. Their findings showed that strategic adjustment of rack orientation and cooling-unit location could improve overall cooling efficiency by up to 30%. This kind of design optimization is especially valuable in environments using advanced or hybrid cooling methods, such as liquid-cooled racks combined with traditional air cooling, where airflow and heat removal are tightly coupled.

2.10. Limitations and Challenges in CFD Applications

Despite the significant benefits of CFD, there are limitations to its application in data centers. Boundary conditions, turbulence models, and complex geometries can all impact the accuracy of simulations. As highlighted by Ahmed et al. (2021), accurately modeling airflow and heat distribution in real-world data centers requires high-quality input data, such as precise server layouts and airflow specifications. In many cases, obtaining such detailed data can be difficult, especially in legacy systems that were not initially designed with CFD analysis in mind.

Moreover, CFD models often rely on assumptions, such as steady-state airflow or uniform heat generation, which may not always align with real-world conditions. Dynamic heat loads, transient airflow, and real-time data fluctuations are not always fully captured, which can lead to discrepancies between simulated and actual results. However, with advancements in real-time monitoring and integration of CFD with sensor-based systems, these challenges are increasingly being addressed, allowing for more accurate and actionable simulations.

Computational Fluid Dynamics (CFD) has established itself as an indispensable tool for optimizing cooling systems and understanding airflow patterns in data centers. Through airflow pattern analysis, scenario testing, and design optimization, CFD allows engineers to proactively address thermal challenges and enhance cooling system efficiency. While challenges such as the need for high-quality input data and dynamic workload modeling remain, the continued development of CFD technology and its integration with real-time monitoring systems promises even greater improvements in data center thermal management. As data centers grow in complexity, the role of CFD in optimizing performance and energy efficiency will continue to be a cornerstone of modern data center design.

2.11. Validation Process of CFD Models

Validation is essential to ensure that CFD results accurately represent real-world conditions. This is typically achieved by comparing simulation outputs with experimental data or operational measurements from data centers (Almoli et al., 2012).

The accuracy and reliability of Computational Fluid Dynamics (CFD) models are critical for generating meaningful insights in data center thermal management. Validation involves comparing the CFD model's predictions with measured values of airflow, temperature, and heat dissipation to assess its accuracy, identify discrepancies, and refine simulation assumptions. Well-validated CFD models increase confidence in design choices and operational decisions based on simulation results.

2.11.1. Importance of Validation

Validation ensures that CFD models reflect the actual thermal and airflow behavior within a data center. Without validation, simulations may produce misleading results, potentially leading to inefficient or ineffective design choices. Key benefits of validation include:

- Improved accuracy by identifying and correcting errors in turbulence models, boundary conditions, and input parameters.
- Enhanced credibility of simulation-based recommendations.
- The ability to calibrate and refine models for future scenario testing.

Recent literature demonstrates that validated CFD models can reliably predict rack inlet temperatures and airflow rates within a margin of 1–3°C compared to experimental measurements (Almoli et al., 2012). As the complexity of data center environments increases, ongoing validation is increasingly important for ensuring robust and actionable simulation outcomes.

2.11.2. Methods of Validation

The three main methods to validate a CFD model according to literature are:

- **Experimental Measurements**
  Validation often begins with direct comparison of CFD results to experimental measurements obtained from physical setups. In data centers, this typically involves recording temperature, velocity, and pressure at multiple points throughout the facility. For example, Ahmed et al. (2021) validated a CFD model of a raised-floor data center by measuring temperature profiles at different heights within the cold and hot aisles, reporting discrepancies of less than 5% between simulations and actual data.
- **Case Studies and Real-World Data**
  An alternative approach uses operational data from live data centers, where integrated sensors continuously monitor temperature, humidity, and airflow. In one study, Tang et al. (2020) validated their CFD model using real-time temperature readings from a hyperscale facility. By calibrating boundary

conditions and turbulence models to match observed data, their simulation achieved a 97% correlation with measured performance.

- **Comparison with Industry Guidelines**
  CFD models are also validated against established benchmarks, such as ASHRAE standards for temperature ranges, airflow rates, and thermal performance in data centers. Ensuring that simulation outputs fall within these recommended ranges helps confirm both compliance and reliability.

Validation is usually an iterative process, with discrepancies between simulation and measurement used to further refine mesh density, boundary conditions, turbulence models, or input data.

2.11.3. Challenges in Validation

Validation is not without challenges, including:

- Real-world data centers exhibit dynamic conditions, including varying workloads, fluctuating heat loads, and transient airflow patterns, making it difficult to capture every variable accurately in CFD models.
- Obtaining precise and comprehensive experimental data can be time-consuming and expensive. In some cases, sensors may be unable to measure certain parameters with the required accuracy or resolution.
- Boundary Condition Assumptions: Simplifications or assumptions made during model setup, such as uniform heat loads or steady-state conditions, may not align with real-world dynamics.

2.11.4. Advancements in Validation Techniques

Recent advancements have helped to overcome these challenges and improve CFD reliability:

- Integration with Machine Learning: AI models can be used to analyze and correct discrepancies between CFD predictions and experimental data, refining simulation parameters and improving accuracy (Kapadia et al., 2021).

- Sensor-Driven Real-Time Validation: Enhanced monitoring systems allow for continuous calibration and validation of CFD models using live data, increasing trust in simulation-based decision making.

What should be concluded about the validation process is that it is an indispensable step in CFD modeling for data centers, ensuring that simulations accurately reflect real-world conditions. By integrating experimental data, real-world observations, and iterative refinement, validated CFD models become powerful tools for optimizing thermal performance, designing effective cooling strategies, and preventing thermal hotspots. Advances in sensor technology and machine learning are poised to further enhance the validation process.

2.12. Industry Standards and Guidelines

Industry standards and guidelines play a pivotal role in establishing benchmarks for the design, operation, and energy efficiency of data centers. Organizations such as the American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) provide a comprehensive framework for thermal management, helping data center operators meet performance, sustainability, and reliability goals. These standards are widely adopted to enhance energy efficiency, ensure operational consistency, and safeguard IT equipment.

ASHRAE Guidelines

According to (ASHRAE, 2022) specifications both recommended and allowable temperature ranges for data centers, and these guidelines are referenced by equipment manufacturers and facility designers worldwide. Importantly, these temperature limits refer to the air temperature measured directly at the inlet of the IT equipment (i.e., the front of servers and network devices)—not simply the general room temperature, air supplied by cooling units, or air returned to cooling systems.

- **Recommended** Rack Inlet Temperature Range:
  18°C to 27°C
  This range is optimized for balancing IT equipment reliability and energy efficiency, and is considered best practice for most modern data centers.

- **Allowable** Rack Inlet Temperature Range (Class A1 equipment):
  15°C to 32°C
  This wider range is intended to accommodate temporary deviations from the recommended conditions without compromising equipment stability or safety (ASHRAE, 2022).

Energy Efficiency Metrics:

ASHRAE promotes key performance indicators for data center energy management, the most widely adopted being Power Usage Effectiveness (PUE) as a key metric for data center energy efficiency, defined by:

$$PUE = \frac{Total\ facility\ Power}{IT\ equipment\ Power}$$

A lower PUE indicates a more efficient data center, as more of the facility's power is used directly by IT equipment rather than by support systems like cooling and lighting (ASHRAE, 2022).

The total facility power is made up of the IT power in addition to two other components: Miscellaneous Power and Cooling, below the subdivisions of each are listed:

- IT power includes:
  - Servers, storage, telco, and other related elements.
- Miscellaneous Power:
  - Batteries, transfer switch, UPS, PDU's, Transformers, Rack distribution units, Braker Panels, Distributed Wiring, lighting, generators.
- Cooling:
  - Chillers, chilled water pumps, cooling towers, condenser Pumps, Dry coolers, CRAHs, Fans, humidifiers, etc.

Impact of ASHRAE Guidelines on Energy Efficiency

Numerous studies and real-world applications demonstrate that adherence to ASHRAE guidelines significantly improves the energy efficiency and operational stability of data centers. For example, Geng et al. (2022) found that implementing ASHRAE's recommended temperature

ranges and airflow containment strategies in a hyperscale data center led to a 20% reduction in cooling energy consumption. Leading technology companies, including Google and Microsoft, have adopted ASHRAE standards to optimize their cooling systems, reporting substantial reductions in operational costs and enhanced sustainability. Moreover, manufacturers such as Vertiv have developed advanced cooling systems—incorporating adiabatic free cooling and supplemental systems—explicitly designed to align with ASHRAE recommendations, further contributing to lower Power Usage Effectiveness (PUE) and reduced energy use.

Other Relevant Standards and Guidelines

While ASHRAE provides the primary reference for thermal management, several other standards are widely adopted to ensure reliability, energy efficiency, and regulatory compliance in data centers.

ISO 50001

ISO 50001 is the globally recognized standard for energy management systems (EnMS), designed to help organizations improve their energy performance in a structured and continuous manner. Applicable across all sectors, including data centers, the standard provides a systematic framework for establishing energy baselines, setting performance targets, and implementing operational controls that promote efficient energy use (International Organization for Standardization [ISO], 2018).

Unlike ASHRAE standards, which focus on specific technical and thermal requirements (e.g., temperature and humidity control), ISO 50001 adopts a holistic view of energy management. It encourages organizations to integrate energy considerations into their overall management practices, including procurement, design, and daily operations. For data centers, this means that energy-intensive systems—such as cooling, power distribution, and IT equipment—are monitored and optimized not in isolation, but as part of a broader energy strategy.

By aligning with ISO 50001, data centers can not only improve energy efficiency and reduce operating costs, but also demonstrate compliance with environmental and sustainability goals.

Uptime Institute Tier Classification

The Uptime Institute Tier Standard is a globally recognized system for classifying data centers based on their infrastructure's reliability and fault tolerance.

- Tier I: Basic infrastructure, limited redundancy; suitable for non-critical use.
- Tier II: Adds redundancy in key components; improved reliability.
- Tier III: Redundant delivery paths for power and cooling; allows maintenance without downtime.
- Tier IV: Fully fault-tolerant with independent, isolated systems; ensures maximum availability (Uptime Institute, 2022).

These tiers help data centers align their thermal and cooling systems with desired uptime requirements. For example, Rahman (2021) found that combining Tier III infrastructure with ASHRAE thermal guidelines enabled mission-critical data centers to maintain high energy efficiency while achieving 99.98% uptime.

European Code of Conduct for Data Centers

The European Code of Conduct for Data Centers, developed by the European Commission, outlines best practices for energy efficiency tailored to the EU context. This code complements ASHRAE guidelines by emphasizing thermal efficiency through strategies such as improved airflow management and optimized cooling configurations. Adoption of the Code has led to energy reductions of up to 30% in participating data centers (European Commission, 2022).

3. Methodology

3.1. Methodology Introduction:

The methodology adopted in this study is structured to systematically analyze and optimize the thermal performance of a high-density data center using Computational Fluid Dynamics (CFD) simulations. All simulations were carried out using ANSYS Fluent, with supporting tools from Vertiv (such as NewHirating and GRS) for cooling load calculations and equipment data integration. The primary objective is to evaluate the impact of insulated cold-aisle containment and other retrofit strategies on hotspot mitigation and energy efficiency, using a data center configuration validated against established literature.

The approach consists of several main phases:

- Developing a detailed geometric model of the data center based on a published and validated reference configuration,
- Mesh generation,
- Setting appropriate boundary and initial conditions that reflect realistic operational scenarios,
- Validating the CFD model by comparing simulation results with published data,
- Simulating the effects of insulated cold-aisle containment and additional retrofitting scenarios, and
- Analyzing performance outcomes using relevant energy and thermal efficiency metrics.

Each step of the modeling process, from mesh generation to scenario analysis, is described in the following subsections.

CFD Platform and Supporting Tools

All computational simulations in this study were performed using ANSYS Fluent, a widely adopted commercial CFD package capable of modeling complex airflow and heat transfer phenomena in data center environments. ANSYS Fluent was selected due to its robust turbulence models, flexibility in handling large computational domains, and extensive validation in the literature for similar thermal management applications (ANSYS, 2014).

Mesh generation and domain discretization were conducted within the ANSYS Workbench environment, allowing for high control over mesh quality, refinement, and zone definitions.

In addition to Fluent, several supporting tools provided by Vertiv were used:

- Hi-rating was employed for preliminary cooling load calculations, enabling accurate specification of heat input and initial sizing of cooling systems.
- GRS software supported equipment selection and liquid cooling load calculations, ensuring that the simulated environment matched real-world operational data.
- Manufacturer-supplied data supplied from Vertiv and Dell was used to specify equipment parameters, cooling unit (CRAH) characteristics, and detailed server heat loads.

Data and boundary condition inputs from Vertiv's tools were integrated into the Fluent simulation environment through standardized import formats, ensuring consistency between real-world specifications and simulation parameters.

All post-processing, including airflow visualization, temperature distribution plots, and calculation of performance metrics (such as PUE), was also carried out within ANSYS Fluent or exported to Excel as required for further analysis.

Model Geometry and Setup

The geometric model developed for this study is directly based on the validated data center described by Hassan et al. (2013), which represents the CQ-University data center at Rockhampton, Australia. This reference case has been extensively studied and experimentally monitored, providing reliable benchmark data for both temperature and airflow distributions.

Domain and Components:

- **Room Dimensions:** 8.05 m × 7.08 m × 2.96 m
- **Raised floor plenum height:** 0.45 m
- **Number of CRAC units:** 3 (downflow type), supplying a total airflow of 11.7 m^3/s with a combined cooling capacity of 42 kW
- **Server racks:** 2 rows, 5 racks per row (10 total), each rack with a heat load of 3.46 kW and airflow rate of 0.23 m^3/s

- **Perforated tiles:** 15 tiles arranged in 3 rows, each with 56% open area; airflow per tile averaged 0.78 m^3/s
- **Hot-aisle/cold-aisle layout:** Racks are placed to create alternating cold and hot aisles, as illustrated in Figure X
- **Other:** Room returns for hot air, no supplemental cooling; all major thermal and airflow parameters set to match the literature reference.

Modeling Approach:

- The domain was constructed in ANSYS Design-Modeler based on the exact dimensions and layout from Hassan et al. (2013).
- Server racks were represented as uniform heat sources, with individual airflow and power settings matching the reference study.
- Perforated tiles and raised floor details were incorporated as boundary surfaces with specified flow rates and open areas.
- CRAH units were modeled with prescribed supply temperatures and airflow to the underfloor plenum.

Assumptions:

- The effect of cable trays and minor obstructions was neglected for computational simplicity.
- Uniform inlet velocity and temperature were set at vent tiles, based on the measured averages reported in the reference.
- The model operates under steady-state conditions to match the setup of Hassan et al. (2013).

A schematic of the modeled data center geometry is shown in the figure below.

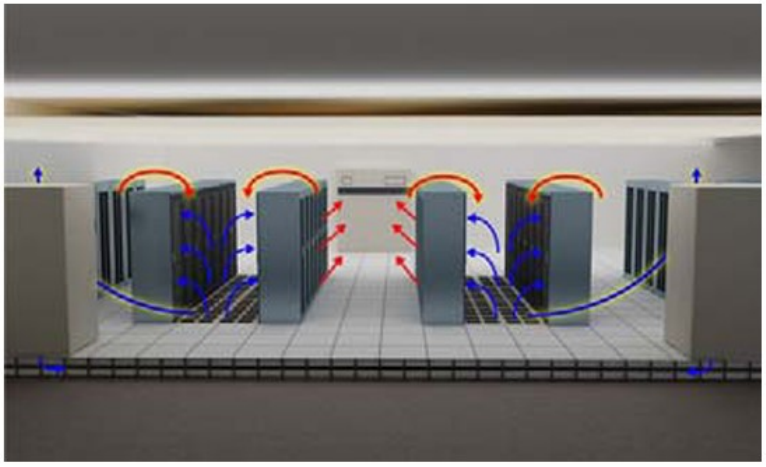

Figure 8. Schematic of the modeled CQ-University data center (adapted from Hassan et al., 2013).

Figure 8 presents the rack arrangement, cold and hot aisle layout, raised floor, and CRAC unit positions of the CQ-University data center in Australia, as modeled in the reference study. This configuration is used as a baseline for validation in the present thesis.

3.2. Mesh Generation and Quality Assessment

The accuracy of CFD simulations relies heavily on the quality and structure of the computational mesh. In this study, the data center domain was discretized using an unstructured tetrahedral mesh generated with ANSYS Meshing (ANSYS 2023 R2). Tetrahedral elements were chosen due to the complex internal geometries of the data center—especially around racks, perforated tiles, and the insulated cold aisle—where structured hexahedral meshing would have been impractical and time-consuming (Sadrehaghighi, 2020). Tetrahedral meshes enable flexible refinement in intricate zones, providing a good balance between computational efficiency and solution accuracy.

Mesh refinement was applied specifically in regions of high velocity and temperature gradients to ensure precise resolution of airflow and thermal patterns. Coarser mesh was used in less critical areas to optimize computational resources. The final mesh contained approximately 135,249 tetrahedral elements, generated with a characteristic element size of 0.2 m. This level of discretization was selected to capture all key thermal and airflow features while maintaining manageable computational times.

Mesh quality was quantitatively assessed using three standard metrics: skewness, orthogonal quality, and aspect ratio. The results are summarized in the Mesh Metrics table below. Mesh quality was quantitatively assessed using three standard metrics: skewness, orthogonal quality, and aspect ratio. The results are summarized in the Mesh Metrics table below:

Table 3. Mesh quality metrics for the final CFD model of the data center.

Mesh Metric	**Minimum**	**Maximum**	**Average**	**Standard Deviation**
Skewness	1.31e-10	0.87208	0.22416	0.12749
Orthogonal Quality	0.12792	1.0	0.77316	0.12432
Aspect Ratio	1.0	10.139	1.8478	0.48526

All mesh quality metrics are well within the accepted ranges for engineering CFD:

- Skewness values are below the 0.94 limit for acceptable cell shape,
- Orthogonal quality is consistently above 0.2, and
- Aspect ratios are within recommended values for three-dimensional simulations (ANSYS, 2014; Ozen, 2014).

A representative view of the generated tetrahedral mesh is shown in the Figure 9.

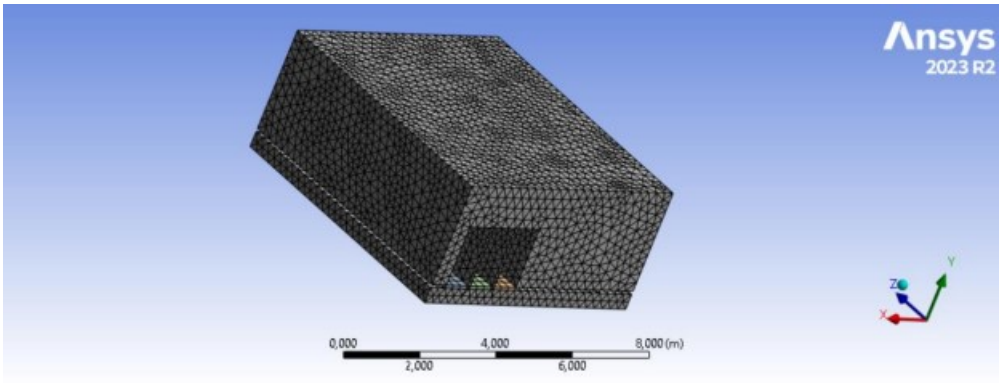

Figure 9. Overall view of the unstructured tetrahedral mesh for the data center domain (ANSYS 2023 R2).

Table 4. Skewness mesh metrics spectrum Ozen (2014).

Excellent	Very good	Good	Acceptable	Bad	Unacceptable
0-0.25	0.25-0.50	0.50-0.80	0.80-0.94	0.95-0.97	0.98-1.00

Table 5. Othogonal Quality mesh metrics spectrum (Ozen, 2014).

Unacceptable	Bad	Acceptable	Good	Very good	Excellent
0-0.001	0.001-0.14	0.15-0.20	0.20-0.69	0.70-0.95	0.95-1.00

2.5 Flow Regime and Reynolds Number

To confirm the nature of airflow in the data center simulations, the Reynolds number (Re) was evaluated. This dimensionless quantity helps classify flow regimes as laminar or turbulent, depending on the relative influence of inertial and viscous forces. It is defined as:

$$Re = \frac{(\rho \times v \times D)}{\mu}$$

Where:

- ρ is the air density (kg/m³),
- v is the characteristic air velocity (m/s),
- D is the characteristic length (m), such as aisle width or tile diameter,
- μ is the dynamic viscosity of air (kg/m·s).

Although the velocity varied across different simulation cases due to changes in operating conditions and airflow configurations, representative calculations showed that all scenarios yielded Reynolds numbers well above the laminar-turbulent threshold of 2100.

For our simulated case, using $\rho = 1.2\, kg/m3, \mu = 1.8 \times 10 - 5\, kg/m.s$, and representative values of velocity and characteristic length shown:

For the base case, the Reynolds number was calculated for the primary cold aisle using:

- $Air\ density, \rho\ =\ 1.2\, kg/m^3$
- $Dynamic\ viscosity, \mu\ =\ 1.8\ \times\ 10^{-5}\, kg/m \cdot s$
- $Mass\ flow\ rate, \dot{m}\ =\ 14.33\, kg/s$
- $Cold\ aisle\ cross-section\ =\ 1.5\, m\ \times\ 2.0\, m\ =\ 3.0\, m^2$
- $Characteristic\ velocity, v = 3.98\, m/s$
- $Characteristic\ length, D = 1.5\, m$

$$Re = \frac{1.2 \times 3.98 \times 1.5}{1.8 \times 10^{-5}} = 398{,}000 > 2{,}100$$

This value is orders of magnitude above the laminar-turbulent transition threshold of 2,100, confirming fully turbulent flow and justifying the use of a turbulence model for all CFD simulations in this study. The flow regime was confirmed to be fully turbulent across all simulated conditions.

As a result, the use of a turbulence model was necessary and consistently applied throughout the CFD analysis.

3.3. Continuity and Governing Equations

The simulations were governed by the fundamental conservation equations of fluid mechanics:

- **Continuity Equation** (mass conservation)

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho V) = 0$$

- **Momentum Equation** (Newton's Second Law for fluids):

$$\frac{\partial (\rho u)}{\partial t} + \nabla \cdot (\rho u V) = -\nabla P + \nabla \cdot \tau + \rho f$$

Here, V is the velocity vector, P is pressure, τ is the viscous stress tensor, and ρ_f represents body forces such as gravity.

Throughout the solution process, the conservation of mass was verified by monitoring the net mass flow at various sections of the model, including the inlet tiles, CRAH outlets, and rack regions. The inflow and outflow mass rates balanced within acceptable margins, indicating that the continuity condition was satisfied. Additionally, the residuals for the continuity equation consistently dropped below 10^{-4}, demonstrating numerical convergence and confirming solution stability.

3.4. Turbulence Modeling

Given the turbulent nature of airflow in all cases, turbulence effects were modeled using the Realizable k–ε turbulence model. This model is widely used for indoor airflow simulations, including data center cooling applications, due to its ability to accurately predict recirculation zones, jet impingement, and separated flow.

The Realizable k–ε model solves two transport equations (ANSYS, 2013):

1. Turbulent kinetic energy (k)

$$\partial/\partial t\,(\rho k) + \partial/\partial x_j\,(\rho k\, u_j) = \partial/\partial x_j\,[\,(\mu + \mu_t/\sigma_k)\,\partial k/\partial x_j\,] + G_k + G_b - \rho\,\varepsilon - Y_M + S_k$$

2. Turbulent dissipation rate (ε)

$$\partial/\partial t\,(\rho\,\varepsilon)\;+\;\partial/\partial x_j\,(\rho\,\varepsilon\,u_j)$$
$$=\partial/\partial x_j\,((\mu\;+\;\mu_t/\sigma_\varepsilon)\,\partial\varepsilon/\partial x_j)+\;\rho\,C_1\;S_\varepsilon-\;\rho\,C_2\,\varepsilon^2/(k\;+\;\sqrt{(\nu\,\varepsilon)})$$
$$+\;C_{1\varepsilon}\,\left(\frac{\varepsilon}{k}\right)C_{\{3\varepsilon\}}\,G_b+\;S_\varepsilon$$
$$where,\qquad C_1\;=\;max\,[\,0.43\,,\eta\,/\,(\eta\;+\;5)\,]$$

$$\eta\;=\;S\;*\;(k\,/\,\varepsilon)$$

$$S\;=\;\sqrt{(2\;*\;S_{\mathrm{ij}}\;*\;S_{\mathrm{ij}})}$$

This model improves upon the Standard k–ε by enhancing prediction accuracy in flows involving rotation and strong streamline curvature, both of which are relevant in data center rack layouts.

