# CSAID Roadmap | Data Center Cooling

Muhammad Atiq Siddiqui

CSAID Roadmap Meeting | Data Center Cooling

June 20th, 2024

# Fermilab Data Centers - Current State

Fermilab's Data Centers are currently located in 2 primary facilities

- Grid Computing Center (GCC)
  - Compute, Networking and Tape Libraries
  - Clean UPS power and refrigerant based air cooling
  - No generator backup - End of Life: 2030
- Feynman Computing Center (FCC)
  - Misc. IT, Critical Servers, Networking, Storage and Tape Libraries
  - Clean UPS power and refrigerant based air cooling
  - Generator backup available - End of Life: 2035

# Capacity

| Total Capacity | Power (UPS) | Cooling | Racks | Rack Density |
|---|---|---|---|---|
| GCC (CRA, CRB, NRA, NRB, TRR) | ~1.62 MW | ~1.7 MW (n+1) | ~160 | 10 KW/Rack |
| FCC (FCC2 and FCC3) | ~1.2 MW | ~1.1 MW (n+1) | ~200 | 5 KW/Rack |

# Goal

The primary goal of our data center facility planning process is to create a sustainable, net-zero, and scalable data center facility infrastructure that supports the evolving computational demands of Fermilab’s scientific mission.

**Strategic Goals:**

- Reduce energy consumption and costs
- Increase cooling capacity to support higher rack power density
- Improve reliability and uptime
- Enhance sustainability and reduce environmental impact

# Key Objectives

**Sustainability:** Minimize the environmental impact of our data center by achieving net-zero emissions through a combination of energy-efficient technologies, renewable energy sources, and innovative cooling solutions.

**Scalability:** Design a flexible infrastructure that can adapt to the growing computational needs of scientific experiments like LHC Run 3, mu2e, SBND, DUNE, and HL-LHC, as well as accommodate future advancements in computational research

**Adaptability:** Continuously assess and update our facility planning strategy to align with the latest developments in scientific research, industry trends, and technological advancements.

**Long-term Vision:** Develop a forward-thinking strategy that extends beyond specific end dates, ensuring our data center remains a valuable resource for the scientific and industrial communities for years to come.