The turbulence model parameters used were based on ANSYS Fluent defaults:

- $C_{1\varepsilon} = 1.44$
- $C_2 = 1.9$
- $\sigma_k = 1.0$
- $\sigma_\varepsilon = 1.2$

These values ensured a stable solution and were consistent with best practices for high-performance indoor cooling simulations.

4. Simulation Scenarios and validation

4.1. Overview of the CFD simulations done

This thesis systematically investigates data center thermal management and airflow distribution through a series of Computational Fluid Dynamics (CFD) simulations, presented in detail from Chapters 4 to 7. Each chapter addresses a specific optimization or analysis objective, reflecting a stepwise approach toward advanced data center design.

- Chapter 4: Model Validation
  Section 4.1 outlines the model validation strategy, focusing on replicating a baseline cold aisle containment scenario. The original configuration is simulated and compared with enhanced setups to verify the accuracy of the CFD model. This includes quantifying improvements in cooling performance, reductions in hot air recirculation, and changes in peak temperature, thus establishing a reliable foundation for further analyses.

- Chapter 5: Simulation Results and Testing various rack arrangements to get the best one. This part explores the impact of higher rack density retrofits. Here, various rack arrangements are systematically studied to reflect the increasing demands of modern IT infrastructure. The simulations evaluate different rack placements and airflow rates, aiming to maintain maximum rack inlet temperatures below 45°C. Special emphasis is placed on how adjustments in airflow and the arrangement of high- and low-density racks affect temperature distribution and overall cooling efficiency.
- Chapter 6: Mega Data Center Case—Normal and Extreme Operation
  This part extends the analysis to a full-scale, high-density data center environment, consisting of 16 pairs of 280 kW modules and totaling 4.48 MW of IT load. A hybrid cooling strategy is employed, combining both air and liquid cooling via CRAH units and CDU systems. Simulations are performed under both normal and extreme operational conditions to assess hotspot formation, system resilience, and overall cooling performance. Proper sizing of CRAH and CDU units is validated using Vertiv's Hirating and GRS 2.1 tools.

- Chapter 7: PUE and Cost Analysis
  The technical study is complemented by a thorough economic analysis in Chapter 7, which quantifies the energy and cost savings associated with optimized airflow management and supply temperature control. Improvements in Power Usage Effectiveness (PUE) and financial benefits achieved through advanced thermal management strategies are presented, demonstrating the value of energy-efficient design and operation.

Collectively, these cases underscore the critical role of CFD in data center optimization, providing a robust, quantitative basis for enhancing thermal performance, energy efficiency, and cost-effectiveness in next-generation computing facilities.

4.2. Boundary Conditions and Assumptions

The assumptions and given for this setup :

• **Room dimensions**: 8.0x7.08x2.96 m

• **Floor area**: 56.99 sq-m

• **Raised floor height:** 0.45m

• **Total flow through the CRAH units:** 14.33 kg/s

• **Supply air temperature**: 15 C

• 10 racks with 3 CRAH modeling

• **Total rack heat load**: 41.12 kW

• **Floor tile opening area:** 56%.

• **Cold aisle dimension:** 1.5x2,5x2 m

• **Pressure drop** (in the CRAHs due to filters, coils and internals): 0.5%

• **Flow rate drop between the inlet and outlet return**: 4.776×0.005= =0.0239 kg/s

• **Mass flow rate at the supply of each CRAH unit**: 4.75 kg/s

Table 6. Boundary Conditions naming in Ansys fluent setup.

Boundary Name	**Type**	**Description**
CRAH_Inlet_1, CRAH_Inlet_2, CRAH_Inlet_3	Mass Flow Inlets	Cold air supply from CRAH units at a fixed temperature.
CRAH_Outlet_1, CRAH_Outlet_2, CRAH_Outlet_3	Mass Flow Outlets	Return hot air to CRAH units, with a 0.5% reduction in mass flow due to pressure losses from filters, coils, and internal components.
Rack_Porous_Zone_1 to Rack_Porous_Zone_10	Porous Zone	IT racks modeled as porous media to account for airflow resistance.
Tile_Porous_Zone	Porous Zone	Perforated tiles with a 56% opening for airflow distribution.
Fan_Outlets (R2, R4, R7, R9)	Pressure Jump Outlet	Fans modeled with a pressure jump curve to represent airflow enhancement.
Walls (Room, Racks, CRAH, Floor, Cold Aisle, Ceiling)	No-Slip Wall	Solid boundaries set with no-slip conditions.

Note that the rack heat loads are described in each chapter due to different retrofitted data centers studies, with different parameters used in each section.

4.3. Validation Approach - Baseline Cold Aisle Containment

In this phase of the study, a previously validated research-based data center configuration (Hassan et al., 2013) was replicated and further enhanced using ANSYS Fluent simulations. The original model, representing the CQ-University data center, was recreated precisely to ensure consistency and facilitate reliable validation. A significant modification, the addition of a cold

aisle containment system, was introduced to assess its effectiveness in improving thermal management and cooling efficiency.

The primary objective was to quantify the cooling improvements resulting from the cold aisle containment by evaluating key performance metrics, such as reductions in hot air recirculation and lower peak rack inlet temperatures and comparing these results directly against the original study and baseline (uncontained) scenario.

Model Comparison:

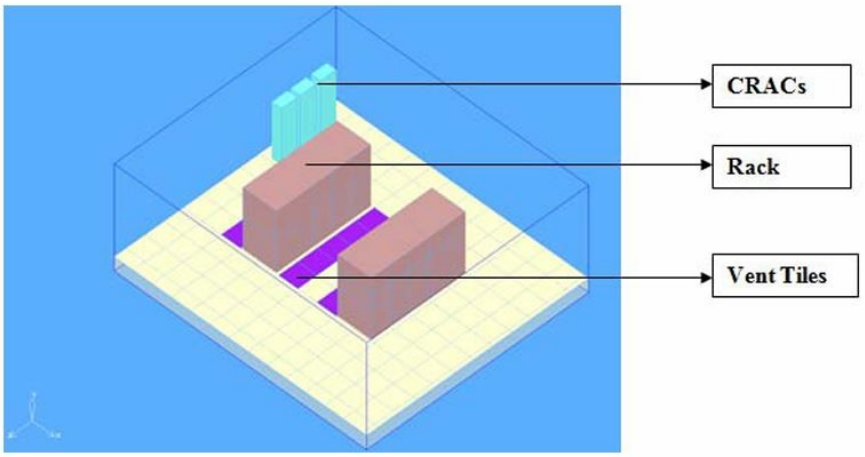

Figure 10. CQ-University Data Center Model missing cold aisle containment (Hassan et al., 2013).

Figure 10 depicts the original CQ-University data center layout as modeled in Hassan et al. (2013). This reference configuration serves as the basis for the current study, which replicates the model and introduces an improvement by adding a cold aisle containment system.

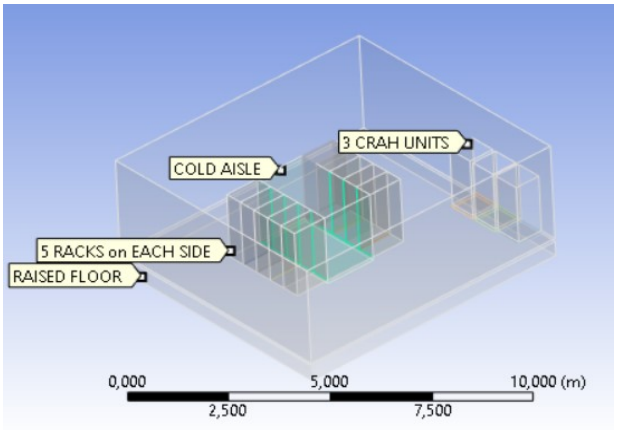

Figure 11. CFD model of the CQ-University data center with added cold aisle containment, developed through this study.

The figure above illustrates the enhanced data center model developed through this study, featuring the implementation of a 2 × 1-meter cold aisle containment structure positioned between the rows of racks. This containment system is specifically designed to fully isolate the supply of cold air from the hot exhaust air leaving the equipment racks. By physically separating the cold aisle from the adjacent hot aisles, the system effectively prevents the mixing of cold and hot air streams, which is a common source of thermal inefficiency in conventional open aisle configurations. As a result, the cold air delivered by the underfloor plenum is directed exclusively to the server inlets, enhancing cooling efficiency and maintaining a more uniform temperature profile across the IT equipment. This modification is expected to improve both thermal management and energy efficiency within the data center environment.

Simulation Parameters and Assumptions:

Table 7. Simulation Parameters and Assumptions with their descriptive values.

Parameter	**Description or Value**
Room Dimensions	8.05 m × 7.08 m × 2.96 m
Floor Area	56.99 m²
Raised Floor Height	0.45 m
Total CRAH Airflow	14.33 kg/s
Supply Air Temperature	15 °C
Rack Configuration	10 racks, arranged in two rows
Rack Heat Load	41.12 kW (total)
Number of CRAH Units	3 units
Floor Tile Opening Area	56%
Cold Aisle Enclosure Dimensions	1.5 m × 2.5 m × 2 m
CRAH Pressure Drop	0.5% (filters, coils, internals)
Flow Rate Drop (Inlet to Return)	0.0239 kg/s
Mass Flow Rate per CRAH Unit	4.7521 s

4.4. Validation Process

The validation involved a two-step process:

1. **Baseline Model Replication:**
    - The original published data center model (without aisle containment) was replicated, matching dimensions, heat loads, and airflow conditions exactly to establish a credible reference point.
2. **Cold Aisle Containment Enhancement:**
    - The enhanced model incorporated insulated containment walls and ceiling around the cold aisle. This structural modification aimed to prevent mixing between the hot and cold air streams, thereby improving airflow distribution and reducing hotspots.

Validation was achieved by benchmarking the outcomes of this enhanced configuration against established data and findings in literature. Specifically, the simulation results were evaluated in terms of maximum temperature reduction and cooling performance improvements. Previous research from the U.S. ENERGY STAR guidelines indicates that adding a cold aisle compartment typically yields a cooling requirement reduction of approximately 10–35%.

Assumptions:

- Uniform heat distribution across racks
- Constant CRAH operating conditions (temperature and airflow)
- Steady-state, incompressible airflow
- Negligible influence of minor obstructions (e.g., cables)

4.5. CFD Simulation and Convergence

To ensure accurate results, the simulation was allowed to run for 500 iterations, achieving convergence of scaled residuals below the threshold (1e-04).

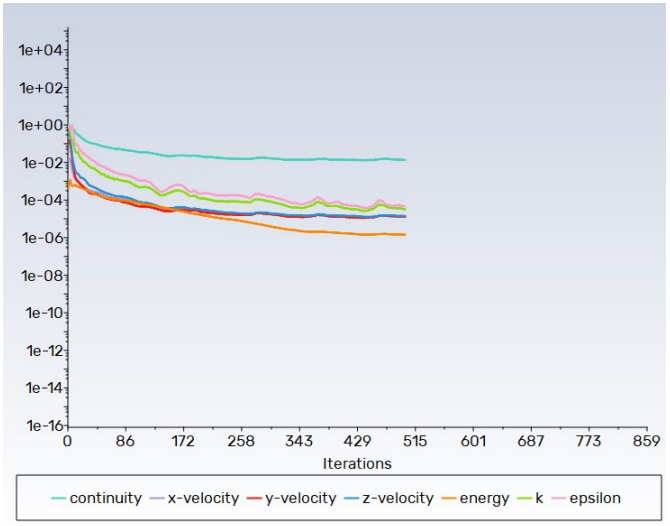

Figure 12. Convergence of scaled residuals for the calculation of the validate baseline case.

4.6. Airflow Analysis

Velocity streamlines obtained from CFD simulations illustrate the airflow pattern within the upgraded data center model. Air supplied from CRAH units flows uniformly through perforated tiles into the contained cold aisle, effectively cooling IT equipment before returning as hot air to the CRAH units. The enhanced containment strategy eliminates hot-air recirculation, significantly enhancing cooling efficiency compared to the original open-aisle model.

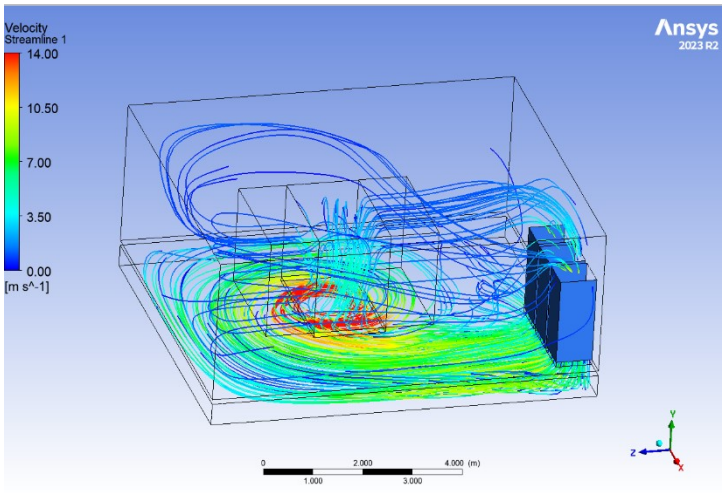

Figure 13. Velocity streamlines of the modeled data center

Comparing the Velocities at the Inlet and Outlet of the CRAH:

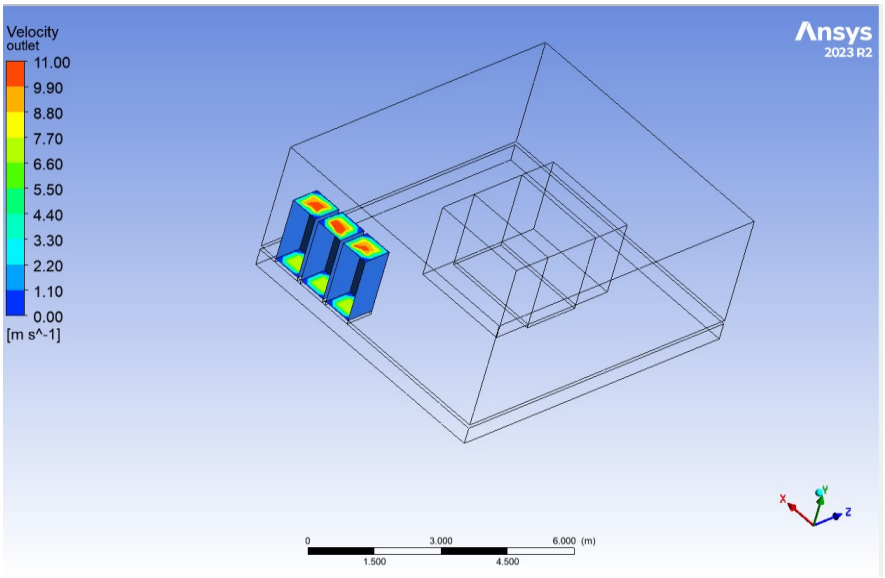

Figure 14. Figure: Velocities at the Inlet and Outlet of the CRAH.

The velocity distribution at the CRAH units provides key insights into airflow behavior within the data center. Higher velocity is observed at the return of the hot air, which is expected due to the flow convergence and the suction effect created by the CRAH units. In contrast, the cold air supply entering the raised floor experiences lower velocities due to the pressure drop occurring across the CRAH unit's internal components, such as filters, coils, and other elements that introduce resistance to the airflow.

4.7. Temperature Distribution and Hot Air Recirculation Analysis

The temperature distribution obtained from the CFD simulation clearly illustrates the effectiveness of the cold aisle containment strategy. The maximum temperature recorded (T_{max}) in the upgraded model was approximately 21°C, significantly lower than the original uncontained setup, where T_{max} reached 29°C.

The below figure, illustrates the temperature contour at 1.5 m height (Y-direction), clearly demonstrating uniform cooling within the cold aisle containment area, effectively minimizing temperature gradients and hotspot formation.

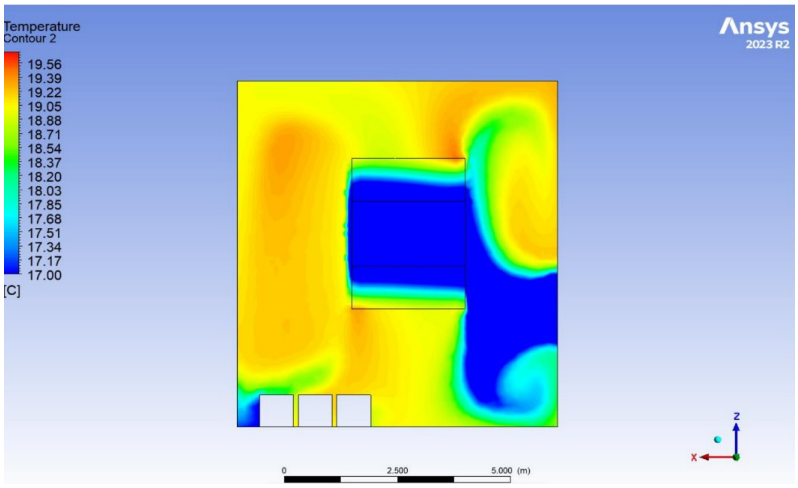

Figure 15. Temperature contour at 1.5 m height in Y-direction.

Moreover, figure 16 shows the mid-section temperature contour in the Z-direction, highlighting effective cooling across the vertical plane of the racks and indicating stable thermal management within the contained cold aisle structure.

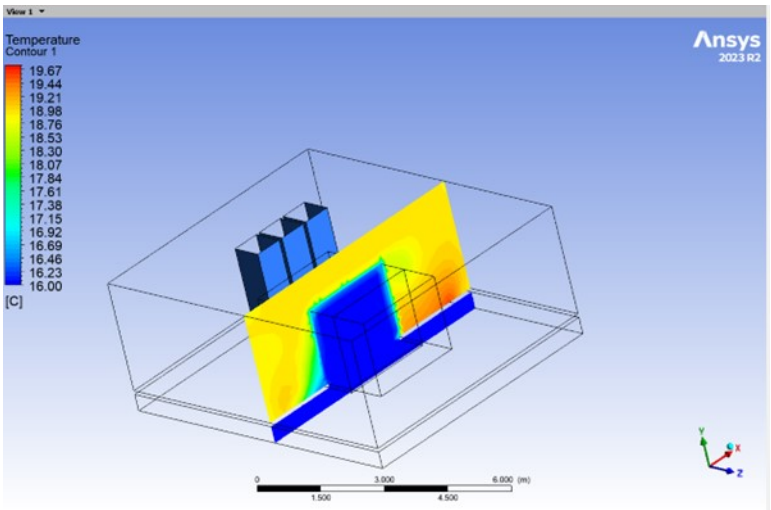

Figure 16. Temperature contour at mid-section in Z-direction.

Similarly, the temperature distribution at the mid-section along the X-direction (Figure 17) emphasizes reduced temperature peaks and significantly lower temperatures at the rack exits,

achieving a maximum temperature of around 19.2°C, well below the original setup's 29°C. This clear reduction reflects improved airflow management and heat removal efficiency.

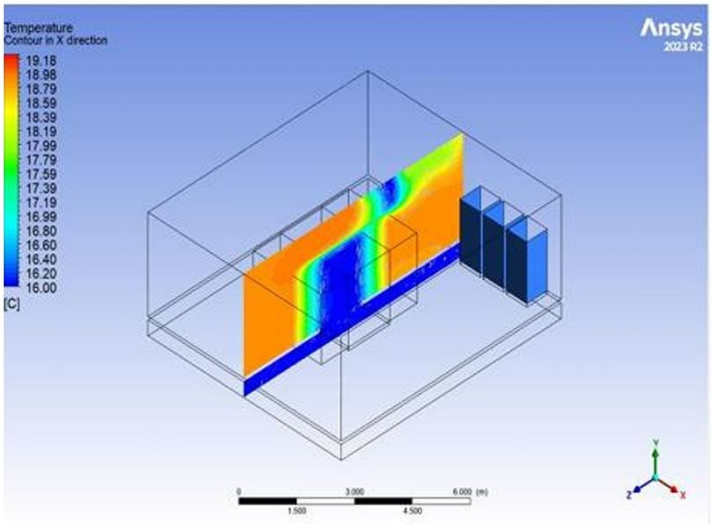

Figure 17. Temperature contour at mid-section in X-direction.

Moving on below in the CFD figures (Figure 18) provide a detailed three-dimensional rendering of temperature distribution, allowing straightforward identification of high-temperature regions. It effectively highlights how the implemented cold aisle containment has substantially reduced temperature hotspots compared to the baseline configuration.

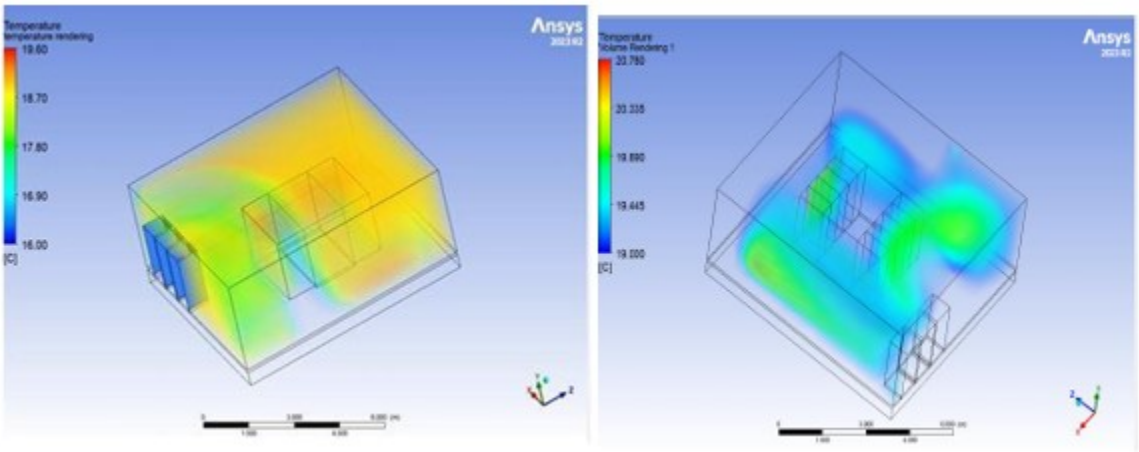

Figure 18. 3D temperature rendering clearly indicating high-temperature zones.

Comparative Analysis of Hot Air Recirculation:

To further illustrate the significant improvement, Figure 19 directly compares the original open-aisle configuration from Hassan et al. (2013) to our modified cold aisle containment design. The baseline configuration clearly demonstrates noticeable hot air recirculation, with substantial mixing between the cold inlet and hot outlet air streams. In contrast, the upgraded setup with containment effectively isolates cold air, virtually eliminating recirculation and significantly enhancing cooling performance.

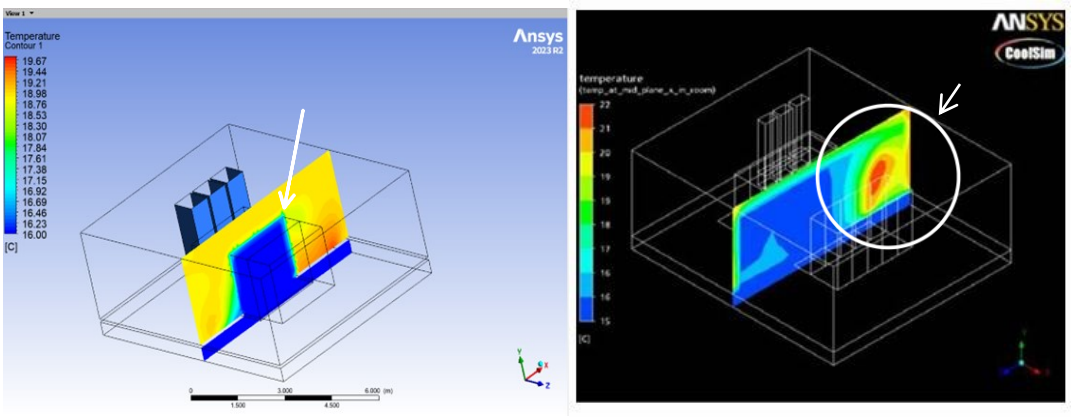

Figure 19. Figure V : Side-by-side comparison of hot air recirculation – enhanced cold-aisle model vs. original CQ-university model's temperature contours.

Side-by-side comparison for figure V, where the left figure is our simulated improved model, and right one is from the source, where within the white circle there is an observable hot hair mixing with the cold air due to the absence of insulated cold aisle between the racks

Higher temperature regions observed in simulation results are predominantly located further from CRAH units, which aligns with expectations due to distance-based airflow efficiency losses. Nevertheless, these temperature values remain significantly below those observed in the original model, validating the effectiveness of the containment strategy in maintaining uniform cooling.

Quantitative Improvement

The implementation of the cold aisle compartment resulted in a **33.79%** reduction in maximum temperature, strongly aligning with industry benchmarks. According to previous studies, notably by ENERGY STAR, incorporating a cold aisle compartment typically achieves a 10–35%

reduction in cooling demands. Our results not only validate these findings but emphasize the practicality and effectiveness of the cold aisle containment approach for substantial thermal performance enhancements in data center environments.

4.8. Pressure Distribution Analysis

As expected, pressure increased slightly within the contained cold aisle environment due to confinement and restricted airflow. These observations align with the literature and are considered acceptable trade-offs for improved cooling performance, shown in the below figure.

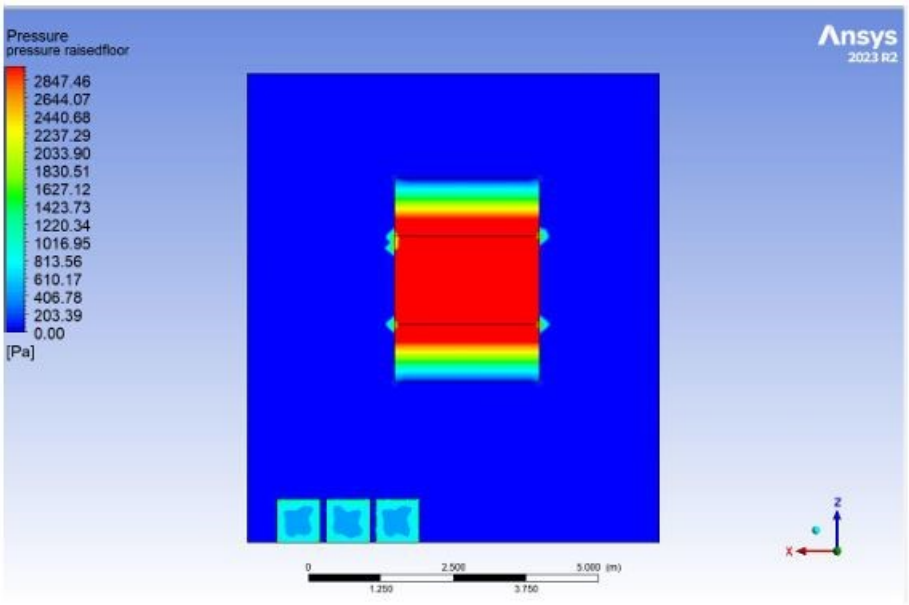

Figure 20. Pressure map of the base case.

.

4.7 Optimization of Energy Efficiency

The primary objective of this optimization phase is to achieve significant reductions in cooling power consumption, while maintaining the target return air temperature at 30 °C in the newly upgraded cold aisle data center. To accomplish this, the study systematically explores and implements targeted operational and design strategies. The optimization process is guided by both simulation results and best-practice recommendations from relevant industry standards. In this section, two key strategies are examined in detail:

1. **Increasing CRAH inlet temperature** (15°C to 18°C)
2. **Reducing total airflow** (14.33 kg/s down to 3 kg/s)

Results for these optimizations were compared with the baseline configuration:

Table 8. Variation of CRAH mass flow rate.

Case	CRAH Inlet T (°C)	Total Mass Flow (kg/s)	Max Temp (°C)
1	16	14.33	20.87
2	16	10.00	21.48
3	16	9.00	22.16
4	16	8.00	23.46
5	16	7.00	25.65
6	16	6.00	28.32
7	16	3.00	29.95 (optimal)
8	16	1.00	48.90

Variation of CRAH inlet temperature:

Table 9. Variation of CRAH inlet temperature, studying the increase in maxium temperature attained in the room.

Case	CRAH Inlet T (°C)	Total Mass Flow (kg/s)	Max Temp (°C)
1	16	14.33	20.87
2	17	14.33	24.80
3	18	14.33	30.06 (optimal)
4	19	14.33	32.05
5	20	14.33	34.02

- **Increasing Inlet CRAH Cooling:** By increasing the inlet CRAH cooling temperature from 15°C to 18°C, we can maintain the same results, implying that the cooling system can handle a higher inlet temperature without sacrificing performance.
- **Decreasing Mass Flow Rate:** Alternatively, reducing the total mass flow rate from 14.33 kg/s to 3 kg/s also achieves the same outcome, suggesting that lowering the mass flow rate can compensate for the increased cooling temperature.

  Both options seem to provide equivalent results for the case study, offering flexibility in how to manage the cooling system with the addition of cold aisle. Both approaches maintain return air temperature around 29–30°C, comparable to the baseline, indicating optimization feasibility.

4.9. Energy and Economic Analysis

To accurately assess the potential energy and cost savings resulting from optimized cooling strategies, the annual energy consumption for each scenario was explicitly calculated using the following assumptions and methodology:

Assumptions

- **Continuous operation:** The data center operates 24 hours per day, 365 days per year (8,760 hours annually).
- **Electricity Price:** The average electricity price for non-household consumers in Italy is taken as €0.1867 per kWh, based on the latest data from the European Commission (June 2024). This value is used consistently throughout the economic analysis.
- **IT Load:** Constant at 41.12 kW for all cases.
- **Baseline cooling power:** 42.01 kW, as determined from CFD simulation and equipment datasheets.
- **Optimizations are based on ENERGY STAR and industry studies:**
    - Raising inlet temperature by 3°C (from 15°C to 18°C) yields a 24.3% reduction in cooling power (based on a typical 4–5% energy saving per 0.56°C/1°F increase).
    - Reducing airflow to 3 kg/s yields a 23% reduction in cooling power.

Calculation Steps

1. **Total Power Consumption (kW):**

$$Total\ Power = IT\ Load\ (kW) + Cooling\ Power\ (kW)$$

2. **Annual Energy Consumption (kWh):**

$$Annual\ Energy\ (kWh) = Total\ Power\ (kW) \times 8{,}760\ hours/year$$

3. **Annual Electricity Cost (€):**

$$Annual\ Cost\ (€) = Annual\ Energy\ (kWh) \times Electricity\ Price\ (€/kWh)$$

4. **Annual Cost Savings (€):**

$$Annual\ Savings\ (€) = Base\ Case\ Cost - Optimized\ Case\ Cost$$

Energy and Cost Calculations

Base Case (Current Configuration):

- IT Load: 41.12 kW
- Cooling Power: 42.01 kW
- Total Power: 83.13 kW
- Annual Energy Consumption: 83.13 × 8,760 = 728,263 kWh
- Annual Cost: 728,263 × €0.1867 = €135,983.01

Case 1: CRAH Inlet Temperature Increased to 18°C

- Cooling Power: 31.81 kW
- Total Power: 41.12 + 31.81 = 72.93 kW
- Annual Energy Consumption: 72.93 × 8,760 = 638,927 kWh
- Annual Cost: 638,927 × €0.1867 = €119,328.63
- Annual Savings: €135,983.01 – €119,328.63 = €16,654.38 (12.24%)

Case 2: CRAH Mass Flow Rate Reduced to 3 kg/s

- Cooling Power: 32.35 kW
- Total Power: 41.12 + 32.35 = 73.47 kW
- Annual Energy Consumption: 73.47 × 8,760 = 643,441 kWh
- Annual Cost: 643,441 × €0.1867 = €120,129.23
- Annual Savings: €135,983.01 – €120,129.23 = €15,853.78 (11.66%)

Table 10. Summary of Energy Consumption Calculations.

Scenario	**Total Power (kW)**	**Annual Energy (kWh)**	**Annual Cost (€)**	**Annual Savings (€)**
Base Case	83.13	728,183	€135,983.01	–
Inlet Temp Increased to 18°C	72.93	638,807	€119,328.63	€16,654.38 (12.27%)
Flow Rate Reduced to 3 kg/s	73.47	643,609	€120,129.23	€15,853.78 (11.62%)

Increasing CRAH inlet temperature to 18°C showed higher savings (12.27%) compared to reducing airflow (11.62%), indicating a slightly preferable option. Both methods substantially reduce annual operating costs compared to the baseline.

The validation study confirms the effectiveness of the cold aisle containment addition, yielding a notable **33.79% improvement** in cooling performance and thermal management efficiency, aligning closely with established literature guidelines. The optimization strategies analyzed—raising inlet temperature and reducing airflow—demonstrated substantial energy savings potential, indicating flexibility in data center cooling operations without sacrificing performance.

5. Impact of Rack Arrangement on Thermal Performance in Retrofitted High-Density Data Centers

Chapter 5 builds upon this foundation by retrofitting the data center to accommodate higher density racks, reflecting the evolving demands of modern IT infrastructure. CFD simulations were employed to examine various rack arrangements and their impact on temperature distribution. The study aimed to determine the optimal rack placement for effective heat dissipation while also identifying the minimum required airflow to maintain a maximum temperature below 45°C. The research explored how adjusting airflow rates and rearranging high- and low-density racks influenced overall cooling efficiency.

5.1. Higher Rack Density Retrofit: Study Overview

As computing demands soar, data centers must adapt to support higher power densities without compromising thermal reliability. In this chapter, we retrofit the previously validated CQ-University model by rearranging and reloading the racks—while keeping the same room geometry and cooling equipment—to reflect a mixed-density environment:

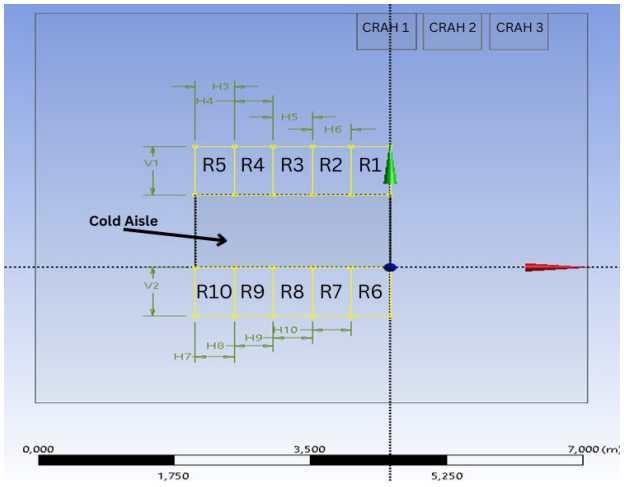

Figure 21. Data center rack arrangement.

- **High-Density Hybrid Racks (R1–R4):** 19.8 kW IT load each, with 15 % served by air cooling (remaining 85% is assumed to be liquid cooled directly, not in the scope of interest for this chapter)

- **Standard Air-Cooled Racks (R5–R10):** 25 kW IT load each, fully cooled by airflow
- **Cooling Infrastructure:** Three CRAH units delivering a total of 14.33 kg/s at 15 °C into the underfloor plenum

This retrofit lets us isolate the effect of internal layout changes on airflow, temperature distribution, and energy use. Our objectives are:

1. **Compare rack arrangements** to find the layout that minimizes hotspots ($T_{max} < 45$ °C).
2. **Determine the lowest mass flow rate** that still prevents overheating.
3. **Balance cooling performance and energy efficiency**, guiding practical retrofit decisions.

We explore three design variations, each simulated until residuals fell below 1×10^{-3}, ensuring fully converged solutions. Results are presented as:

- **Summary Tables** of rack loads, airflow, and peak temperatures
- **Mass Balance Checks** validating continuity at CRAH inlets/outlets and perforated tiles
- **Thermal Contours & Streamlines** highlighting airflow patterns and hotspot locations

5.1.1. Initial Rack Arrangement and Baseline Performance

- **Layout & Loads**

Table 11. Rack type and loads for baseline case.

Rack ID	Type	IT Load (kW)
R1–R4	Hybrid (15 % air-cooled)	19.8
R5–R10	Standard air-cooled	25.0

Notes:

- All three CRAH supply ports deliver 14.33 kg/s at 15 °C to the underfloor plenum.
- Perforated tiles have 56 % open area.

- For the Hybrid 132kW High Density racks – only the air-cooled side bas studied, where the rest of the power is assumed to be taken out by liquid cooling (not studied in this part).

Simulation convergence: Residuals < 1e-3 within 500 iterations.

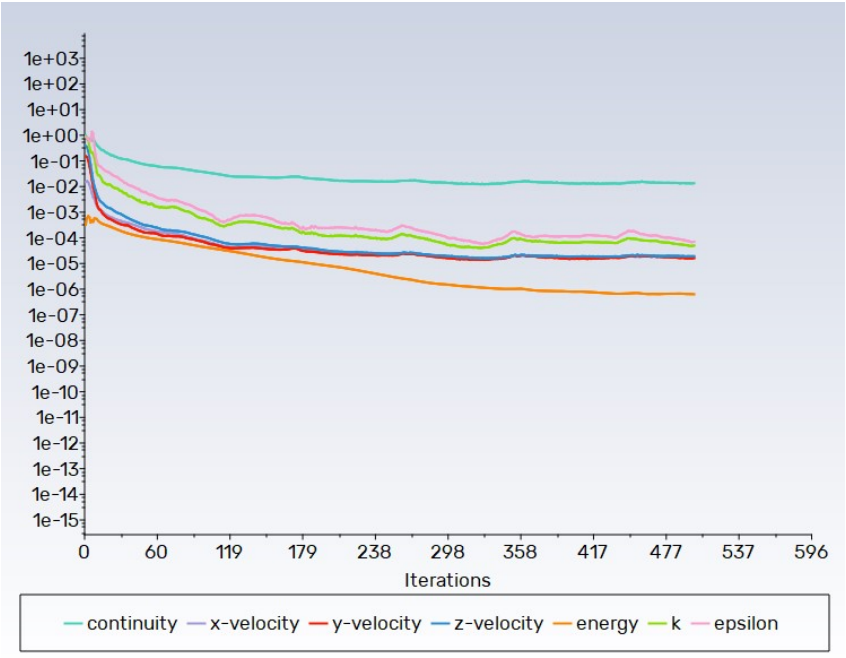

Figure 22. graph of scaled residuals, proving convergence of the solution case.

We ran through fluent a solution consisting of 500 iteration and saw that the scale residuals converged to values lower than 1e-04 suggesting that the solution is fully converged, and no greater number of iterations are needed for this model.

Mass Balance Analysis

Table 12. boundary conditions mass flow rate, direction and description table.

Boundary/Location	**Mass Flow (kg/s)**	**Direction**	**Description**
inlet_crah1	4.75	Inlet	CRAH 1 air supply to plenum
inlet_crah2	4.75	Inlet	CRAH 2 air supply to plenum
inlet_crah3	4.75	Inlet	CRAH 3 air supply to plenum

Total CRAH Inlets	14.25	Inlet	Total supply from CRAHs
interior_inflow_tiles	14.40	Inlet	Air entering through raised floor tiles
interior_outflow_tiles	14.40	Outlet	Air leaving through raised floor tiles
outlet_crah1	4.97	Outlet	Return air from CRAH 1
outlet_crah2	5.07	Outlet	Return air from CRAH 2
outlet_crah3	5.03	Outlet	Return air from CRAH 3
Total CRAH Outlets	15.07	Outlet	Total air returned to CRAHs

Mass Balance Equation:

$$\sum \dot{m}_{in} = \sum \dot{m}_{out}$$

Where:

- $\dot{m}_{in}$ = total mass flow rate entering the control volume (kg/s)
- $\dot{m}_{out}$ = total mass flow rate leaving the control volume (kg/s)

$$\dot{m}_{inlets\ (CRAH)} = \dot{m}_{outlets\ (tiles)} = \dot{m}_{outlets\ (CRAH\ return)}$$

14.25 kg/s (inlet) ≈ 14.40 kg/s (tiles) ≈ 15.07 kg/s (return)

- **Mass Balance at the Inlet and Outlet of CRAH Units:**
  A minor loss in mass flow is observed due to pressure losses across the filters, coils, and internal components of the CRAH units. This effect has been considered in the simulation to ensure accurate modeling of real-world performance.
- **Mass Balance at the Inlet and Outflow of Perforated Tiles:**
  The net difference between the inlet and outflow through the perforated tiles is negligible,

indicating a well-balanced airflow distribution at this checkpoint. This further supports the validity of the CFD model.

All major domain boundaries (CRAH inlets/outlets and raised floor tiles) show very close agreement for inflow and outflow, verifying that the simulation satisfies mass conservation and confirming the validity of your CFD setup.

Thermal study simulation

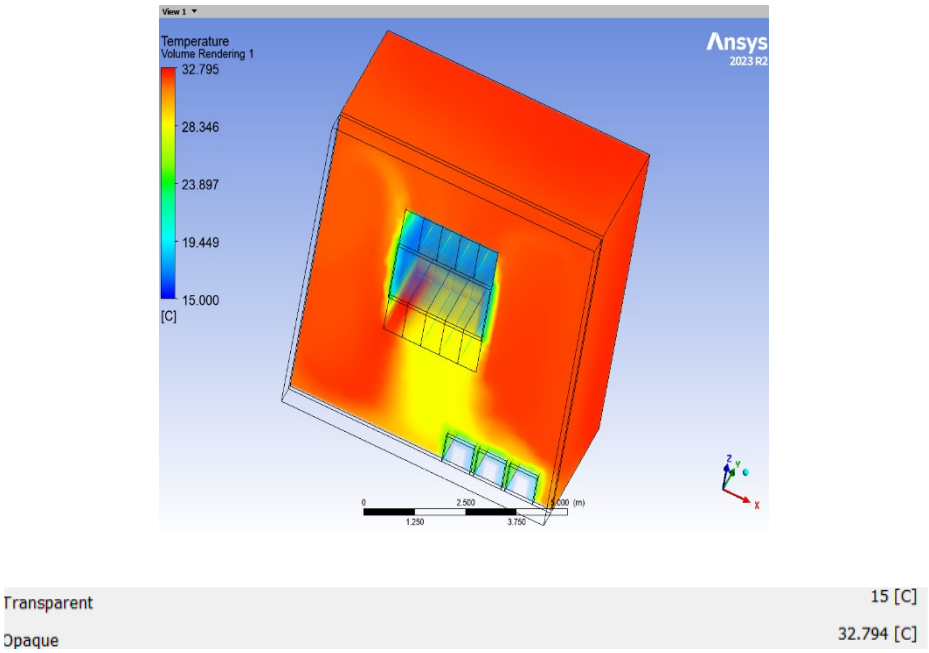

Figure 23. volume thermal rendering of baseline arrangement.

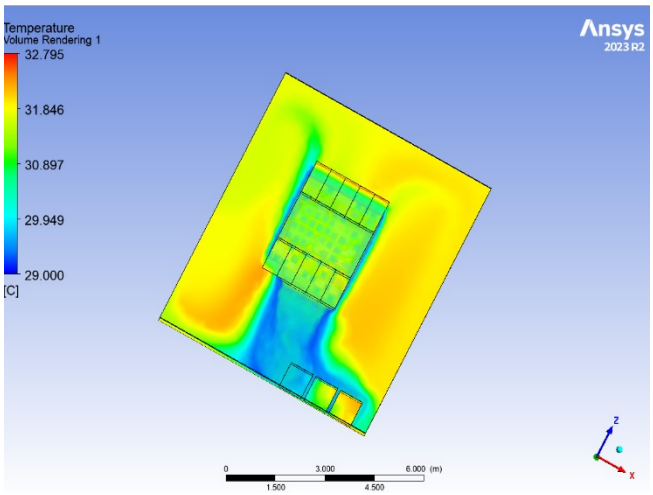

Figure 24. high temperature concentration map for baseline arrangement.

From the above figures (figure 23 and 24), the thermal observations show that some elevated temperatures are behind the high-density rack cluster. However, the maximum temperature is within the range recommended by ASHRAE (under 45 C), thus not indicating any hotspot formation; nevertheless, it still possible to decrease the energy usage by decreasing the mass flow rate and get near the allowable range, this will be studied in the following section 5.1.2.

5.1.2. Airflow Optimization Iterations

To optimize cooling energy use, the total mass flow rate was reduced in steps, maintaining the same rack arrangement.

Table 13. CRAH Mass Flow Rate Iteration Results and Maximum Temperature Outcomes.

Iteration	**Total CRAH Flow (kg/s)**	**Flow per CRAH (kg/s)**	**T_{max} (°C)**	**Meets Limit?**
Baseline	14.33	4.76	32.8	Yes
1	10.00	3.33	38.2	Yes
2	9.51	3.17	41.8	Yes
3	8.00	2.67	46.2	**No**

Through iterated simulations in ANSYS, the above table was generated. Concluding that case 2 is the optimal case with the following conclusions:

- Optimal mass flow: 9.51 kg/s (T_{max} = 41.8 °C), offering ~37% reduction in airflow and fan power vs. baseline.

- Flow below 9.51 kg/s results in overheating (>45 °C).

Results of optimal case (2):

Figures from Post-CFD:

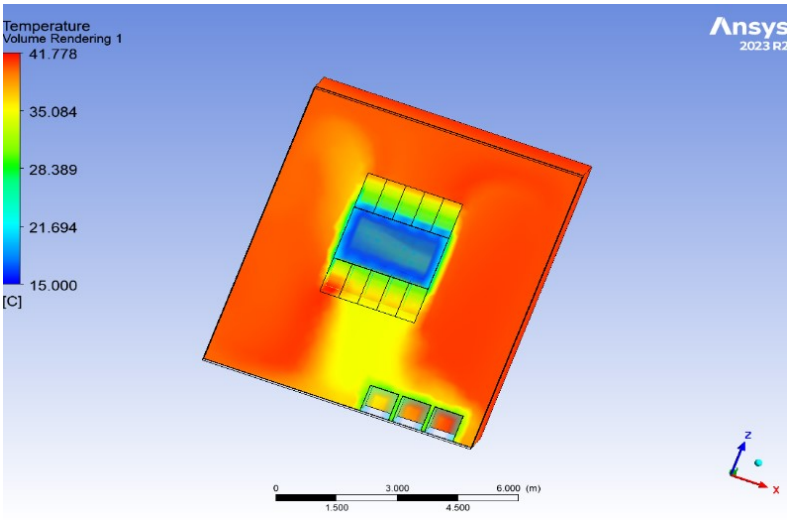

Figure 25. Temperature rendering of case 2 simulation.

Figure 25 displays a three-dimensional temperature rendering of the data center, visually illustrating the thermal distribution and identifying hot regions throughout the room.

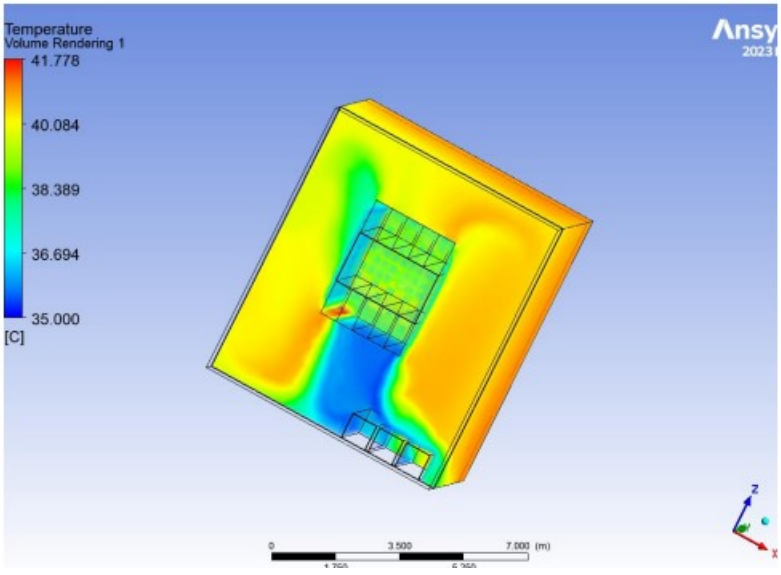

Figure 26. High temperature map for case 2 (optimal).

Moreover, Figure 26 focuses on the high-temperature map, pinpointing the locations and extent of elevated temperatures under the optimized cooling scenario.