# Roadmap Overview

## High-Level Roadmap

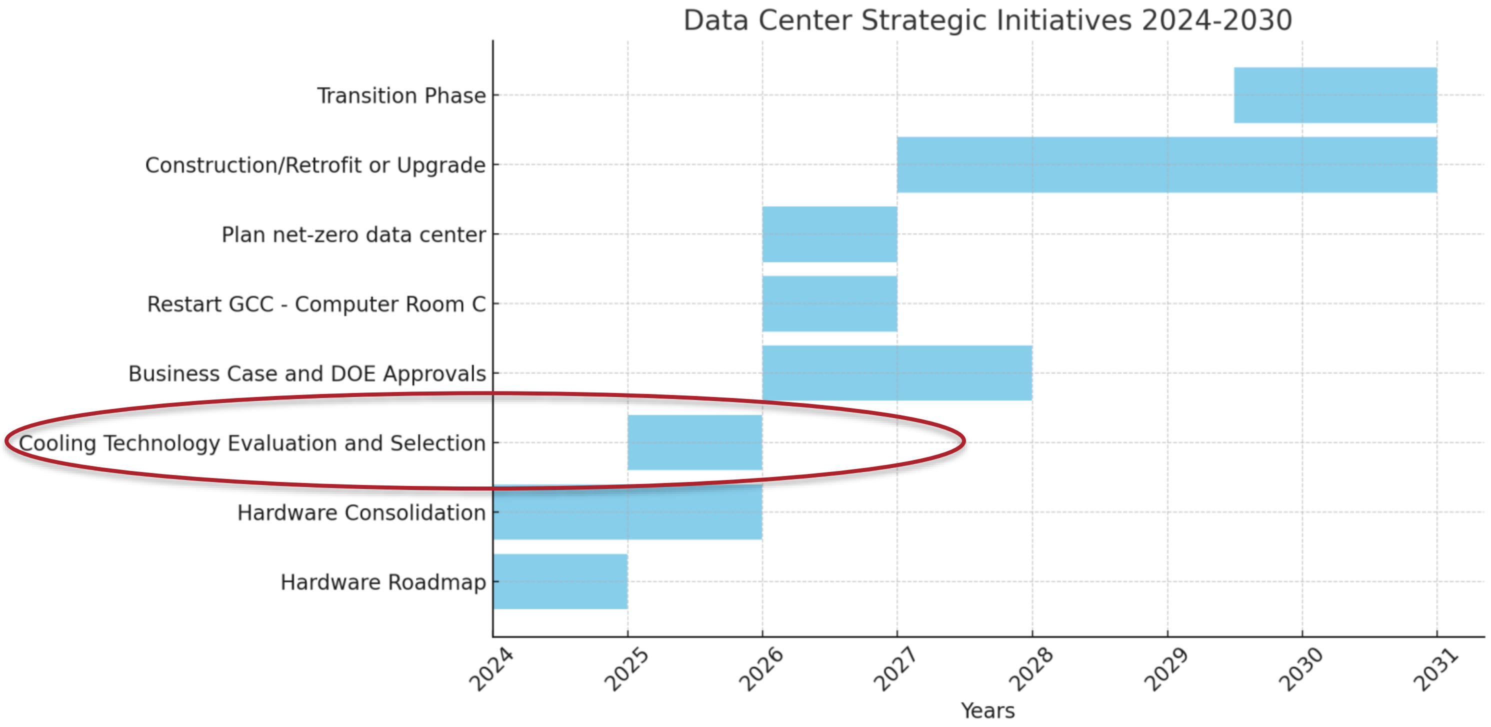

# Data Center Cooling – Importance

## Importance of Efficient Data Center Cooling

- **Critical for Equipment Reliability:** Prevents equipment malfunction and failure due to overheating.
- **Maximizes Performance:** Ensures optimal operating temperatures for maximum computational output.
- **Extends Equipment Lifespan:** Reduces wear and tear caused by heat stress.
- **Energy Efficiency:** Lowers operational costs and environmental impact.

# Data Center Cooling – Challenges

**Increasing Heat Density:** High-density compute environments generate more heat, pushing traditional air-cooling to its limits.

**Innovative Cooling Solutions:** Adoption of liquid cooling and immersion cooling for more effective heat dissipation.

**Cost:** Implementation of advanced cooling technologies can be costly

# Moore's Law and the Thermal Challenge

- Moore's Law states that the number of transistors in a dense integrated circuit doubles approximately every two years.
- This exponential growth in transistor density leads to increased processing power, but also results in higher power consumption and heat generation.
- Modern CPUs generate a significant amount of heat, especially in high-performance data center environments.