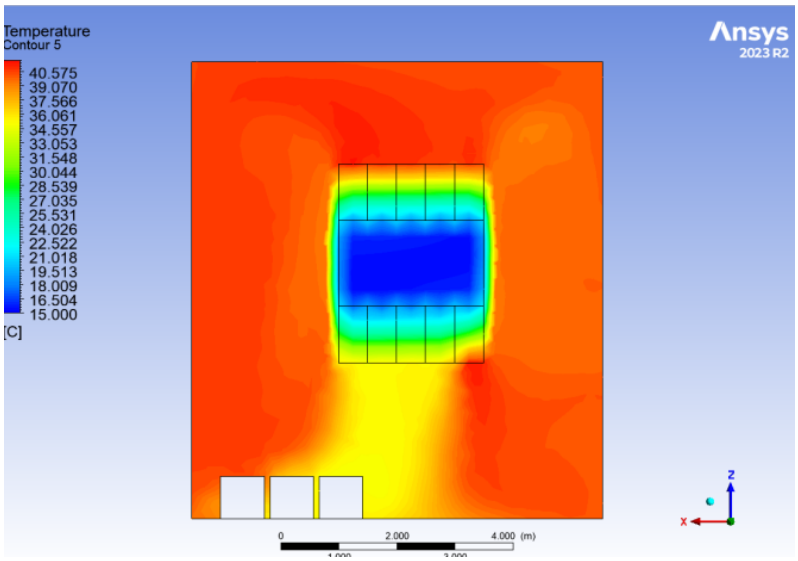

Figure 27. Temperature contour at the mid-section in Y axis (1.5m) for case 2.

With the last figure, providing a temperature contour at the mid-section (Y = 1.5 m), offering a detailed cross-sectional view of the temperature gradients across the cold aisle.

The various figures above show as documented, the maximum temperature that could be observed to be 41.78 C. This increase is directly attributed to the reduction in the airflow rate supplied by the CRAH units, leading to a lower cooling capacity.

Selection of Optimal Case and Further Arrangement Testing

After analyzing the three cases with varying cold air flow rates, the primary objective is to select the optimal configuration that ensures the lowest possible flow rate while maintaining a Tmax below 45°C to prevent hotspot formation. Based on the results, Case 2 has been chosen as the most efficient setup, as it meets these thermal management criteria.

To further optimize the thermal performance of the data center, additional simulations will be conducted by altering the arrangement of high and low-density racks. This study aims to identify the most effective alignment of racks through both visual (temperature maps, contours) and numerical (Tmax, uniformity) analysis, ensuring improved cooling efficiency and airflow distribution.

5.1.3. Comparison of Alternative Rack Layouts

Using the optimal airflow (9.51 kg/s), alternative rack arrangements were simulated to further improve cooling distribution and reduce hotspots.

Table 14. Results for different rack arrangement cases.

Case	**High-Density Racks**	**Standard Racks**	**T_{max} (°C)**	**ΔT_{max} vs. Baseline**
Baseline	R1–R4 together	R5–R10	41.78	—
Case B	R6–R9 together	R1–R5, R10	41.75	– 0.03
Case C	Alternating (R2, R4, R7, R9)	others	41.14	– 0.64

Best result: Alternating high- and low-density racks (Case B) reduced T_{max} to 41.14 °C and provided the most uniform temperature distribution with minimal hotspots.

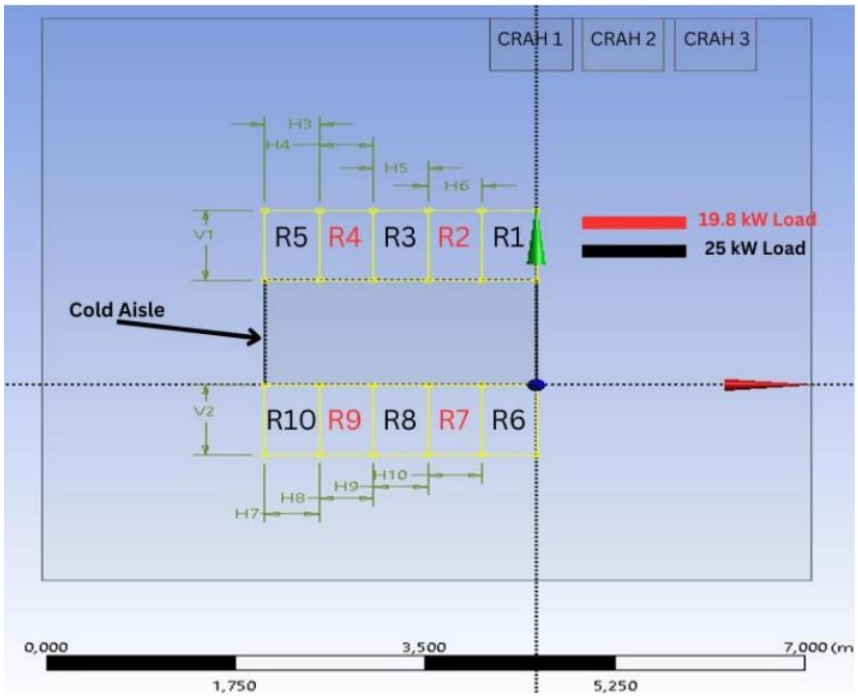

Figure 28. best rack arrangement case C (alternating configuration between low and high density racks).

CFD Ansys Results of the Alternating Configuration:

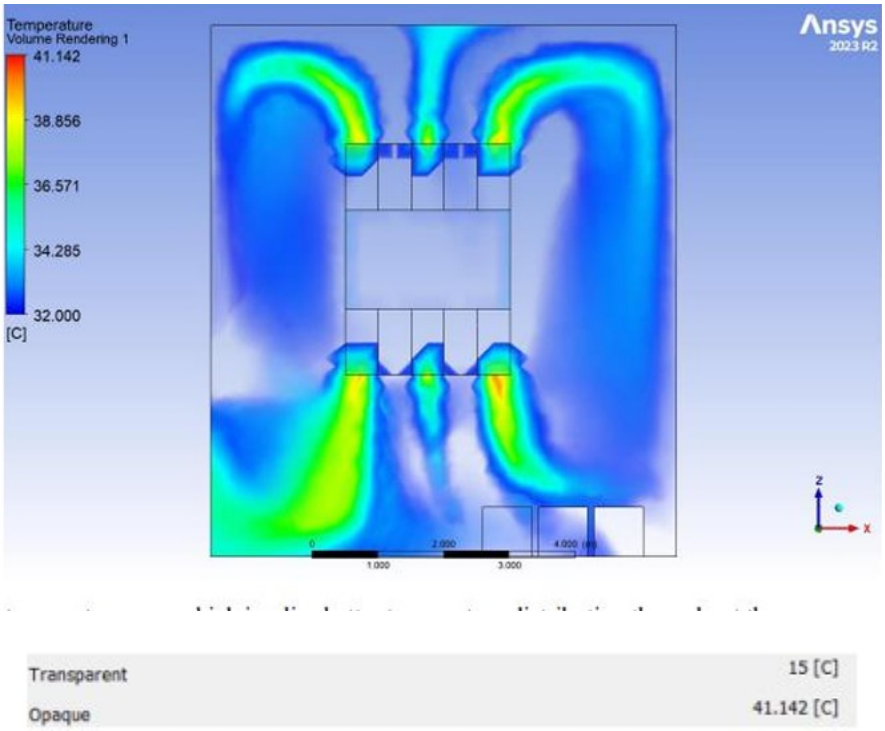

Figure 29. Temperature map of the alternating configuration.

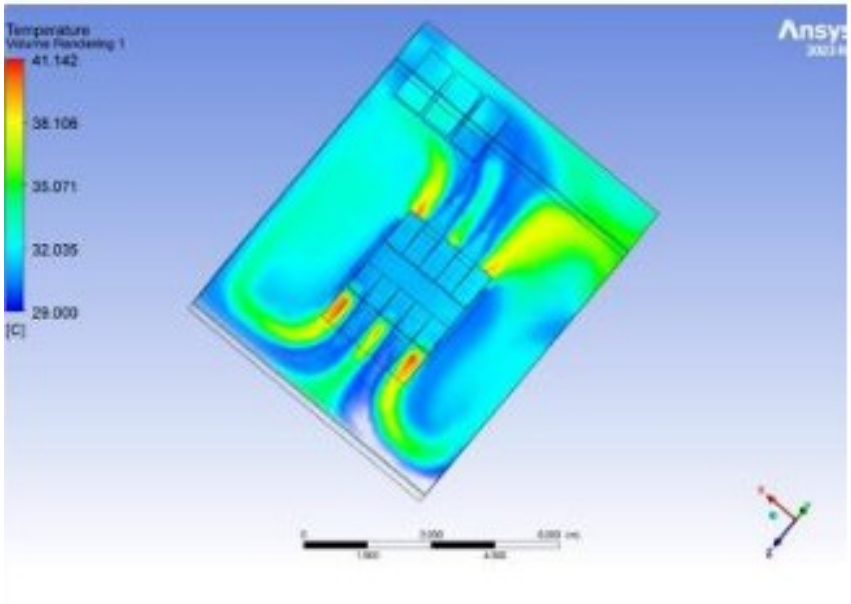

Figure 30. Temperature map of the alternating configuration in isometric view.

The modified rack arrangement in Case C (alternating case) resulted in an improved thermal performance compared to the baseline case, which serves as the base case for comparison. The maximum temperature (Tmax) in Case C was recorded at 41.142°C, which is 0.636°C lower than the Tmax observed in base case. Additionally, analyzing the temperature maps, high-temperature maps, and contour plots reveals a more uniform temperature distribution with fewer high-temperature zones and reduced hotspot concentration. This indicates that the updated rack configuration contributes to better airflow management and improved cooling efficiency within the data center.

6. Full-Scale High-Performance Computing (HPC) Mega Data Center Simulation: CFD Analysis of Normal and Extreme Operation

In this part, a detailed CFD study of the full-scale data center was conducted. The facility consists of 16 pairs of 280 kW setups, totaling 4.48 MW of IT heat load. A hybrid cooling strategy was employed, incorporating both liquid and air cooling, with CRAH units and CDU systems distributing airflow and managing temperature regulation. Two operational scenarios were investigated: normal conditions and extreme working conditions, assessing their impact on hotspot formation and cooling system performance. To ensure proper sizing of CRAH and CDU units, Vertiv's Hi-rating tool and GRS 2.1 tool were utilized for precise cooling capacity calculations. The 280-kW module configuration was inspired by high-density deployments using Dell PowerEdge XE8640 rack servers, with technical specifications referenced from the manufacturer's datasheet (Dell Technologies, 2024; see Appendix B1, B2). Appendix B2 provides a schematic diagram of the internal components of the server used as the basis for the IT load modeling.

6.1. Section Introduction and Modeling Approach

Building on the optimal alternating rack arrangement from Section 5.1, the data center model was scaled up by a factor of 16 to represent a full-scale mega data center, totaling 4.48 MW of IT load (16 modular pairs at 280 kW each).

In this section, a Computational Fluid Dynamics (CFD) analysis is conducted to evaluate and optimize the cooling strategies for a full-scale hybrid data center with a total IT load of 4.48 MW. The facility is modeled by scaling up a 280 kW retrofit module sixteen times, reflecting a modular design approach that aligns with real-world hyperscale data center construction.

To efficiently study system-level performance, a single 280 kW unit—comprising ten server racks, three CRAH units, and an underfloor plenum—is simulated in detail, with results extrapolated to represent the entire 16-module facility.

A hybrid cooling strategy is implemented, utilizing both Computer Room Air Handling (CRAH) units for traditional air cooling and Cooling Distribution Units (CDUs) for targeted liquid

cooling. *The CDU, a specialized heat exchanger, circulates coolant (typically water or water-glycol) through rear-door heat exchangers or cold plates at high-density racks. The absorbed heat is transferred to the main chilled water loop, where it is ultimately rejected to the outside environment via air-cooled chillers or rooftop heat exchangers*.
Two operational scenarios are analyzed:

- **Normal Operation:** Standard IT loads and cooling parameters.
- **Extreme Operation:** Elevated IT loads and increased cooling demands.

The objective is to assess airflow distribution, thermal performance, and energy efficiency under both scenarios, identify potential hotspots, and determine optimal cooling strategies. Cooling system sizing and design validation are performed using Vertiv's Hi-Rating and GRS 2.1 tools, ensuring the solutions are robust and industry compliant.

The CDU modeled in this study is based on proprietary design data provided by Vertiv. Due to the sensitive and confidential nature of this information, specific details, naming conventions, and technical diagrams cannot be publicly disclosed. The unit was sized and configured using Vertiv's GRS tool to accurately reflect real-world performance for high-density cooling applications. For transparency and reproducibility, all key operational parameters used in the simulations are summarized in Table 17.

6.2. Hotspot Definition and Temperature Limits

For the purposes of this study, a hotspot is operationally defined as any region within the data center where the local air temperature exceeds **50 °C**. This threshold is selected based on established reliability guidelines, as prolonged exposure to such elevated temperatures can compromise the performance and lifespan of IT equipment (ASHRAE Technical Committee 9.9, 2021). Identifying and mitigating these hotspots is therefore essential for maintaining data center uptime and minimizing the risk of unexpected hardware failures.

In addition to the hotspot definition, the maximum allowable return air temperature—that is, the temperature of air returning to the Computer Room Air Handler (CRAH) units—is set at **45 °C.** Maintaining this limit is critical for two reasons: first, it ensures that heat is effectively removed

from the data hall; second, it prevents the CRAH units from being subjected to excessive thermal loads, which could otherwise reduce their operational efficiency and increase energy consumption. Adhering to these temperature thresholds supports both the thermal management and energy efficiency objectives of the data center.

6.3. Data Center Geometry and Boundary Condition

6.3.1. Geometry setup:

The computational model represents a retrofitted data center module with the following characteristics, which is then scaled up to represent the full 4.48 MW facility:

- **Room Dimensions:** 8.05 m (L) × 7.06 m (W) × 2.96 m (H)
- **Raised Floor Height:** 0.45 m (underfloor air distribution)
- **Server Racks:** 10 racks per module (2.0 m × 0.5 m × 1.0 m each), arranged in two rows (alternating high/low-density as established in previous optimization)
- **Cooling Infrastructure:** 3 CRAH units per module (1.75 m height, 0.75 m × 0.75 m footprint)
- **Perforated Floor Tiles:** 56% open area for efficient airflow to the cold aisle
- **Airflow Direction:** Z-axis; first five racks face –Z, second five face +Z

6.3.2. Boundary Conditions:

- **Normal Operation:** Boundary conditions set to reflect standard IT and cooling loads (see table 15 and 17)
- **Extreme Operation:** Adjusted mass flow, IT loads, and cooling temperatures to reflect high-stress, peak-load scenarios (see table 15 and 17).

6.3.3. Summary of Ansys Boundary Conditions Used in CFD Simulation:

Table 15. Summary of Ansys Boundary Conditions Used in CFD Simulation.

Boundary Name	**Type**	**Description**
CRAH_Inlet_1, CRAH_Inlet_2, CRAH_Inlet_3	Mass Flow Inlets	Cold air supply from CRAH units at a fixed temperature.
CRAH_Outlet_1, CRAH_Outlet_2, CRAH_Outlet_3	Mass Flow Outlets	Return hot air to CRAH units, with a 0.5% reduction in mass flow due to internal losses.
Rack_Porous_Zone_1 to Rack_Porous_Zone_10	Porous Zone	IT racks modeled as porous media to account for airflow resistance.
Tile_Porous_Zone	Porous Zone	Perforated tiles with a 56% opening for airflow distribution.
Fan_Outlets (R2, R4, R7, R9)	Pressure Jump Outlet	Fans modeled with a pressure jump curve to represent airflow enhancement.
Walls (Room, Racks, CRAH, Floor, Cold Aisle, Ceiling)	No-Slip Wall	Solid boundaries set with no-slip conditions.

Note: Exact values for each boundary condition are provided separately for normal and extreme scenarios.

6.3.4. The datacenter configuration model:

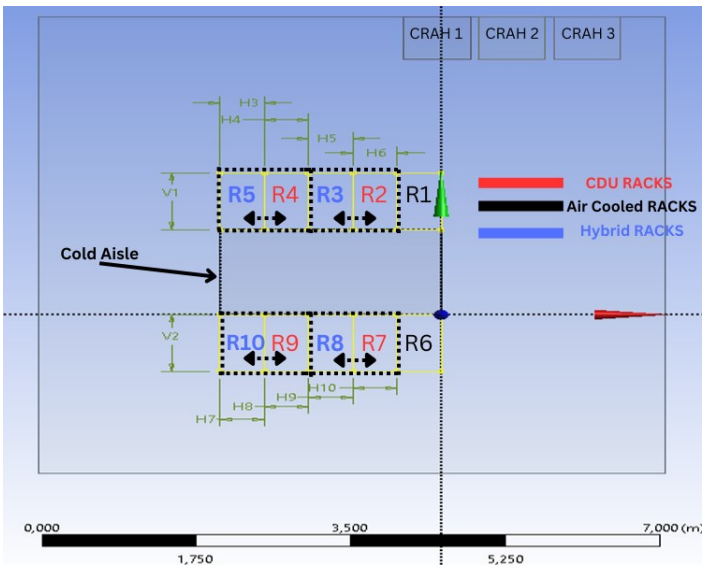

Figure 31. Modeled retrofit setup model of 280kW data center.

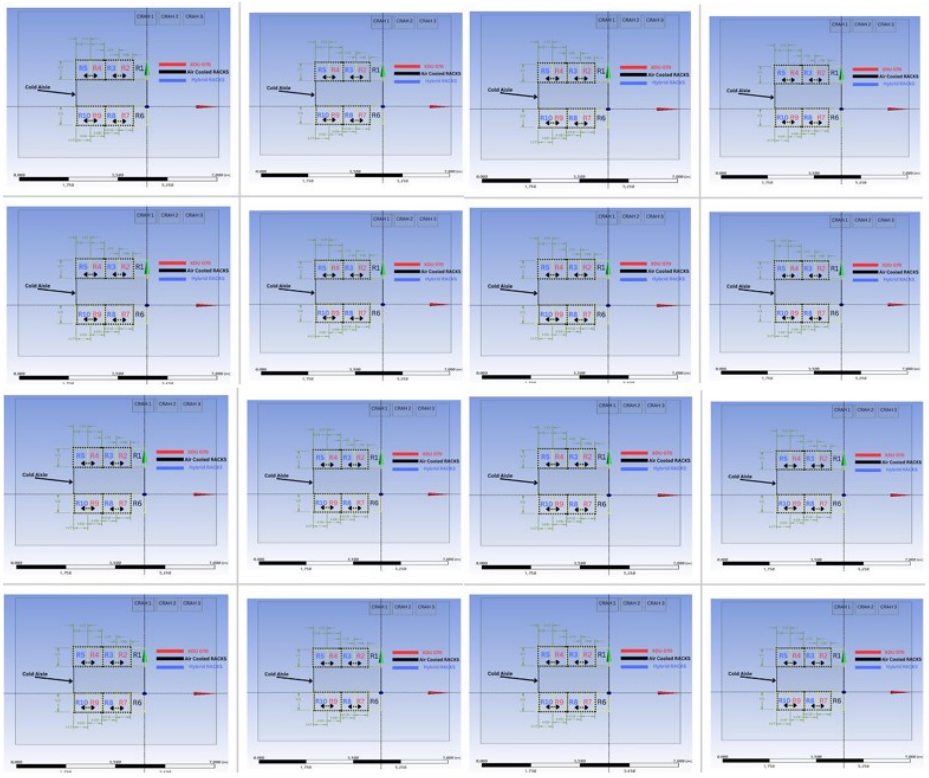

Figure 32. Scaling up the data center to the 16 X the simulated model.

The above figure shows:

- Total module IT load: 280 kW
- Scaling: 16 identical modules = 4.48 MW total facility IT load

The below table lists each rack number with its type and IT load that will be cooled by CRAH supply units, for both normal and extreme cases:

Table 16. lists the rack number, cooling type, and corresponding IT load (in kW) as used in both the normal and extreme scenarios. Values are updated for each scenario as needed.

RACK number	**Type**	**IT Load**
R1	AIR COOLED	**26 kW**
R2	XDU 070	**40 kW**
R3	Hybrid	**17 kW**
R4	XDU 070	**40 kW**
R5	Hybrid	**17 kW**
R6	AIR COOLED	**26 kW**
R7	XDU 070	**40 kW**
R8	Hybrid	**17 kW**
R9	XDU 070	**40 kW**
R10	Hybrid	**17 kW**

Note: In hybrid racks, a portion of the load is removed by liquid (CDU), the remainder by air (CRAH).

Boundary values for airflow rates, supply and return temperatures, and pressure losses are specified for both normal and extreme cases (summarized in operational tables below).

Operational parameters table - for normal and extreme operation cases:

Table 17. Key operational parameters for normal and extreme operating cases.

Parameter	**Normal Operation**	**Extreme Operation**	**Notes**
Total CRAH Airflow (kg/s)	19.94	15.99	Distributed across 3 CRAHs
CRAH Supply Air Temp (°C)	23.5	23	At underfloor plenum entry
Return Air Temp Limit (°C)	45	45	Thermal design target
Hotspot Temp Threshold (°C)	50	50	For hotspot definition
CDU Liquid Cooling Fraction (%)	81.47	81.47	% of IT load removed by CDU
Air-Cooled IT Load (%)	18.57	18.57	% of IT load cooled by air
CRAH Unit Pressure Loss (%)	0.5	0.5	Internal component losses
IT Load per Module (kW)	280	280	Per 10-rack module
Facility IT Load (kW)	4480	4480	16 × module

Table 17 summarizes the principal operational parameters applied in the CFD simulations for both the normal and extreme scenarios. It details airflow rates, temperature setpoints, and load distributions across the facility, providing a direct comparison of the baseline and stressed conditions.

6.4. Cooling System Sizing and Fan Data

- Summarize use of Vertiv's Hi-Rating (CRAH) and GRS 2.1 (CDU) tools.
- Present manufacturer datasheet snippets (CDU and CRAH at normal & extreme).
- Short explanation of the fan performance curve implementation in Fluent.

To ensure robust and reliable thermal management in both normal and extreme operating scenarios, the cooling system components were sized using manufacturer-recommended engineering tools provided by Vertiv. Two main software platforms were utilized:

- **Hi-Rating Tool**: Used to select and size the Computer Room Air Handling (CRAH) units according to the required airflow rates, supply/return temperatures, and internal pressure drops. The tool provided detailed performance data for CRAH operation under various boundary conditions, allowing precise input for CFD simulation.
- **GRS 2.1 Tool**: Used to configure and size the Cooling Distribution Units (CDUs) responsible for liquid cooling of the hybrid racks. This software outputs key parameters including flow rates, required heat transfer, pressure losses, and power consumption for the CDU fans.

Key Manufacturer Data Inputs

- The CDU specification data (heat exchange performance, operating pressure, and fan power) were obtained for both normal and extreme conditions, ensuring accurate representation of hybrid cooling behavior in the simulation.
- For each scenario, the datasheet outputs for the selected CRAH and CDU units were used to define inlet temperatures, total airflow rates, and fan characteristics in Ansys Fluent.

Fan Performance Curve Implementation

- The fan performance for the CDU units was implemented in Fluent by digitizing the manufacturer's fan curve and applying it as a piecewise-linear profile for the relevant pressure jump boundaries (see Table 15).
- This method ensures that airflow through hybrid racks (R2, R4, R7, R9) matches real-world behavior, capturing the influence of the CDU fan on cooling performance and pressure distribution.

Table 18. Key Operational Parameters for CDU and CRAH Units (per 280 kW Module, Normal and Extreme Cases).

Parameter	**CDU (Normal)**	**CDU (Extreme)**	**CRAH (Normal)**	**CRAH (Extreme)**
Cooling Technology	Liquid-to-air	Liquid-to-air	Air-cooled	Air-cooled
Cooling Capacity (kW)	40.0	40.0	302.8	291.3
Airflow Rate (m³/h)	8,947	8,500	58,569	47,038
Supply Air Temp (°C)	23.5	23.0	23.6	22.6
Return Air Temp (°C)	—	—	39.0	42.0
Discharge Air Temp (°C)	37.2	37.0	—	—
Inlet Fluid Temp (°C)	41.0	41.0	20.0 (glycol)	20.0 (glycol)
Outlet Fluid Temp (°C)	31.1	31.0	32.0 (glycol)	32.0 (glycol)
Power Input, Fans (kW)	0.7	0.5	12.92	7.12
Power Input, Total (kW)	0.9	0.6	13.02	7.22
Fan Static Pressure (Pa)	51.9	50.0	422	287
Number of Fans	7	7	4	4
Supply Air RH (%)	—	—	71.8	89.5
Fluid Flow (l/s)	—	—	6.56	6.46
ESP Air (Pa)	—	—	20	20

Key input parameters for each cooling system were extracted using the manufacturer's specialized engineering software tools—GRS 2.0 for the Cooling Distribution Unit (CDU) and

HiRating for the Computer Room Air Handler (CRAH) units. These tools provided detailed performance data, including recommended flow rates, operating pressures, temperature setpoints, and fan power characteristics for both normal and extreme operating conditions. The use of manufacturer-validated data ensures that the modeled boundary conditions and equipment behavior in the simulation closely match real-world performance. Complete datasheets and output summaries for the cooling units are included in Appendix A1 and Appendix A2, respectively, for reference and validation purposes.

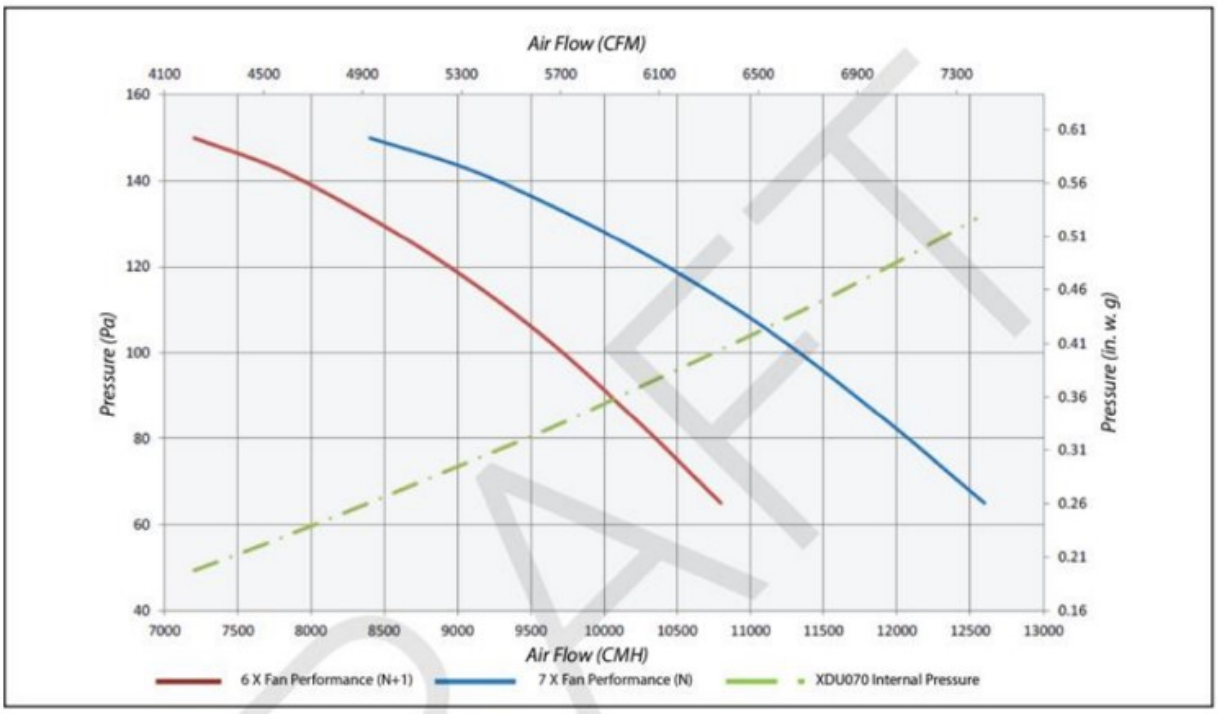

Figure 33. Fan characteristics curves for the CDU. Retrieved from: Vertiv™ Coolant Distribution Unit (Liquid to Air) Application and Planning Guide.

Regarding the fans of the CDU unit the fan curve was based on the manufacturer's fans with the following points taken to have a graph input for the fan boundary condition:

Table 19. Fan Performance Data: Airflow and Pressure Relationship.

Airflow (CFM)	**Pressure (Pa)**	**Airflow (m^3/s)**
7100	150	3.35
8000	135	3.77
8800	120	4.15
9500	100	4.48
10200	85	4.81
10800	65	5.10
11000	60	5.19

After inputting the points extracted in the table in Fluent, the following fan curve was got which is in-line with the manufacturer's specifications by using the "Piecewise-Linear Profile".

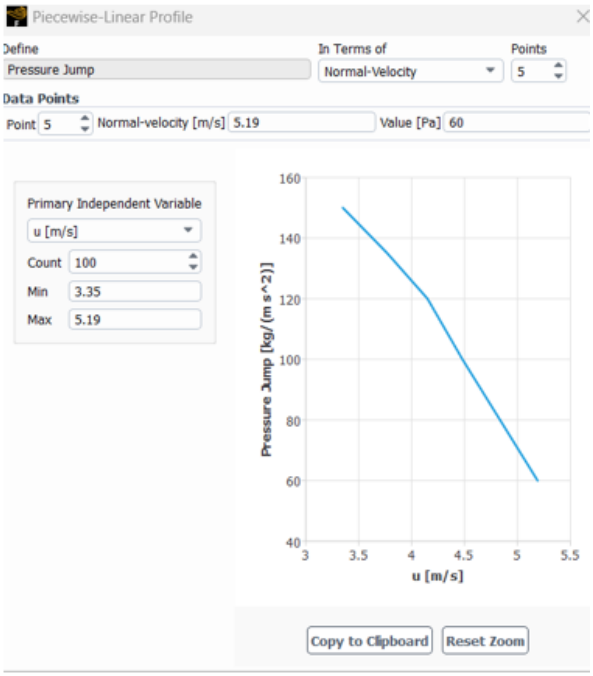

Figure 34. Fan performance curve inputted in Fluent.

With the specified fan inputs, the CDU units are modeled to replicate the manufacturer-supplied airflow performance data. This ensures that, from an airflow perspective, the simulated behavior of the CDU units in the model accurately reflects real-world operation as documented by the equipment manufacturer.

6.5. CFD Simulation and Mass Balance Verification (both cases)

Overview of CFD Model Setup

The CFD simulations for both the normal and extreme operating cases were conducted using ANSYS Fluent. The computational domain, mesh structure, and boundary condition setup were kept consistent with the approach described previously in Section 3.2. All relevant boundary conditions including those for the CRAH inlets, outlets, IT racks, and perforated tiles, were assigned in accordance with the methodology detailed earlier. A steady-state solver was utilized, and the Realizable k–ε turbulence model was selected to ensure accurate prediction of indoor airflow patterns and recirculation zones. This consistent modeling framework ensures that comparisons between the normal and extreme cases are robust and attributable solely to changes in operational parameters.

Convergence Criteria

Numerical convergence was assessed by monitoring the **scaled residuals** of continuity, momentum, and energy equations. The solution for each scenario was deemed converged when all residuals dropped below 1×10^{-3} and remained stable for at least 100 additional iterations.

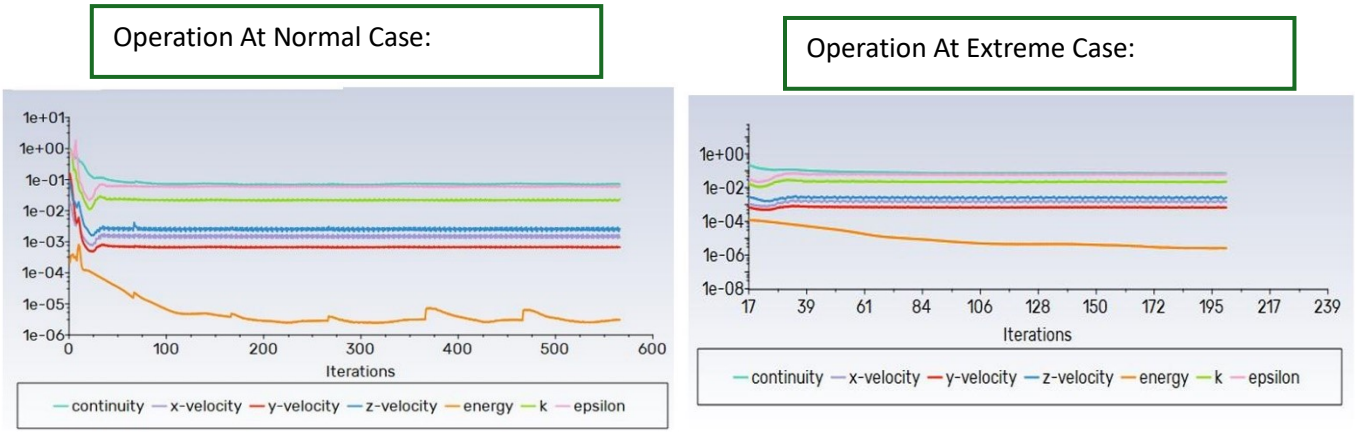

Figure 35. Scaled residuals for normal and extreme operation cases.

The scaled residual graphs in Fluent indicate convergence when they consistently decrease and level off at low values. This means the numerical solution has stabilized, and further iterations will not significantly change the results. Since our residuals meet standard convergence criteria, we can confidently rely on the obtained data for airflow, heat dissipation, and hotspot analysis.

Mass Balance Verification

To verify the physical validity of each simulation, mass balance checks were performed at all critical boundaries which include:

- Total mass flow in (CRAH Inlet)
- Total mass flow through racks and tiles
- Total mass flow out (CRAH Outlet)

The mass balance equation is:

$$\sum \dot{m}_{inlet} \approx \sum \dot{m}_{outlet}$$

The bellow table summarizes the total mass flow rates at the main boundary surfaces (CRAH inlets, tiles, rack outlets, and CRAH outlets) for both normal and extreme operating scenarios, as obtained from the CFD simulation. The net difference between inflow and outflow at each boundary is minimal, confirming mass conservation in all cases.

Table 20. Mass Flow Rate Summary at Critical Boundaries (per 280 kW module).

Scenario	**CRAH Inlet (kg/s)**	**Tiles In (kg/s)**	**Average Racks Out (kg/s)**	**CRAH Outlet (kg/s)**	**Net Balance (kg/s)**
Normal Operation	19.94	19.98	19.98	19.93	0.01
Extreme Case	15.99	15.99	16.02	15.98	0.03

In both cases, the simulated mass flow rates at all boundaries are closely balanced (A discrepancy within 1% is considered acceptable for engineering CFD studies). Confirming that the solution respects the conservation of mass. This validation step reinforces the reliability of the CFD results for subsequent thermal and airflow analysis.

Results – Normal Operation

Mass conservation is applied between the CRAH units' cold air supply in the raised floor and the rack inlets and outlets. This step verifies that the airflow distribution remains consistent throughout the system, reinforcing the validity of the simulation results.

For the racks:

$$\sum \dot{m}_{rack\ inlets} = \sum \dot{m}_{rack\ outlets}$$

The net value was calculated by Fluent and was zero, and since there is no mass accumulation is steady stated we could say that:

$$\sum (\dot{m}_{inlet} - \dot{m}_{outlet}) = 0$$

For the whole flow:

$$\dot{m}_{CRAH\ supply} = \sum \dot{m}_{rack\ inlets}$$

$$\dot{m}_{CRAH} = 19.94 \text{ Kg/s}$$

$\dot{m}_{rack}$ = 1.412+3.035+1.392+3.003+1.380+1.407+3.020+1.386+3.021+1.403 = 19.980 Kg/s

This confirms that mass is effectively conserved across the system boundaries, indicating that the simulation setup correctly accounts for all airflow inputs and outputs. As a result, the mass balance validation reinforces the credibility of the numerical results and supports the overall reliability and accuracy of the study's findings.

Airflow Patterns and Velocity Distributions: Normal Case

The velocity streamline plots shown below reveal the detailed airflow distribution within the data center under normal operating conditions. Cool air from the CRAH units is supplied into the underfloor plenum, then delivered through perforated tiles to the cold aisles. The alternating rack configuration and the presence of CDU-equipped racks produce well-organized flow paths, minimizing recirculation and promoting effective cooling across all racks.

Figure 36 presents the velocity streamlines of the airflow throughout the room under normal working conditions. The flow clearly originates from the CRAH units, enters the underfloor plenum, and rises through the perforated tiles into the cold aisles. The alternating rack layout, in combination with the CDU units, contributes to a coherent and directed airflow pattern that reduces undesirable recirculation zones.

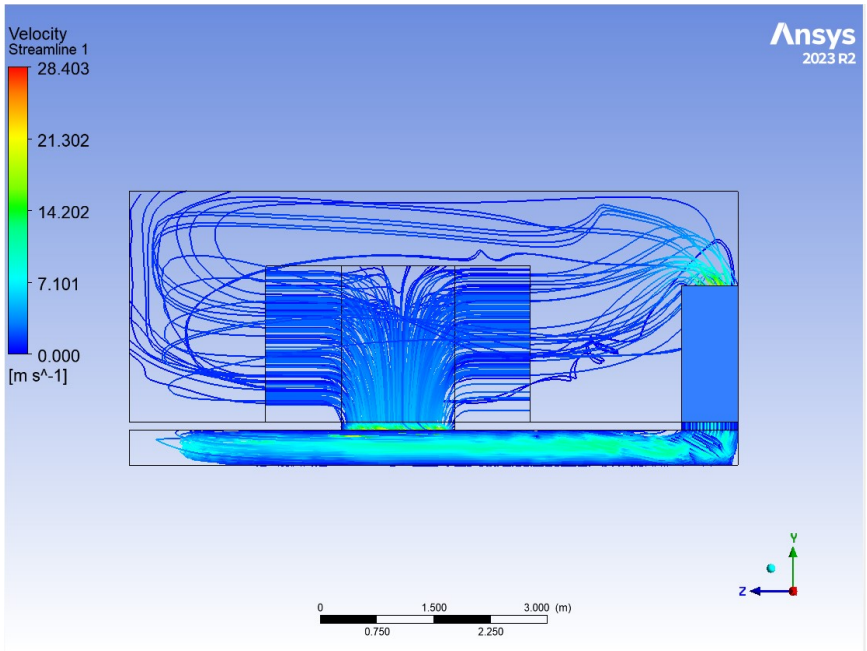

Figure 36. Normal working condition velocity streamlines.

To further emphasize the flow behavior around the CDU-equipped racks, Figure 37 provides an isometric view of the same streamlines. This perspective highlights the intensified airflow around racks R2, R4, R7, and R9, where CDUs are installed. These localized enhancements ensure that high-density racks receive sufficient cooling, and the flow pattern remains undisturbed across adjacent zones by the fans boundary condition.

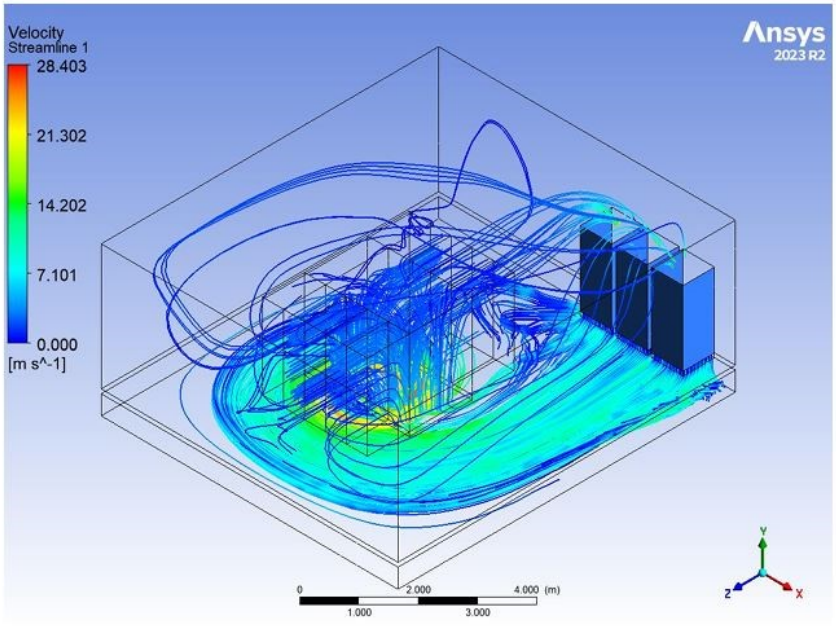

Figure 37. Isometric view of streamlines highlighting enhanced flow through CDU racks.

Next, Figure 38 depicts the velocity contours at a mid-height plane (Y = 1.6 m), offering insight into the lateral distribution of air as it moves through and around the racks. High velocity regions are clearly visible in front of racks aligned with open tiles, confirming that cold air is effectively reaching the cold aisle.

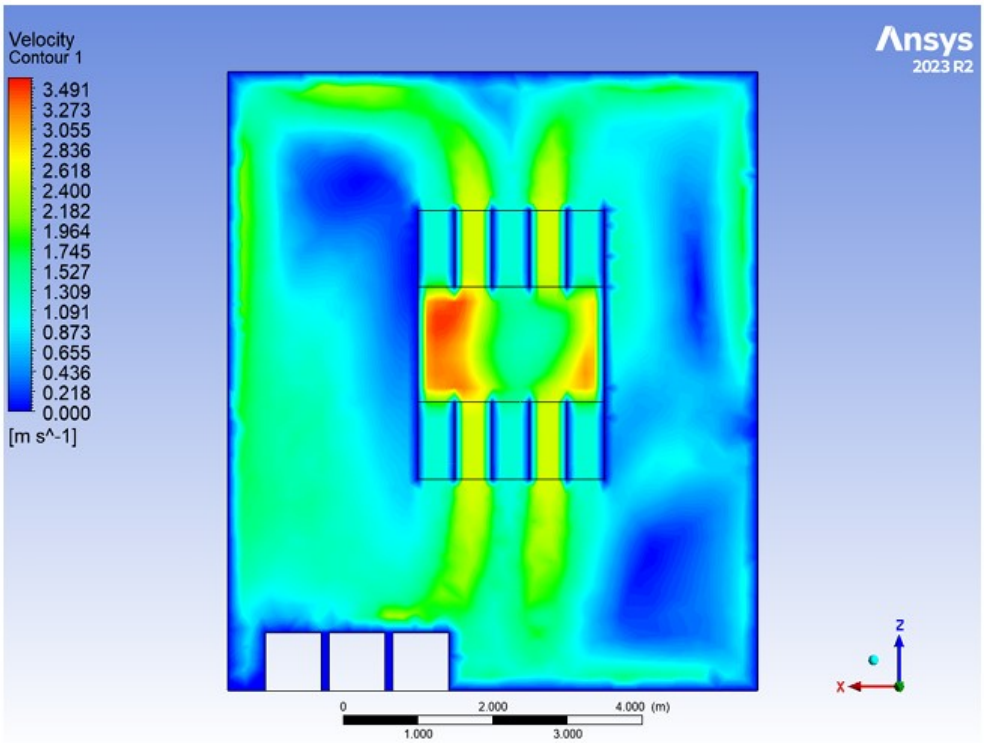

Figure 38. Velocity contours at representative at mid-height (Y=1.6 m).

A complementary view is provided in the below figure, which shows the velocity contours at a higher plane (Y = 2.3 m), corresponding to the return air level at the CRAH units. The increase in velocity at this level reflects the upward movement of warm return air, especially near the CRAH return zones.

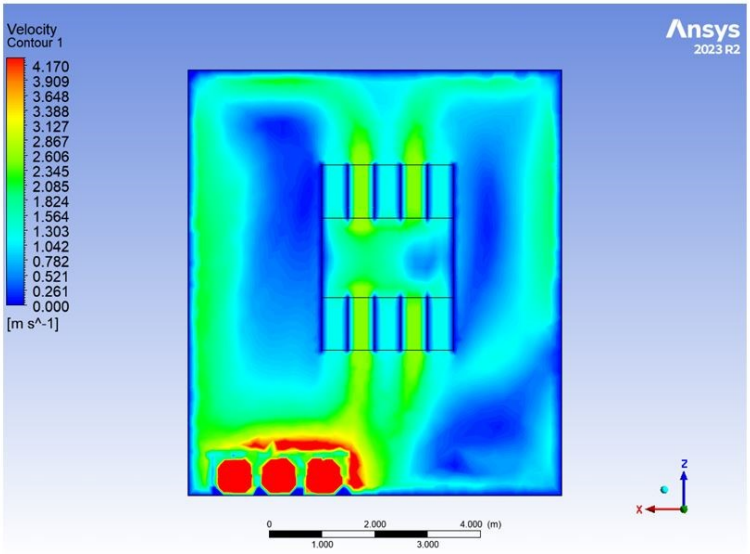

Figure 39. velocity contour at Y=2.3 m (level of the CRAH return air).

To better analyze airflow through individual racks, Figures 5 and 6 isolate the outflow velocity contours for racks 1–5 and 6–10, respectively. These figures allow for direct comparison of cooling performance across the rack row, where enhanced flow is again visible around CDU-equipped racks. The figures validate the CFD model's prediction of differential cooling rates based on equipment density and flow assistance.

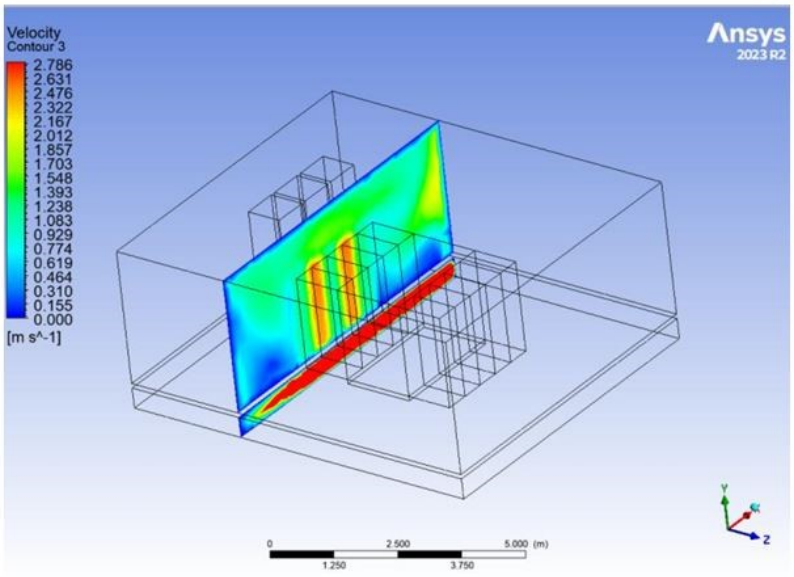

Figure 40. Outflow velocity contours for racks 1-5.

From the above figure (Figure 40) we realize that the rack numbers R2 and R4 have the highest velocity with respect to this row of racks which is what we could expect due to the additional fans presented at these CDU liquid-to air-cooled racks.

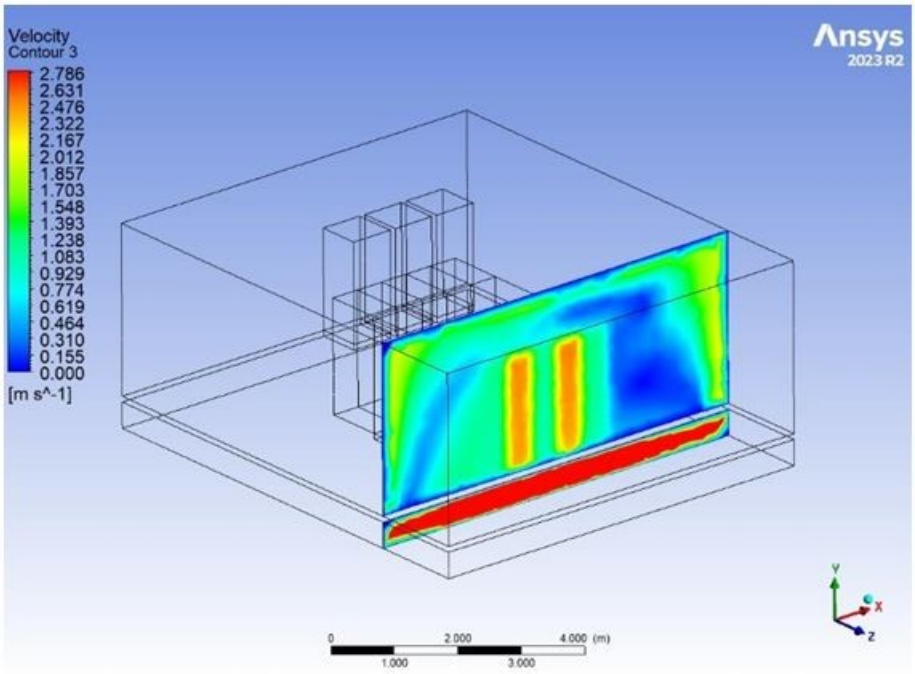

Figure 41. Outflow velocity contours for racks 6-10.