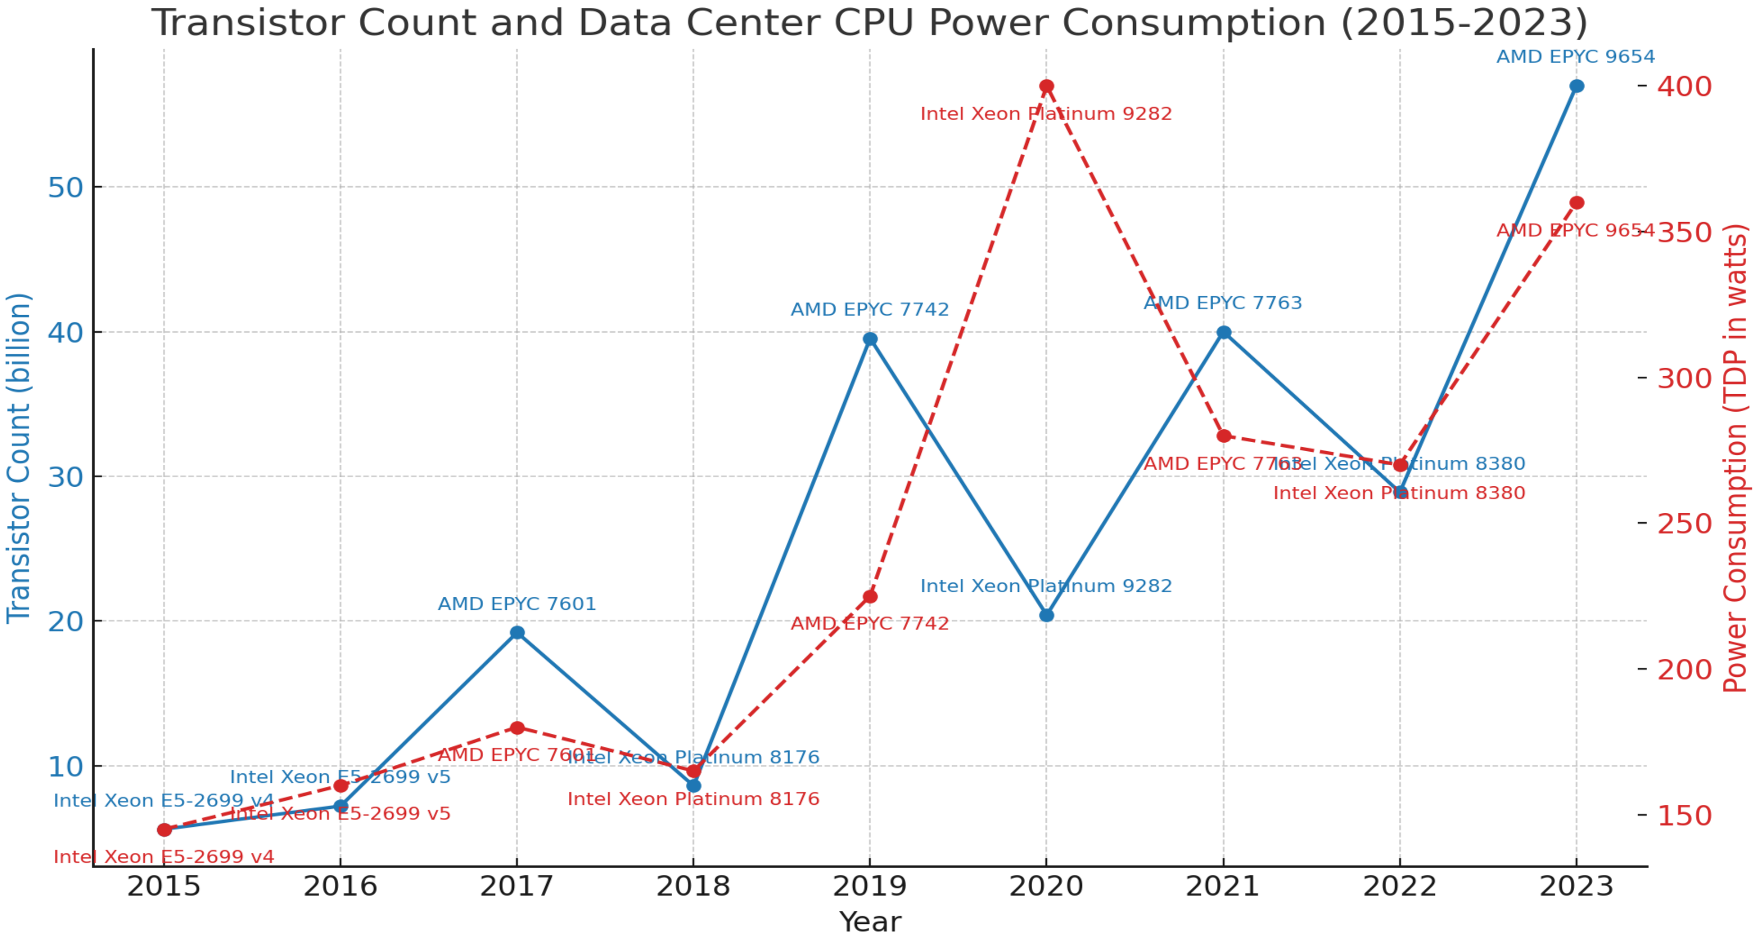
Transistor Count and Data Center CPU Power Consumption (2015-2023)
Transistor Count (billion)
Power Consumption (TDP in watts)
Year
10
20
30
40
50
150
200
250
300
350
400
2015
2016
2017
2018
2019
2020
2021
2022
2023
Intel Xeon E5-2699 v4
Intel Xeon E5-2699 v5
AMD EPYC 7601
Intel Xeon Platinum 8176
AMD EPYC 7742
Intel Xeon Platinum 9282
AMD EPYC 7763
Intel Xeon Platinum 8380
AMD EPYC 9654

# Why Air Cooling is Hitting Limits:

- Traditional air cooling methods are becoming less effective in managing the thermal demands of high-performance CPUs/GPUs.
- Air cooling struggles to dissipate heat effectively in dense server racks due to limited airflow and space constraints.
- Relying on fans for air cooling can lead to excessive noise levels, which is undesirable in data centers.