Similarly, the above figure (Figure 41) shows that that the rack numbers R7 and R9 have the highest velocity with respect to this row of racks which is what we could expect due to the additional fans presented at this CDU liquid-to-air cooled racks.

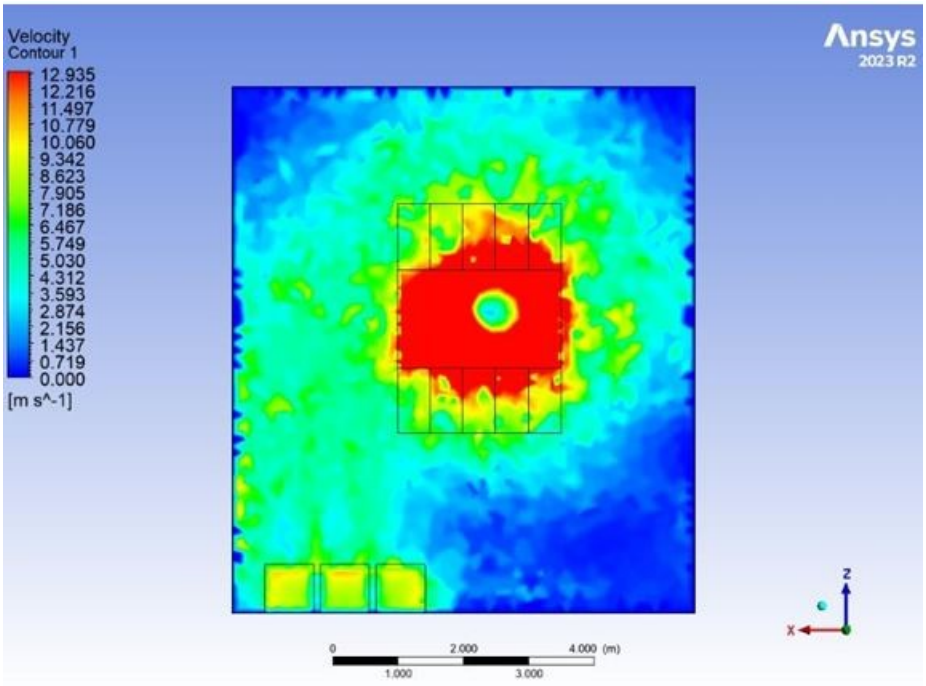

Figure 42. velocity contour concentrating at the maximum velocity spots, open tiles location.

Moving on to Figure 42, the maximum velocity spot is observed at the opening of the tiles, where the airflow reaches 12.935 m/s. This is considered the peak velocity in the study and occurs at a predictable location due to the direct supply of pressurized cold air from the underfloor plenum through the perforated tiles. At these points, the velocity is naturally highest as the flow transitions from a confined plenum space into the open room environment, encountering minimal resistance and high momentum as it enters the cold aisle.

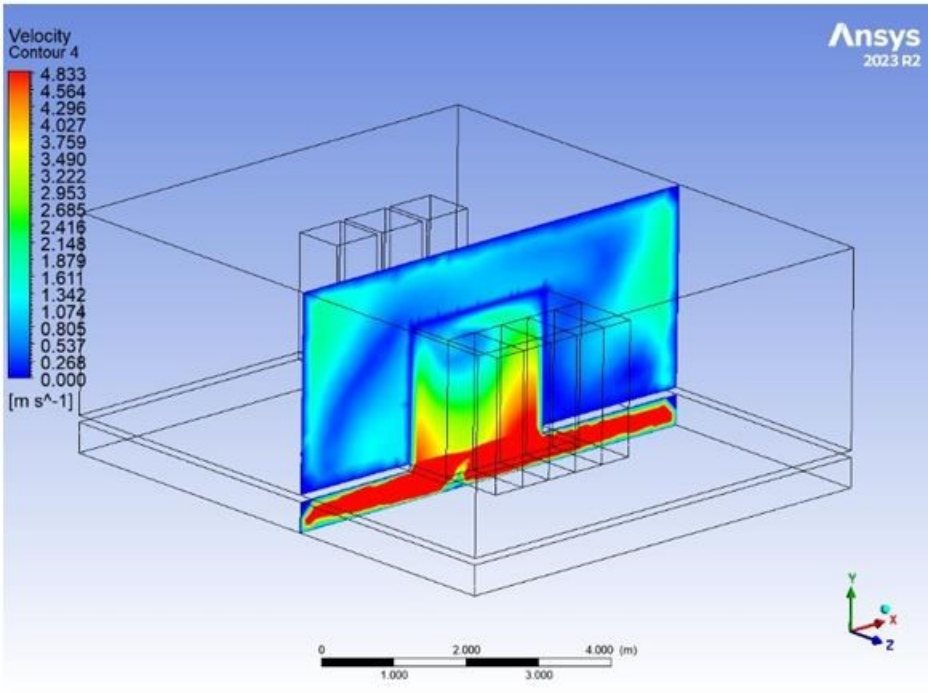

Figure 43. Velocity contour midplane in Z direction.

Figure 43 displays a cross-sectional velocity contour of the data center, highlighting the airflow behavior from the underfloor plenum through the perforated tiles and into the cold aisle region. While the previous figure emphasized the localized velocity peak at tile openings, this view provides a broader understanding of how the supplied cold air propagates upward and recirculates within the room. The red region along the raised floor confirms high-velocity jets existing the tiles, gradually slowing as they mix with surrounding air. Notably, the central cold aisle maintains a relatively uniform upward flow, whereas the periphery exhibits velocity decay and lateral dispersion, suggesting possible inefficiencies or early mixing. This visualization reinforces the importance of tile placement and underfloor pressure management in sustaining directed and laminar flow toward IT equipment inlets.

Thermal Analysis: Normal Case

The thermal performance of the data center under normal operating conditions was evaluated using temperature contours and volume renderings from the CFD simulation. This assessment

focuses on the effectiveness of the cooling strategy, the identification of potential hotspots, and comparison with manufacturer data (Vertiv HIRATING and GRS tools).

CRAH and Return Air Temperatures

The average cold air supply temperature at the CRAH outlets and the average return air temperature were compared against unit design specifications (datasheet available in the Appendix). The simulation produced a return air temperature of approximately **39 °C**, matching the expected value for the selected cooling configuration.

Maximum Temperature and Hotspot Analysis: Outline

- **Maximum rack outlet temperature** ($\mathbf{T_{max}}$) was recorded at the exits of R1 and R6, reaching **43.97 °C**.
- **Temperature volume renderings** and **high-temperature zone maps** (see Figures 44 and 45) visually demonstrate that all rack outlets and return air temperatures remain below the hotspot threshold of **50 °C**.
- **Temperature contours** at mid-height (Y = 1.6 m) and at rack outlets confirm a uniform temperature distribution across the data center, with no evidence of localized overheating.

Rack-Level Temperatures:

The table below summarizes the simulated average inlet and outlet temperatures for each rack:

Table 21. Average Inlet and Outlet Temperatures of CRAH Units and Racks (Normal Case).

Location	Inlet Temp (°C)	Outlet Temp (°C)
CRAH 1	23.50	39.41
CRAH 2	23.50	39.20
CRAH 3	23.50	39.54

R1	24.88	41.78
R2	24.54	39.68
R3	24.37	35.93
R4	24.67	39.93
R5	24.50	35.61
R6	24.87	41.81
R7	24.61	39.79
R8	24.61	35.48
R9	24.65	39.85
R10	24.44	41.78

The results in Table 21 show that the cold aisle supply temperatures from the CRAH units are highly uniform (23.5 °C). Rack inlet temperatures range from 24.37 °C to 24.88 °C, reflecting slight variations due to local airflow and rack positioning. The maximum rack outlet temperature (T_{max}) is 41.81 °C (R6), well below the critical threshold of 50 °C. Outlet temperatures at other racks are also consistently maintained below risk levels, indicating that the cooling system efficiently handles the IT load under normal operating conditions. The average return air temperature is approximately 39 °C, aligning closely with the designed return temperature of the cooling units.

These findings confirm that the cold aisle containment strategy and current CRAH/CDU operation effectively manage thermal loads and prevent hotspot formation in the data center.

ANSYS CFD Thermal figures:

The temperature data obtained from the CFD simulation were extracted at all major inlets and outlets within the data center. Table 21 summarizes the area-weighted average temperatures at the CRAH supply, rack inlets, and rack outlets for the normal operating scenario. These quantitative results, together with the temperature contour plots shown in the below figures,

provide detailed verification of the summary findings stated in the table. The data confirms that both supply and return air temperatures are maintained within the desired ranges, and that the maximum rack outlet temperatures do not exceed the design threshold.

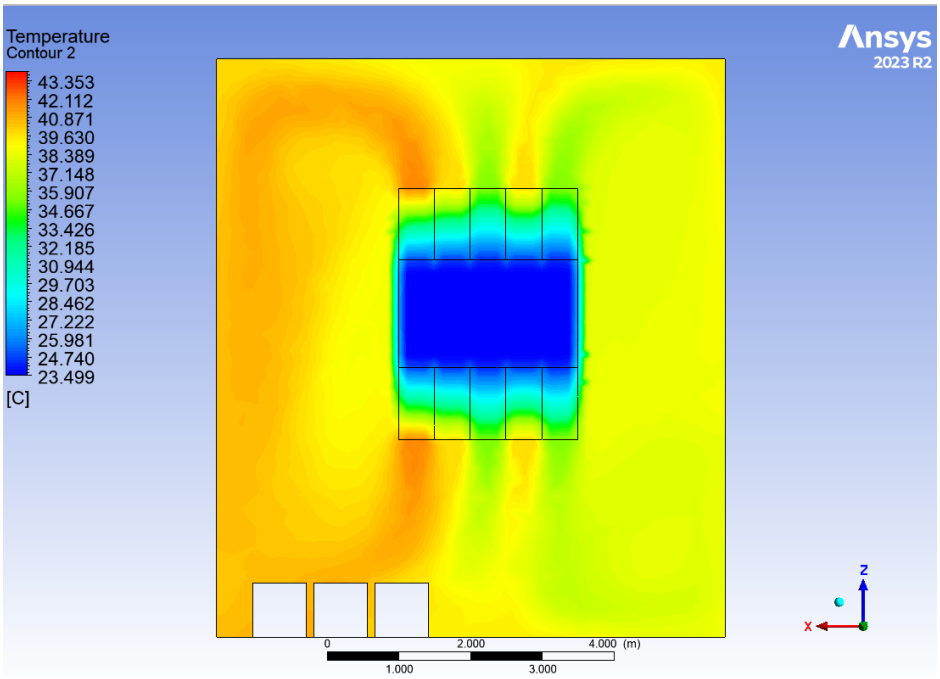

Figure 44. contour at Y= 1.6m (mid section in height) CRAHs running on normal conditions with flow of 19.94 Kg/s (6.64 each).

From the temperature contours and the rendering of the volume inside the room analyzed, the maximum temperature (T_{max}) reached in the room is observed at the exit of **R1 and R6**, with a temperature of **43.973°C**.

Temperature volume rendering:

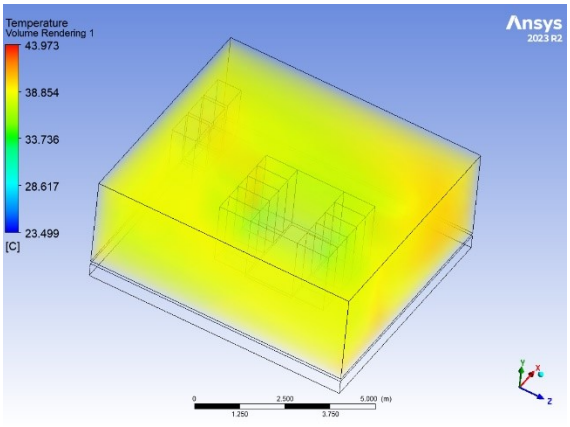

Figure 45. Temperature volume rendering top view.

Figure 45 offers a top view of the temperature volume rendering, providing a comprehensive spatial overview of the thermal field inside the data center. This visualization highlights how heat is distributed across the room, particularly within the aisles and near the rack exhausts.

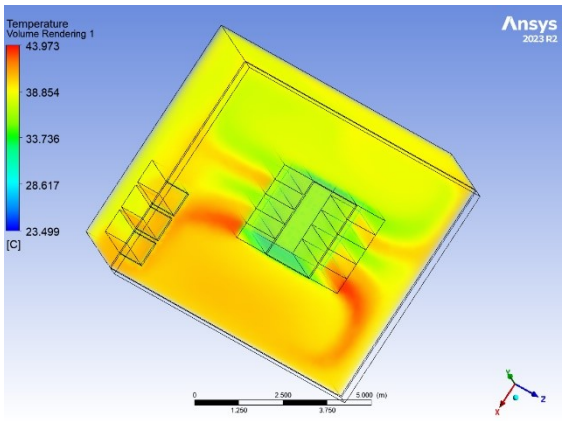

Figure 46. Temperature volume rendering isometric view from the bottom.

Moving on to figure 46, it displays an isometric view of the same temperature volume rendering, this time from the bottom. This perspective emphasizes the vertical temperature gradients and helps to identify how hot air rises and stratifies within space.

Focusing on the most critical thermal regions, Figure 47 presents a high temperature map, isolating and visually representing only those areas where the temperature exceeds 36 °C. This

targeted view allows for quick identification of potential hot spots that may threaten equipment reliability.

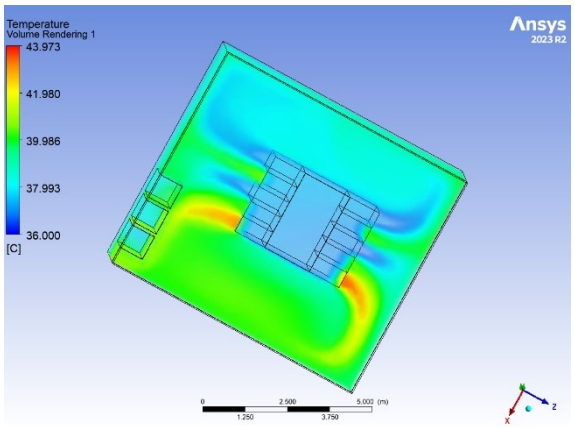

Figure 47. high temperature map (showing temperature above 36 C).

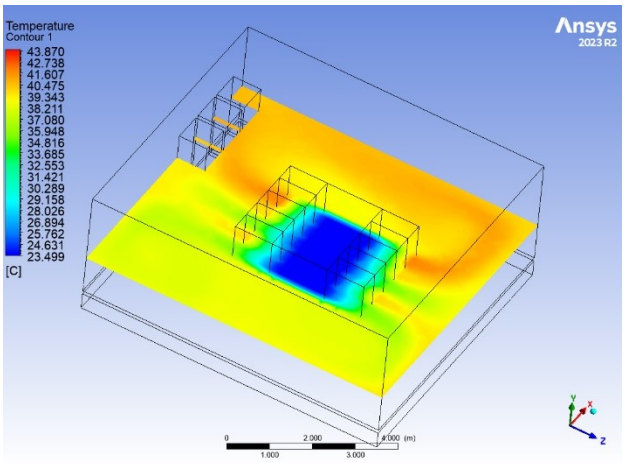

Figure 48. temperature contour at the mid-height of the data center.

Figure 48 shows the temperature contour at the mid-height plane of the data center, further illustrating how temperature varies laterally across the space. This cross-sectional view complements the volume renderings by clarifying the extent and distribution of moderate and elevated temperatures.

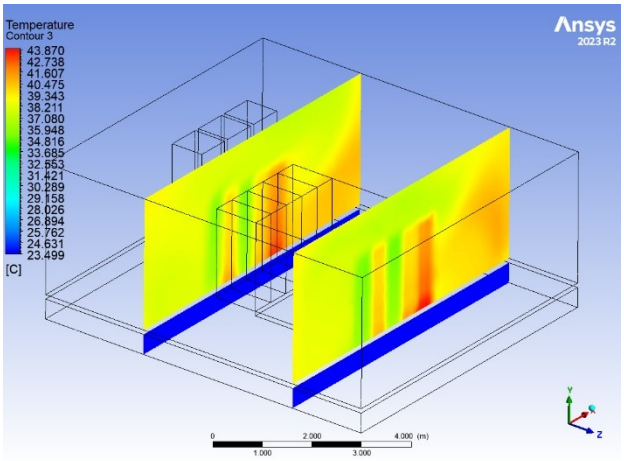

Figure 49. temperature contours at the outflows of the racks.

Figure 49 displays the temperature contours at the outflows of the server racks for the normal operating case. The highest temperatures are observed at the air-cooled racks (R1 and R6), which is primarily due to the lower airflow passing through these units compared to the CDU-equipped racks. The CDU racks, equipped with active fans, draw higher volumes of air, resulting in more effective heat removal and lower exhaust temperatures. Despite these localized temperature peaks—reaching up to 43.87 °C—no region exceeds the 50 °C hotspot threshold. This confirms that, even at the most thermally loaded points, the cooling strategy effectively prevents the formation of critical hotspots.

Conclusion for the Normal Operating Case

The results demonstrate that the current cooling configuration efficiently manages the data center's heat load, with all critical temperatures remaining safely below risk thresholds. The highest temperatures are observed at the air-cooled racks, primarily due to their lower airflow compared to CDU-equipped racks, yet even these values do not exceed the defined hotspot criterion of 50 °C. The temperature distribution throughout the room is relatively uniform, and the return air temperature aligns closely with manufacturer-specified performance, which serves to validate both the CFD model and the adopted cooling strategy.

Overall, the normal operating case satisfies the established thermal requirements. The maximum recorded temperature remains below 50 °C, confirming that no critical hotspots develop within

the environment. Furthermore, the average return air temperature is consistently maintained below the 45 °C threshold, demonstrating that the cooling system effectively supports heat dissipation and airflow management across all rack types under standard operational conditions. These findings confirm the reliability and adequacy of the current cooling approach for maintaining data center performance and equipment safety.

6.6. Results – Extreme Operation

This section analyzes the thermal and airflow performance of the retrofitted data center model under **extreme working conditions**, as defined in the methodology. All simulation parameters and boundary conditions were adjusted to reflect the most challenging operational scenario, including reduced airflow, higher inlet air temperature, and increased IT load. The objective is to evaluate whether the cooling strategy can maintain all temperatures within safe operational thresholds and to identify potential hotspots that may arise.

Airflow Patterns and Mass Balance

As in the normal case, velocity streamline plots (see Figures 51 and 52) illustrate airflow behavior through the raised floor, perforated tiles, and alternating rack arrangement under extreme conditions.

The mass flow rates at all major boundaries were checked, and the net inflow/outflow was confirmed to be balanced (see Section 5.3.4), validating the CFD results.

Temperature Distribution and Hotspot Analysis

Table 22 presents the simulated area-weighted average inlet and outlet temperatures for the CRAH units and each rack during extreme operation:

Table 22. Simulated area-weighted average outlet temperatures for each rack and CRAH return during extreme operation.

Rack	**Outlet Temp (°C)**
R1	**58.47**
R2	46.64
R3	43.94

R4	46.06
R5	46.86
R6	**58.48**
R7	44.06
R8	46.05
R9	44.08
R10	46.64
CRAH1 return temperature	46.32
CRAH2return temperature	46.41
CRAH3 return temperature	47.04

Under extreme working conditions, initial CFD simulation results (using a total CRAH mass flow of **13.33 kg/s** and an inlet temperature of **23 °C** (value recommended by datasheet) showed thermal challenges.

The maximum rack outlet temperature (T_{max}) reached **58.47 °C** at R1 and R6 (see above Table), which is well above the defined hotspot threshold **of 50 °C**. This result indicates that the standard configuration cannot maintain safe operation and requires optimization.

Theoretical Cooling Calculation and Iterative Solution

To address these hotspots, the required mass flow rate for non-CDU racks (since CDU rack mass flow rate is already pre-set by the fans) was theoretically calculated using the equation:

Equation used to calculate the mass flow rate:

$$\dot{m} = \frac{Q}{C_p \,.\, \Delta T}$$

Where:

- $\dot{m}$: mass flow rate in kg/s
- Q : Heat load in $W = 26{,}000$ W.
- C_p : Specific heat capacity of air $= 1005 \frac{J}{kg}.K$

- ΔT : Temperature rise $= 50 - 23 = 27\ ℃$

With the above stated the calculated mass flow rate for non-CDU racks is: $\dot{m}$ **= 0.958 kg/s.** For partially air-cooled hybrid racks (17 kW), the required flow rate is **0.63 kg/s** per rack. With this in mind, iterative CFD simulations were conducted, increasing the CRAH mass flow rate stepwise.

Optimized Configuration and Final Results

After several iterations, increasing the total CRAH mass flow rate from 13.33 kg/s to 15.99 kg/s or 47,020 m³/h (5.33 kg/s per CRAH unit) resulted in:

- T_{max} (R1 and R6): 48.7 °C (below the 50 °C hotspot threshold).
- Average air return temperature: 43 °C (below the 45 °C limit).
- All rack outlets and returns are now within acceptable thermal limits.

Mass Flow Rates at CRAH and Rack Inlets

Table 23 summarizes the final mass flow rates at the **CRAH unit inlets** and each of the **10 server racks** for the optimized extreme scenario.

Table 23. Mass flow rates per rack for the optimized "Extreme case" scenario.

Rack	Mass Flow Rate (kg/s)
R1	0.998
R2	2.619
R3	0.958
R4	2.569
R5	0.979
R6	0.995
R7	2.607

R8	0.900
R9	2.610
R10	0.995

Temperature Distribution at Rack Outlets and CRAH Returns

The table below presents the area-weighted average static temperatures at the outlet of each rack and at the CRAH return outlets after optimization

Table 24. Area-weighted average static temperatures by rack and CRAH .

Rack	**Outlet Temp (°C)**
R1	49.73
R2	39.26
R3	35.65
R4	41.99
R5	39.95
R6	49.76
R7	39.78
R8	41.96
R9	41.90
R10	39.26

CRAH Unit	**Inlet Temp (°C)**	**Outlet Temp (°C)**
CRAH 1	23.50	42.90
CRAH 2	23.50	42.92
CRAH 3	23.50	43.02

These tabulated CFD results clearly demonstrate that, after optimizing the CRAH airflow, the data center can operate safely under extreme load conditions without exceeding critical temperature thresholds or causing hotspot formation.

Key Results and Compliance

- Maximum possible average air return temperature: 43 °C (*below the 45 °C limit*).
- Maximum outlet temperature: T_{max} at R1 and R6 = 48.7 °C (*below the 50 °C hotspot threshold*).
- All rack outlets and return air temperatures remain within safe operational limits.

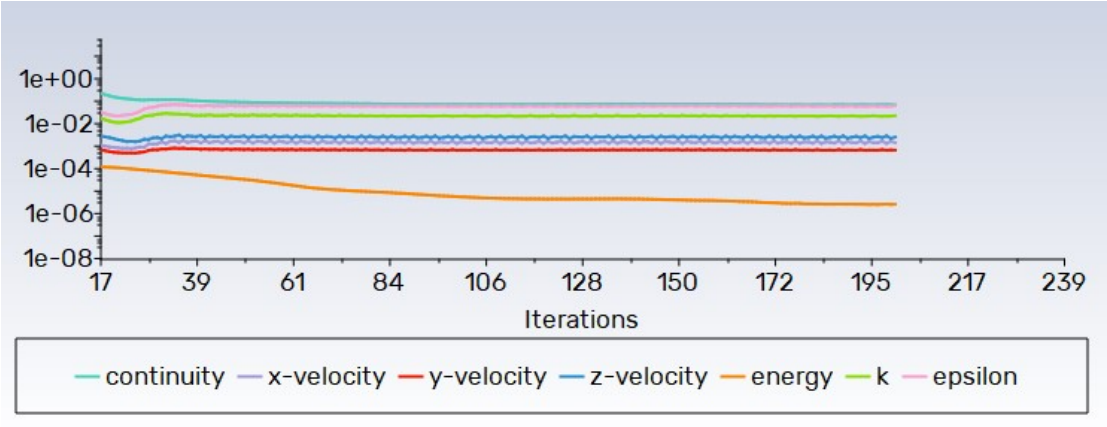

Figure 50. Scaled residuals graph for the simulation of The Extreme case.

Figure 50 presents the evolution of scaled residuals for all major solution variables—continuity, velocity components (x, y, z), energy, turbulence kinetic energy (k), and epsilon—across 239 solver iterations. The plot demonstrates steady convergence, with all residuals falling below the standard threshold of 1×10^{-3}, and the energy residual reaching values near 1×10^{-6}. This consistent reduction and stabilization of residuals confirm the robustness and numerical stability of the simulation results for the extreme operating scenario.

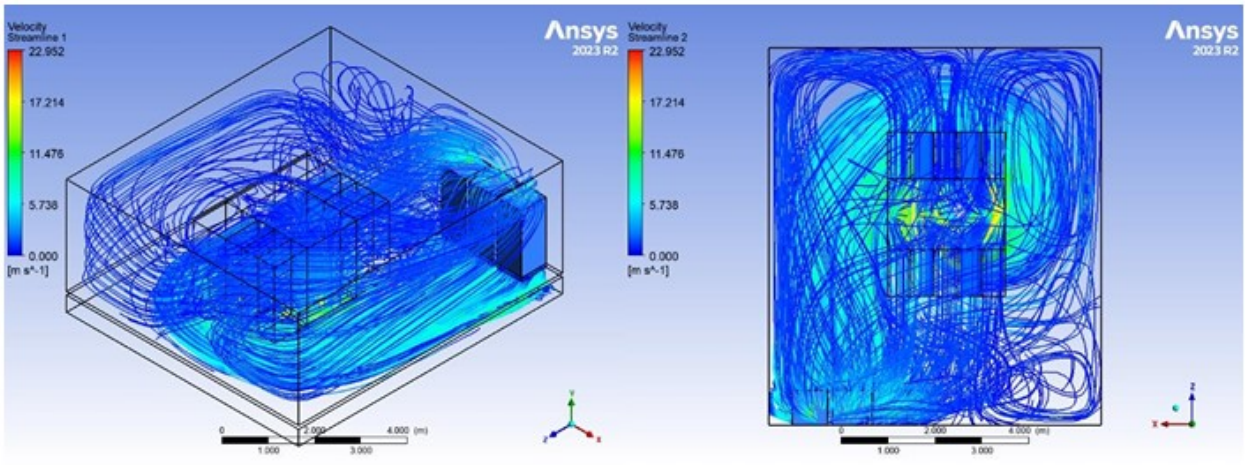

Figure 51. velocity streamlines from different views at Extreme Case.

Observing Figures 51, 52 the flow of air from CRAH units in the underfloor to pass through the racks and return to in CRAH back are shown. The flow is increased in the CDU racks due to the presence of the fans there, which could be seen by higher streamline concentrations, similar to the Normal Case, but at a higher velocity due to mass flow rate increase.

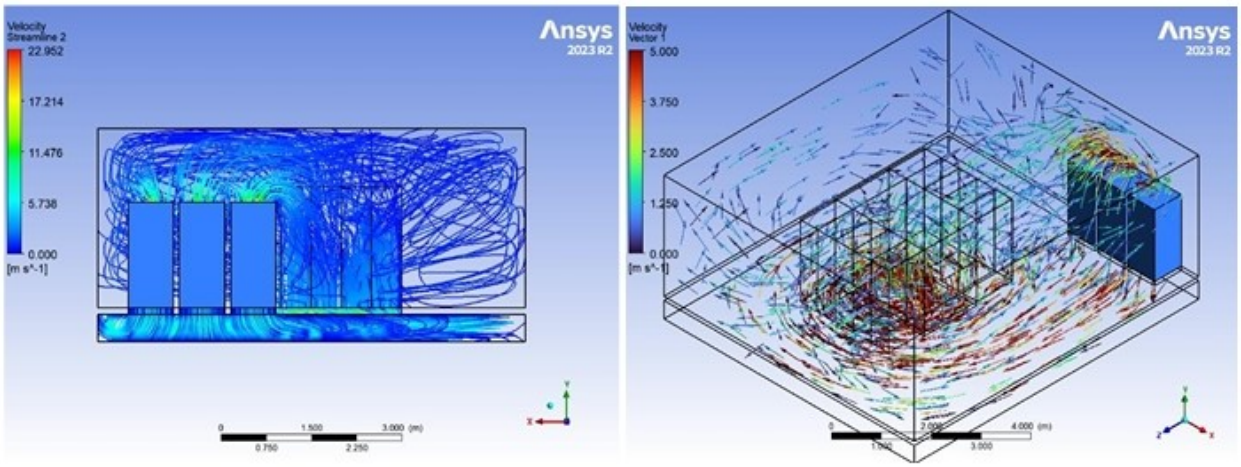

Figure 52. Velocity streamlines from different views at Extreme case.

Figure 53 displays the velocity contour just above the raised floor plenum (0.25 m), under extreme operating conditions. The visualization shows a pronounced velocity peak directly beneath the rack area, indicating the concentrated supply of cold air from the underfloor tiles.

The highest velocities are observed near the tile openings, confirming efficient delivery of air to the cold aisle.

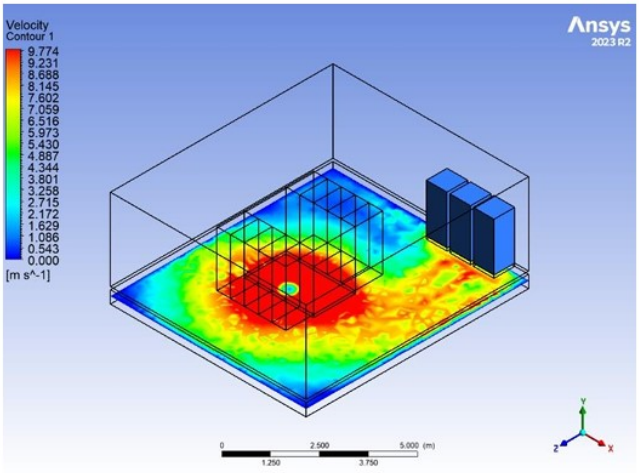

Figure 53. underfloor velocity contour at extreme case (0.25m height)

Moreover, Figure 54 presents the velocity distribution at the central plane of the data center (1.6 m above ground). The contour highlights elevated airflow within the cold aisle region, as well as the transition of air from supply to return pathways. The persistence of defined velocity channels at this height suggests effective upward movement of cold air towards the server inlets, with localized high-velocity zones directly aligned with the main rack rows.

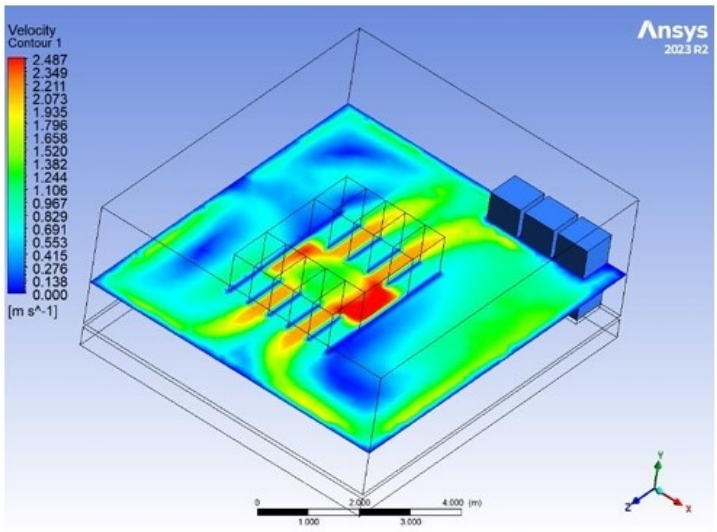

Figure 54. Velocity contour at mid-height (1.6m)

Finally, by observing the next figure (Figure 55) it could be understood that the velocity field near the ceiling, at the CRAH return level (2.5 m) is maximum (around 3 m/s per CRAH). The velocity is generally lower and more diffuse compared to the lower sections, reflecting the mixing and deceleration of air as it rises. The highest velocity region is concentrated near the CRAH return area, indicating efficient extraction (suction) of warm air from the data hall and its direction toward the cooling units.

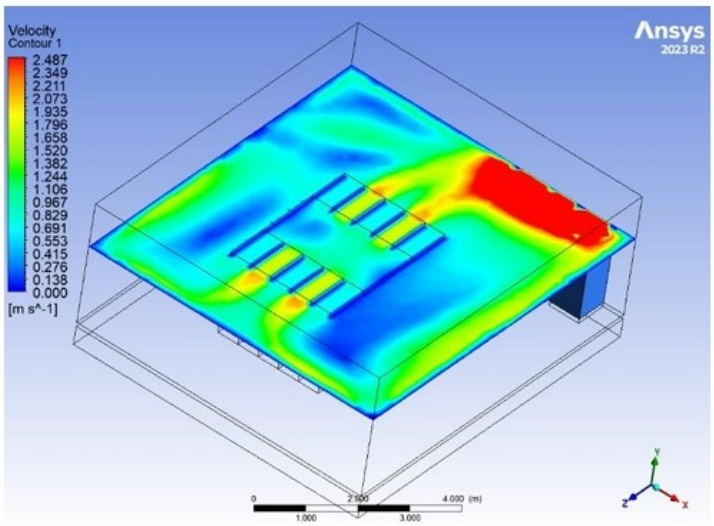

Figure 55. Velocity contour at the top-level of the CRAH return (2.5m from base ground)

Velocity Streamlines and Contours Analysis

The most critical observation from the velocity streamline figures is that there is no recirculation of air from the outlet to the inlet of the racks. This is primarily due to the presence of a pressurized underfloor plenum and a cold aisle containment system, ensuring that cold air is efficiently directed towards the IT equipment without unwanted mixing with the warmer return air.

Additionally, we observe a proper airflow pattern where the CRAH units supply cold air into the raised floor, which is then evenly distributed through the perforated tiles into the cold aisles. The return air is effectively sucked back by the CRAH units from above at approximately the same flow rates, maintaining a balanced and controlled cooling process.

Another significant observation is the airflow behavior through the CDU racks (R2, R4, R7, R9). These racks exhibit a more steady and higher flow rate, which can be attributed to the presence of dedicated fans within the CDU units. These fans generate a consistent pressure jump, as defined by the fan curve, ensuring stable and efficient cooling for the liquid cooled sections of the hybrid racks.

These findings confirm that the air distribution strategy in the data center is functioning effectively, minimizing hotspots and optimizing cooling performance.

Temperature Thermal Analysis Using Contours, Maps, and Post-CFD Results

To evaluate the thermal performance of the data center under extreme conditions, various temperature contours, thermal maps, and post-CFD visualizations have been analyzed. These results provide insights into hotspot formation, airflow distribution, and overall cooling efficiency, ensuring the system operates within the designed temperature limits.

The first temperature rendering was utilized to generate a temperature map across the entire room volume, allowing for the identification of the highest temperature (T_{max}) and its exact location.

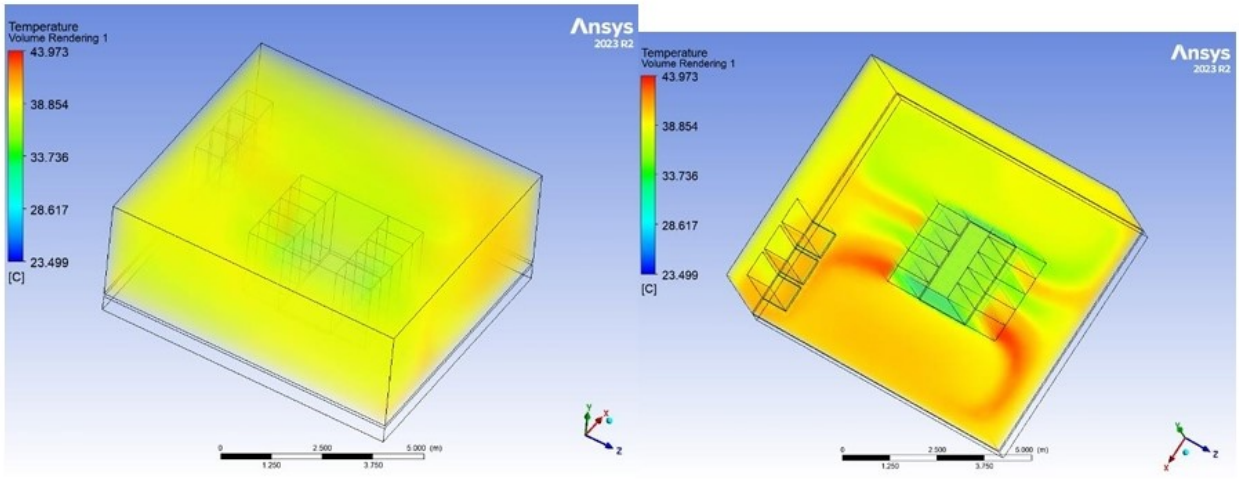

Figure 56. temperature rendering of the volume (extreme case).

T_{max} is recorded at 43.973°C, observed at the exit of racks 1 and 6. However, these are not considered hotspots as the temperature remains within the acceptable limit of 45°C.

To further analyze the higher temperature zones within the data center, a high-temperature map has been generated and is presented below. As previously stated, the elevated temperature regions are primarily concentrated around the exit zones of racks R1 and R6, confirming the expected thermal behavior in these areas. This visualization provides a clearer understanding of the airflow patterns and heat dissipation across the room.

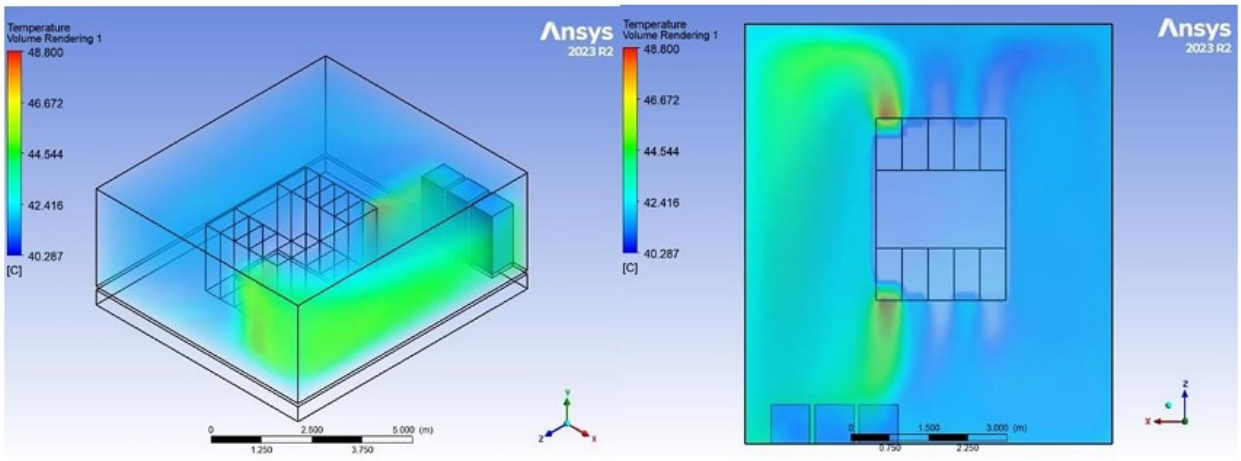

Figure 57. High temperature maps (extreme case).

This high-temperature map (Figure 57) further confirms the thermal distribution within the data center, aligning with the expected exit temperatures at R1 and R6.

Moving on to Figure 58, it presents temperature contours at the rack exits along the Y-axis. Here, the temperature at the hottest rack exhausts again approaches 48 °C, while adjacent racks show lower values, generally in the range of 34–42 °C. This cross-sectional view demonstrates the non-uniformity of the thermal distribution, making it possible to visually compare the effect of rack position and load on temperature rise within the aisle.

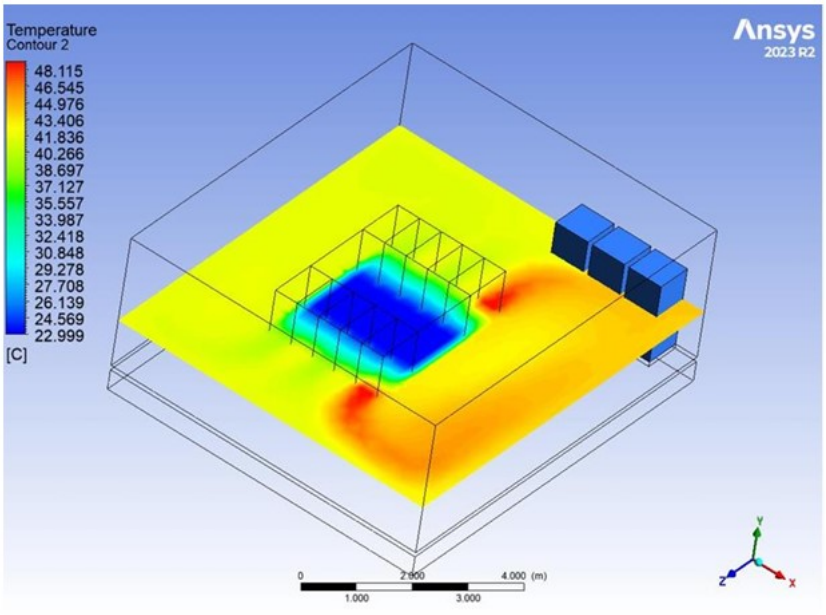

Figure 58. Temperature contours at the rack exits along the Y-direction (Extreme Case).

Figure 59 shows the temperature distribution at the rack exits along the Z-direction. The highest recorded temperatures are concentrated at the central rack rows (comparable to the Normal case), especially in the region spanning R1 to R6, with maxima near 48 °C. Away from these racks, temperatures decrease progressively, with side zones maintaining values as low as 23–28 °C. This indicates that the extreme case produces significant thermal gradients across the data center.

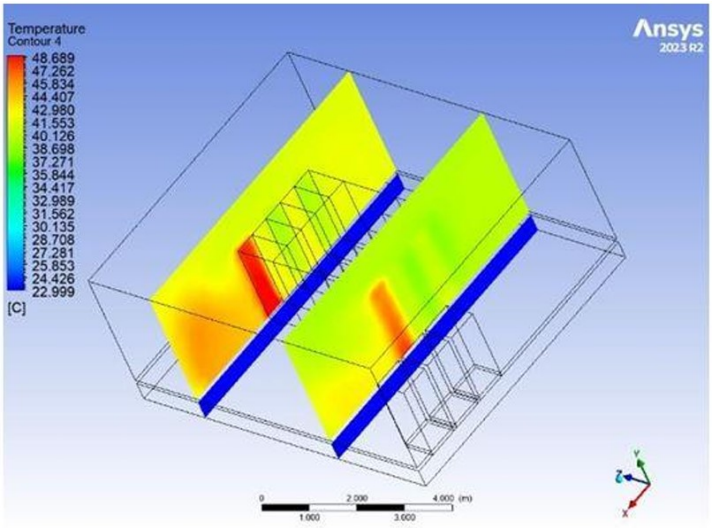

Figure 59. Temperature contours at the rack exits along the Z-direction (Extreme Case).

Temperature Contour Analysis and Hotspot Mitigation

The temperature contour figures, presented from multiple viewpoints, vividly illustrate the thermal distribution within the data center model. These contours enable a detailed assessment of airflow patterns, temperature gradients, and the identification of critical heat accumulation zones. In the initial simulation, high-temperature regions were evident at racks R1 and R6, where local temperatures exceeded the defined hotspot threshold of 50 °C, indicating a risk to equipment reliability.

To address these hotspots, the mass flow rate supplied by the cooling system was increased. The resulting updated temperature contours reveal a markedly more uniform temperature distribution across the entire data center. The previously problematic hotspots at R1 and R6 were effectively eliminated, with maximum temperatures at these critical locations reduced to within acceptable operational limits. As a result, the cooling efficiency of the system improved significantly, with all areas maintaining temperatures below the established threshold and the return air temperature stabilized near 45 °C.

These findings confirm that optimizing airflow rates is a highly effective strategy for enhancing thermal management and preventing localized overheating in high-density data center environments. Such targeted interventions not only safeguard IT equipment but also contribute to the overall energy efficiency and reliability of the facility.

7. Power Usage Effectiveness (PUE) and Economic Analysis

7.1. Power Usage Effectiveness (PUE) calculations

Starting with the assumptions taken for the Power Usage Effectiveness and Economic Analysis assumption taken we have:

Assumptions for Economic Analysis

Electricity Price:
The average electricity price for non-household consumers in Italy is taken as €0.1867 per kWh, based on the latest data from the European Commission (June 2024). This value is used consistently throughout the economic analysis.

Continuous Operation:
It is assumed that the data center operates continuously at a stable power level, 24 hours per day, 365 days per year (8,760 hours annually). There are no scheduled shutdowns, partial load periods, or seasonal variations in IT or cooling loads.

Isolation from External Environment:
The data center room is considered fully isolated, with no thermal interaction with the external environment. Consequently, the cooling load is assumed to be independent of seasonal changes (i.e., there is no difference in cooling performance or energy use between summer and winter).

Additional Common Assumptions for Data Center Economic Analysis

- Equipment Efficiency: All cooling equipment (CRAH, CDU, fans, pumps) is assumed to operate at the efficiency levels specified by the manufacturer, with no degradation due to aging, fouling, or part-load operation.
- No Unexpected Downtime: It is assumed that there are no unplanned outages, major failures, or maintenance shutdowns affecting the IT or cooling systems during the year.

- Constant Load Profile: IT and cooling loads are considered constant at the design value for the entire year. Variability due to user demand, batch jobs, or maintenance windows is neglected.
- Utility Grid Stability: The electricity supply is assumed to be stable, with no interruptions or voltage fluctuations that could impact energy consumption or operational performance.
- Cooling System Sizing: Cooling system sizing is based on peak design load and does not account for potential future expansions or load increases
- Ambient Conditions: Since the data center is isolated, ambient external temperature and humidity are assumed to have no impact on internal environmental conditions.
- Exclusion of Additional Loads: Only IT equipment and primary cooling system power consumption are considered. Additional loads such as lighting, security systems, or office support equipment are not included in the energy or cost calculations.

Annual Energy Consumption Calculation:

Annual energy consumption is calculated as follows:

Annual Consumption (kWh) = Total Power (kW) × 24 × 365

This approach assumes constant power draw throughout the year.

Where:

- **Normal Case:**

$$4{,}817.6kW \times 8{,}760hours = 42{,}202{,}937.6kWh/year$$

- **Extreme Case:**

$$4{,}591.04kW \times 8{,}760hours = 40{,}215{,}510.4kWh/year$$

- **Electricity Cost Calculation:**

$$Annual\ Cost\ (€) = Annual\ Consumption\ (kWh) \times Electricity\ Price\ (€\ /\ kWh)$$

PUE Calculation Summary:

By applying the definition of Power Usage Effectiveness (PUE), as introduced in Section 2.12 (Energy Efficiency Metric), the PUE for each scenario was calculated. The results are summarized in the following table:

Table 25. PUE calculation results for the Extreme and Normal cases.

	Normal Operation	**Extreme Operation**
IT Load (kW)	4,480	4,480
CDU Power (kW)	57.6	38.4
CRAH Power (kW)	208.32	72.64
Total Facility Power (kW)	4,745.92	4,591.04
PUE	**1.06**	**1.021**

Interpretation of PUE (ASHRAE, 2020):

- PUE = 1.0 → Perfect Efficiency (ideal, unattainable)
- PUE < 1.5 → Highly Efficient Data Center
- PUE ≈ 1.6 – 2.0 → Typical Efficiency
- PUE > 2.0 → Inefficient (high cooling and power losses)

With the above interpretation described by ASHRAE, both the "Normal" and "Extreme" cases could be listed as highly efficient data center designs. Moreover, the next table would show the percentage improvement of PUE when employing the "Extreme" case instead of the "Normal".

Table 26. Comparison of PUE values for normal and extreme cases, with calculation of percentage improvement achieved in the extreme scenario.

	Normal	Extreme	% Improvement
PUE	1.06	1.021	3.68%

The implementation of optimized airflow management and targeted cooling adjustments led to a measurable improvement in the data center's energy efficiency. Specifically, the Power Usage Effectiveness (PUE) was reduced from 1.06 under normal operating conditions to 1.021 under the optimized (extreme) scenario, representing a 3.68% increase in overall efficiency. This improvement indicates that a greater proportion of total facility power is being utilized directly by IT equipment, with less energy expended on cooling and auxiliary systems. As a result, the

data center not only meets thermal requirements but also achieves substantial reductions in operational costs and environmental impact, supporting long-term sustainability goals.

7.2. Economic Analysis: Power Consumption and Cost Savings

Annual Consumption and Cost Analysis

To fully assess the benefits of the optimized cooling strategies, it is essential to complement the technical results with an economic analysis. Evaluating annual energy consumption and corresponding operational costs provides a clearer perspective on the real-world impact of improvements made to data center efficiency. By comparing both the normal and optimized (extreme) operating scenarios, this analysis quantifies the monetary savings that can be achieved through reduced power usage. Understanding these cost implications is crucial for data center operators, as energy expenses represent a significant portion of total operating costs. The following section details the methodology and results of the annual consumption and cost analysis for each scenario.

Table 27. Comparison of annual power consumption and electricity costs for the full-scale data center under normal and extreme cooling operation.

	Normal	**Extreme**
Total Power (kW)	4,817.60	4,591.04
Annual Energy (kWh)	42,202,938	40,215,510
Electricity Price (€ / kWh)	0.1867	0.1867
Annual Cost (€)	7,882,950	7,508,190

Table 28. Annual cost savings.

	VALUE (€)	**% SAVINGS**
COST REDUCTION (NORMAL VS. EXTREME)	374,760	4.75%

Key Points

- Optimized extreme-case cooling reduces annual energy costs by €374,760, representing a 4.75% savings versus normal conditions.
- CFD-guided airflow management and system tuning directly translate to measurable financial and environmental benefits.
- This level of efficiency is only possible with careful simulation, monitoring, and ongoing optimization.

The economic analysis, summarized in Tables 27 and 28, demonstrates that implementing the optimized, extreme-case cooling scenario leads to substantial cost savings for the data center. Specifically, annual energy expenditure is reduced by approximately €374,760, representing a 4.75% decrease compared to normal operations. These results highlight how CFD-guided airflow management and system tuning can directly translate technical improvements into measurable financial and environmental benefits. Notably, achieving this level of efficiency requires the use of advanced simulation techniques, ongoing monitoring, and continuous optimization, underscoring the importance of data-driven strategies in modern data center management.

8. Conclusion

This thesis presents a comprehensive investigation into optimizing the thermal management and energy efficiency of high-density data centers, using advanced Computational Fluid Dynamics (CFD) modeling, detailed scenario analysis, and real-world validation. The work is motivated by the rapid expansion of global data center capacity, rising IT power densities, and the critical need to mitigate thermal hotspots while reducing operational costs.

The study began with an extensive literature review, identifying key cooling challenges such as recirculation, thermal stratification, and hotspot formation—common in legacy systems with inadequate airflow management. Prior research highlighted the effectiveness of cold/hot aisle containment, optimized rack arrangements, and data-driven airflow strategies, which have demonstrated energy savings of 10–35%. These insights guided the development of both the baseline and retrofit scenarios.

A validated CFD model of the CQ-University data center (Hassan et al., 2013) was developed using ANSYS Fluent 2023 R2, accurately replicating airflow and temperature distributions based on published experimental data. The baseline simulation, lacking containment, revealed a maximum rack outlet temperature of ~29 °C and significant mixing between hot and cold air, especially in high-load areas.

To improve performance, insulated cold aisle containment was introduced. CFD results showed a 33.8% reduction in maximum temperature (T_{max}), bringing it to ~21 °C—well within recommended IT equipment limits. This confirmed the elimination of recirculation and hotspots, aligning closely with ENERGY STAR and industry benchmarks. The impact of this improvement was further quantified through two energy-saving strategies: (1) raising the CRAH inlet temperature and (2) reducing total CRAH mass flow rate. Both were simulated with the goal of reaching a room return temperature similar to the baseline T_{max} (~29 °C). Results showed that increasing the CRAH inlet temperature from 15 °C to 18 °C reduced annual energy costs by €26,812.80 (12.27%), while reducing mass flow from 14.33 kg/s to 3 kg/s yielded savings of €25,372.20 (11.62%). Both maintained safe thermal conditions and proved effective in the retrofitted setup.

To support future high-density applications, the model was further upgraded to simulate hybrid air-liquid cooling for denser racks. Several layout configurations and supply airflow scenarios were analyzed. The alternating high-/low-density rack arrangement produced the most uniform temperature distribution, with $T_{max} \approx 41.14$ °C, while individual hybrid racks managed air-cooled loads up to 19.8 kW and total rack loads up to 132 kW. Carefully lowering total supply airflow to ~9.5 kg/s still ensured thermal safety while reducing energy usage.

The optimized configuration was then scaled to a high-performance computing (HPC) facility with a total IT load of 4.48 MW (16 pairs of 280 kW racks). Using equipment specifications from Vertiv's Hi-rating and GRS tools, the cooling infrastructure was sized and simulated for both normal and extreme operational scenarios. Under normal conditions, the system achieved a PUE of 1.06, meeting all thermal targets without hot-spots. Under extreme loads, an iterative adjustment of airflow and CRAH inlet temperature (increased to 23 °C) kept T_{max} under 50 °C and average return temperatures below 45 °C, complying with ASHRAE TC 9.9 guidelines and reaching approximately the threshold (reaching 49 °C for T_{max} and 43 °C at the return temperature of CRAH).

Economically, these enhancements proved substantial. The optimized extreme-case scenario achieved a PUE of 1.021—3.68% better than the normal case—translating to ~€374,760 in annual energy savings based on 2024 electricity tariffs for non-household consumers in Italy. This represents a 4.75% reduction in operational expenses—an important margin as energy prices rise and sustainability becomes increasingly critical.

Beyond technical and financial benefits, this work underscores the value of CFD as a predictive, high-resolution tool for design and decision-making in modern data centers. Unlike empirical methods, CFD enables scenario testing, design validation, and performance forecasting under variable load and environmental conditions, eliminating trial-and-error in costly physical deployments.

In summary, this thesis validates the technical and economic advantages of CFD-driven optimization for data center cooling and offers a practical framework for engineering next-generation, energy-efficient digital infrastructure. By integrating advanced simulation tools, best-practice containment, hybrid cooling strategies, and validated manufacturer data, this work provides actionable guidance for achieving energy efficiency, equipment reliability, and sustainability.

9. Future Recommendations and Energy Reuse Potential

As data centers continue to grow in both capacity and energy demand, future planning must not only focus on optimizing internal thermal performance, but also on transforming waste heat into a valuable resource. One of the most impactful strategies is the reuse of heat rejected by IT equipment, particularly high-density racks, for external or internal applications. Rather than rejecting this thermal energy to the environment, it can be recovered and utilized for district heating systems, industrial processes, or even building-level heating and domestic hot water production. This form of energy cascading represents a key enabler of the circular energy economy and supports long-term carbon neutrality goals.

The effectiveness of heat recovery largely depends on the temperature and quality of the waste heat. While traditional air-cooled systems often operate at relatively low exhaust temperatures that are challenging to reuse efficiently, the adoption of liquid cooling technologies, specifically liquid–liquid Cooling Distribution Units (CDUs), enables higher outlet temperatures and more efficient heat transfer. These CDUs remove heat directly from hotspots using liquid coolant and transfer it to a secondary loop, which can be connected to a heat exchanger for energy reuse. In addition to reducing reliance on chilled water systems, this method significantly improves thermal density, allowing for greater cooling with less space and energy input.

Furthermore, incorporating AI-driven control strategies into the thermal management framework offers additional opportunities for optimization. Machine learning algorithms can be trained on real-time and historical data to predict computational workloads, anticipate thermal loads, and adjust cooling setpoints dynamically. This results in more responsive and granular control, minimizing overcooling and underutilization of cooling capacity. In scenarios involving free cooling or heat recovery, AI can also optimize thermal routing and load shifting, ensuring that the timing and location of heat rejection align with external reuse needs.

From a systems integration perspective, future sustainable data centers should be designed with thermal energy harvesting interfaces, modular CDU configurations, and AI control layers from the outset. Coupling these technologies will not only reduce total Power Usage Effectiveness

(PUE) but also contribute to a new metric of Carbon Usage Effectiveness (CUE), enabling measurable reductions in emissions. As global regulations and energy costs evolve, these advancements will likely become essential, not optional, for data centers seeking to operate at scale while meeting environmental standards.

APPENDIX A

CRAH cooling unit specifications at Normal Operation:

PW70GEHP3

Open Door Extended Height,Downflow Down,HT Coil,EC Fan Advanced - HP

Unit inlet air temperature	39.0	°C	Fluid	ETHYLENE GLYCOL 30%	
Unit inlet air relative humidity	30.0	%	Inlet fluid temperature	20.0	°C
Unit airflow	58569	m³/h	Outlet fluid temperature	32.0	°C
ESP	20	Pa	Unit fluid flow	6.56	l/s
Sea level	0	m	Unit power supply	400 V/3 ph/50 Hz	

Unit performances

Unit	PW70GEHP3		Unit power input	13.02	kW
Net total cooling capacity	289.8	kW	Unit Net Sens EER	22.30	
Net sensible cooling capacity	289.8	kW	Unit Net Total EER	22.30	
nSHR	1.00		Internal filter class (EN16890 std)	ePM10 50%	
Gross total cooling capacity	302.8	kW	Unit air pressure drop	422	Pa
Gross sensible cooling capacity	302.8	kW	Internal filter air pressure drop	61	Pa
Supply air temperature	23.6	°C	Coil air pressure drop	219	Pa
Supply air relative humidity	71.8	%	Depth	1050	mm
Room SPL (@ 2m, f.f)	71	dB(A)	Width	3200	mm
			Height	2350	mm
			Weight	1225	kg

CW Coils

Quantity	1	n°	Fluid pressure drop coil+connections	383	kPa
Unit fluid flow	6.56	l/s	Valve pressure drop	9	kPa
Unit fluid side pressure drop	392	kPa			
Regulation valve model	R2050-40-S4				

CW Fans

Quantity	4	n°	Operating Ampere	4 x 4.90	A
Type	EC Fan Advanced - HP		Full load Ampere	4 x 5.6	A
Power supply	400 V/3 ph/50 Hz		Locked rotor Amp.	4 x 0.1	A
Power input	4 x 3.23	kW	Room fan modulation (%)	97.0	%

Table A. 1. CRAH unit Normal operational conditions specification, supplied by Vertiv's Hirating tool.

CRAH cooling unit specifications at Extreme Operation:

PW70GEHP3

Open Door Extended Height,Downflow Down,HT Coil,EC Fan Advanced - HP

Unit inlet air temperature	45.0	°C	Fluid	ETHYLENE GLYCOL 30%	
Unit inlet air relative humidity	30.0	%	Inlet fluid temperature	20.0	°C
Unit airflow	39166	m³/h	Outlet fluid temperature	32.0	°C
ESP	20	Pa	Unit fluid flow	6.31	l/s
Sea level	0	m	Unit power supply	400 V/3 ph/50 Hz	

Unit performances

Unit	PW70GEHP3		Unit power input	4.54	kW
Net total cooling capacity	286.5	kW	Unit Net Sens EER	59.70	
Net sensible cooling capacity	271.2	kW	Unit Net Total EER	63.10	
nSHR	0.95		Internal filter class (EN16890 std)	ePM10 50%	
Gross total cooling capacity	291.0	kW	Unit air pressure drop	208	Pa
Gross sensible cooling capacity	275.7	kW	Internal filter air pressure drop	30	Pa
Supply air temperature	23.1	°C	Coil air pressure drop	111	Pa
Supply air relative humidity	98.0	%	Depth	1050	mm
Room SPL (@ 2m, f.f)	61	dB(A)	Width	3200	mm
			Height	2350	mm
			Weight	1225	kg

CW Coils

Quantity	1	n°	Fluid pressure drop coil+connections	357	kPa
Unit fluid flow	6.31	l/s	Valve pressure drop	8	kPa
Unit fluid side pressure drop	365	kPa			
Regulation valve model	R2050-40-S4				

CW Fans

Quantity	4	n°	Operating Ampere	4 x 1.79	A
Type	EC Fan Advanced - HP		Full load Ampere	4 x 5.6	A
Power supply	400 V/3 ph/50 Hz		Locked rotor Amp.	4 x 0.1	A
Power input	4 x 1.11	kW	Room fan modulation (%)	69.0	%

Table A. 2. CRAH unit Extreme operational conditions specification, supplied by Vertiv's Hirating tool.

APPENDIX B

Dell PowerEdge XE 8640 Technical Specifications:

Feature	Technical Specifications**	
Processor	Two 4th Generation Intel Xeon Scalable processors with up to 56 cores per processor Two 5th Generation Intel Xeon Scalable processors with up to 64 cores per processor	
Memory	• 32 DDR5 DIMM slots, supports RDIMM 4 TB max, speeds up to 5600 MT/s • Supports registered ECC DDR5 DIMMs only	
GPU	• 4 NVIDIA HGX H100 80GB 700W SXM5 GPUs, fully interconnected with NVIDIA NVLink technology	
Storage controllers	• Internal controller: PERC H965i • Internal boot: Boot Optimized Storage Subsystem (NVMe BOSS-N1): HWRAID 1, 2 x M.2 SSDs • Software RAID: S160	
Drive bays	Front Bays • Up to 8 x 2.5-inch SATA/SAS/NVMe SSD max 122.88 TB • Up to 8 x E3.S NVMe SSD max 122.88 TB	
Power Supplies	2800W Titanium 200-240 VAC or 240 VDC, redundant, hot swap 3200W Titanium 277 VAC or 260-400 VDC redundant, hot swap	
Cooling Options	• Air cooling for Processors • Liquid Assisted Air Cooling for GPU **Note:** No facility water to rack required.	
Chipset	Intel® C741 chipset	
PCIe	2 CPU configuration: Up to 4 PCIe slots (4 x16 Gen5)	
Embedded NIC	2 x 1 GbE	
Network Options	1 x OCP 3.0 (x8 PCIe lanes)	
Fans	• Up to 6 Standard (STD) fans installed in mid tray • Up to 5 High performance (HPR) gold grade fans installed on the front of the system • All are hot swap fans	
Dimensions and weight	• Height — 174.3 mm (6.86 inches) • Width — 481.91 mm (18.97 inches) • Depth — 901.4 mm (35.48 inches) with bezel — 865.54 mm (34.07 inches) inch without bezel • Weight — 61.4 kg (135.36 pounds)	
Form Factor	4U rack server	
Embedded Management	• iDRAC9 • iDRAC Direct • iDRAC RESTful API with Redfish • iDRAC Service Module	
Bezel	Optional LCD bezel or security bezel	
OpenManage Software	• CloudIQ for PowerEdge plug in • OpenManage Enterprise • OpenManage Enterprise Integration for VMware vCenter • OpenManage Integration for Microsoft System Center	• OpenManage Integration with Windows Admin Center • OpenManage Power Manager plugin • OpenManage Service plugin • OpenManage Update Manager plugin
OpenManage Integrations	• BMC Truesight • Microsoft System Center • OpenManage Integration with ServiceNow	• Red Hat Ansible Modules • Terraform Providers • VMware vCenter and vRealize Operations Manager
Security	• Cryptographically signed firmware • Data at Rest Encryption (SEDs with local or external key mgmt) • Secure Boot • Secured Component Verification (Hardware integrity check)	• Secure Erase • Silicon Root of Trust • System Lockdown (requires iDRAC9 Enterprise or Datacenter) • TPM 2.0 FIPS, CC-TCG certified, TPM 2.0 China NationZ
Ports	Front ports • 1 x iDRAC Direct (Micro-AB USB) port • 1 x USB 2.0 • 1 x VGA	Rear ports • 1 x USB 2.0 • 1 x USB 3.0 • 1 x VGA • 1 x RJ45 iDRAC9 ethernet port
Operating Systems and Hypervisors	• Canonical Ubuntu Server LTS • Red Hat Enterprise Linux For specifications and interoperability details, see Dell.com/OSsupport.	
OEM-ready version available	From bezel to BIOS to packaging, your servers can look and feel as if they were designed and built by you. For more information, visit Dell.com > Solutions > OEM Solutions.	

NOTE: ** indicates Additional features Coming Soon.

Discover more about PowerEdge servers

Figure B1. Excerpt from the Dell PowerEdge XE8640 rack server datasheet, showing technical specifications relevant to this study (Dell Technologies, 2024).

Inside the system:

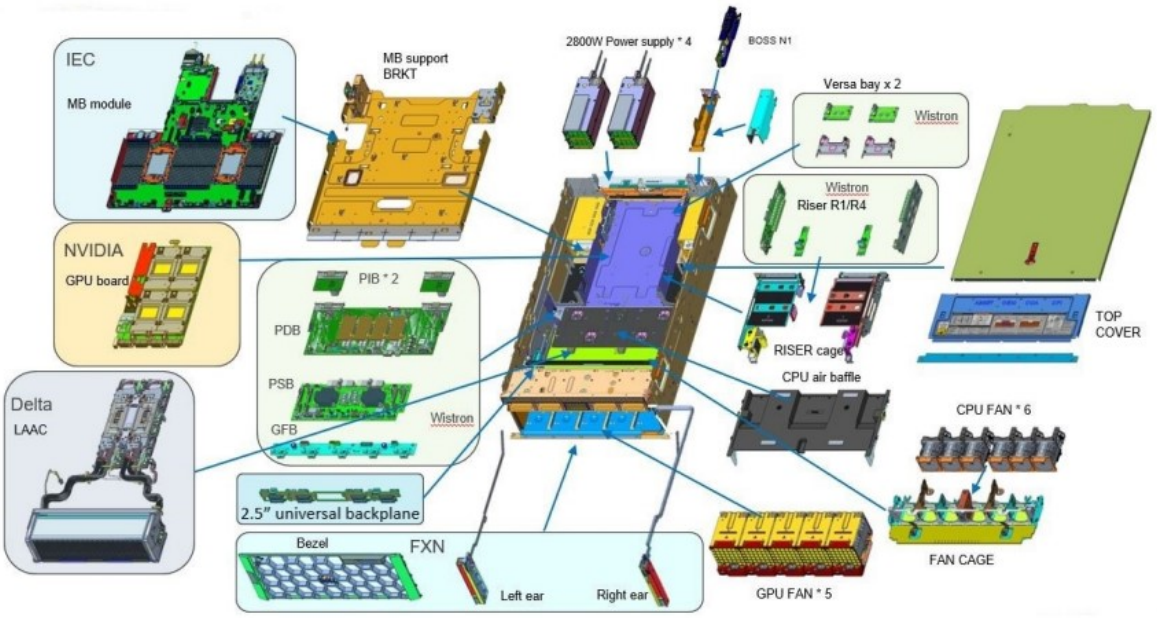

Figure B2. Schematic diagram of a Dell PowerEdge rack server, illustrating key components relevant to the IT load modeling in this study.

References:

Ahmed, M., et al. (2021). *Challenges in real-world CFD simulations for data center cooling*. *International Journal of Thermal Sciences, 65*, 211–223.

Ahmed, M., Singh, R., & Sharma, P. (2021). *Dynamic cooling systems for data centers: Challenges and opportunities*. *Energy Engineering Journal, 48*(1), 25–35.

Almoli, A., De La Cruz, I., Thompson, H., & Raghavan, V. (2012). *Computational fluid dynamics modelling of high-density data centres: Airflow analysis for heat management*. *Applied Energy, 98*, 39–48. https://doi.org/10.1016/j.apenergy.2012.03.038

American Society of Heating, Refrigerating and Air-Conditioning Engineers. (2012). *Air cooling of computer equipment* (Datacom Series). ASHRAE.

American Society of Heating, Refrigerating and Air-Conditioning Engineers. (2012). *Datacom equipment power trends and cooling applications* (2nd ed.). ASHRAE.

American Society of Heating, Refrigerating and Air-Conditioning Engineers Technical Committee 9.9. (2018). *Data-centre power-equipment thermal guidelines* (5th ed.). ASHRAE.

American Society of Heating, Refrigerating and Air-Conditioning Engineers Technical Committee 9.9. (2021). *Thermal guidelines for data processing environments* (5th ed.). ASHRAE.

ANSYS. (2013). *ANSYS Fluent theory guide* (Release 15.0). Retrieved from http://www.pmt.usp.br/academic/martoran/notasmodelosgrad/ANSYS%20Fluent%20Theory%20Guide%2015.pdf

ANSYS. (2014). *ANSYS Fluent theory guide* (Release 15.0). ANSYS Inc.

Balaras, C. A., & Argiriou, A. A. (2002). *Infrared thermography for building diagnostics*. *Energy and Buildings, 34*(2), 171–183.

Bash, C., Patel, C., & Sharma, R. (2015). *Data-center cooling configurations: CFD-based assessment of efficiency and reliability*. Hewlett Packard Enterprise.

Cho, J., & Kim, B. S. (2011). *Evaluation of air management system's thermal performance for superior cooling efficiency in high-density data centers*. *Energy and Buildings, 43*(8), 2145–2155. https://doi.org/10.1016/j.enbuild.2011.04.025

Cho, J., & Kim, K. (2009). *Thermal analysis of a high-density data centre for minimising the occurrence of hot spots*. *Energy Conversion and Management, 50*(6), 1478–1488. https://doi.org/10.1016/j.enconman.2009.03.006

Dell Technologies. (2024). *Dell PowerEdge XE8640 server: Technical specifications* [Data sheet]. https://www.dell.com/support/manuals/xe8640

ENERGY STAR. (2019). *Move to a hot-aisle / cold-aisle layout* [Case study]. U.S. Environmental Protection Agency.

ENERGY STAR. (2024). *Data center energy efficiency strategies*. U.S. Environmental Protection Agency. https://www.energystar.gov/

European Commission. (2022). *2022 Best practice guidelines for the EU Code of Conduct on Data Centre Energy Efficiency*. Joint Research Centre. https://e3p.jrc.ec.europa.eu/publications/2022-best-practice-guidelines-european-code-conduct-data-centres

Geng, Q., Li, M., & Wu, J. (2022). *Integrating renewable energy into data center cooling systems: A life-cycle assessment*. *Renewable Energy Journal, 125*, 457–471.

Hassan, N. M. S., Khan, M. M. K., & Rasul, M. G. (2013). *Temperature monitoring and CFD analysis of data centre*. *Procedia Engineering, 56*, 551–559. https://doi.org/10.1016/j.proeng.2013.03.159

International Organization for Standardization. (2018). *ISO 50001:2018 – Energy management systems — Requirements with guidance for use*. https://www.iso.org/standard/69426.html

Jones, P., Geng, Q., & Zhao, L. (2022). *Environmental impacts of data center cooling systems: A global perspective*. *Environmental Research Letters.*

Kapadia, N., Malik, A., & Zhou, Q. (2021). *Dynamic thermal load management using smart cooling systems*. *IEEE Transactions on Sustainable Computing, 6*(4), 455–467.

Kingma, D. P., & Welling, M. (2014). *Auto-encoding variational Bayes*.

Liu, X., Zhang, T., & Chen, Y. (2019). *Airflow optimization in raised-floor data centers*. *International Journal of HVAC&R Research, 25*(2), 143–159.

Nakao, H., Yamada, T., & Suzuki, S. (2019). *Optimization of data-center airflow management using CFD*. *Applied Thermal Engineering, 150*, 1022–1032.

Ozen, M. (2014). *Meshing workshop*. Ozen Engineering Inc.

Rahman, S. (2021). *Optimizing thermal reliability in Tier III data centers through ASHRAE guideline compliance*. *International Journal of Data Center Management, 9*(2), 45–57.

Sadrehaghighi, F. (2020). *A comparison of meshing strategies for complex data center CFD models*. *Journal of Thermal Science and Engineering Applications, 12*(6), 061019. https://doi.org/10.1115/1.4047542

Schmidt, R. R., & Iyengar, M. K. (2018). *Airflow configurations for high-density data-center clusters*. *Applied Thermal Engineering, 92*, 278–289.

Shehabi, A., Smith, S. J., Hubbard, A., Newkirk, A., Lei, N., Siddik, M. A. B., Holecek, B., Koomey, J. G., Masanet, E. R., & Sartor, D. A. (2024). *2024 United States Data Center Energy Usage Report* (Report No. LBNL-2001637). Lawrence Berkeley National Laboratory. https://doi.org/10.71468/P1WC7Q

Sullivan, D., Rambo, J., & Pulvino, L. (2008). *High-density data centres: Case studies and best practices*. Uptime Institute.

Tang, J., et al. (2020). *CFD-based optimization of airflow distribution in raised-floor data centers*. *Energy Efficiency Journal, 23*(4), 267–276.

Uptime Institute. (2022). *Tier Standard: Topology*. https://uptimeinstitute.com/Tier

Wong, E. T. T., Chan, M. C., & Sze, L. K. W. (2015). *Spreadsheet modeling of data center hotspots*. In M. Khosrow-Pour (Ed.), *Computer engineering: Concepts, methodologies, tools, and applications* (pp. 1207–1219). IGI Global. https://doi.org/10.4018/978-1-4666-5888-2.ch115

Yamada, K., et al. (2021). *Optimizing data center design and cooling efficiency using CFD simulations*. *Journal of HVAC&R Research, 27*(2), 45–58.