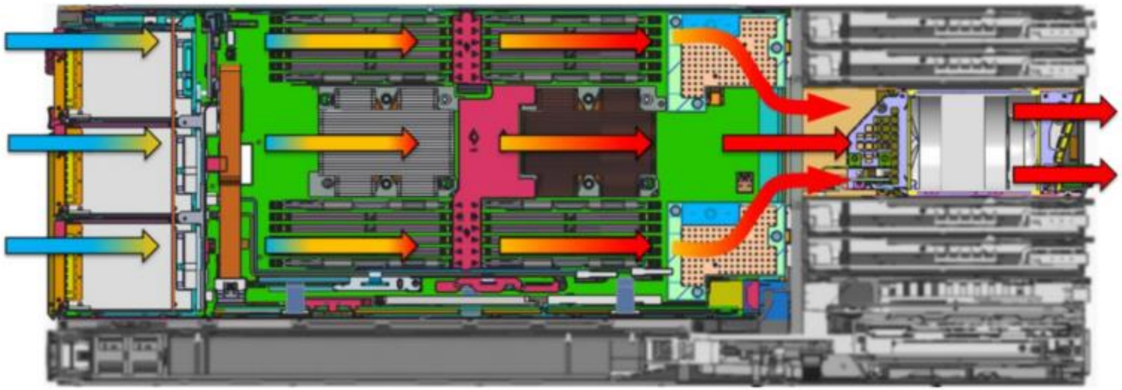

# Data Center CPU Comparison (2021-2024)

| CPU Family | Year Released | Typical TDP Range (W) | Cooling Requirement |
|---|---|---|---|
| **Intel Xeon** | | | |
| Ice Lake (3rd Gen) | 2021 | 150-270 | Air/Liquid |
| Sapphire Rapids (4th Gen) | Early 2024 | 120-350 | Air/Liquid |
| Emerald Rapids (4th Gen Refresh) | Late 2024/Early 2025 | 120-350 | Air/Liquid |
| **Granite Rapids (5th Gen)** | **2025** | **~500** | **Likely Liquid** |
| **AMD EPYC** | | | |
| Milan (3rd Gen) | 2021 | 120-280 | Air/Liquid |
| Genoa (4th Gen) | Late 2022 | 170-360 | Air/Liquid |
| Bergamo (4th Gen) | Late 2022 | 170-360 | Air/Liquid |
| Genoa-X (4th Gen) | Early 2024 | 200-360 | Air/Liquid |
| **Turin (5th Gen)** | **2024** | **~400** | **Likely Liquid** |

# Data Center GPU Comparison (2013-2024)

| GPU Family (Nvidia) | Year Released | Typical TDP Range (W) | Cooling Requirement |
|---|---|---|---|
| Tesla K40 | 2013 | 235 | Air |
| Tesla K80 | 2014 | 300 | Air |
| Tesla P100 | 2016 | 300 | Air |
| Tesla V100 | 2017 | 300 | Air |
| Tesla T4 | 2018 | 70 | Air |
| Tesla A100 | 2020 | 400 | Air/Liquid |
| H100 | 2022 | 700 | Liquid |
| B200 | 2024 | 1000 | Liquid |

# The Liquid Cooling Solution

- Liquid cooling offers a more efficient and scalable solution to address the thermal challenges of modern data centers.
- By utilizing liquids with high thermal conductivity, like water or specialized coolants, liquid cooling systems can transfer heat away from CPUs much more effectively than air.
- There are several types of liquid cooling systems available, each with its own advantages and considerations.

# Direct-to-Chip (D2C) Cooling

A cold plate is attached directly to the CPU die, allowing coolant to flow through microchannels and absorb heat directly from the source

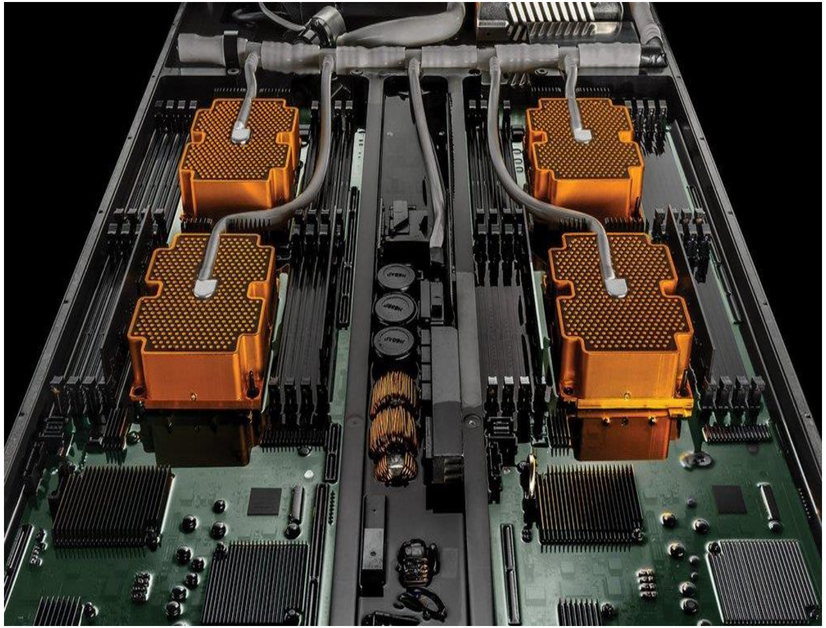

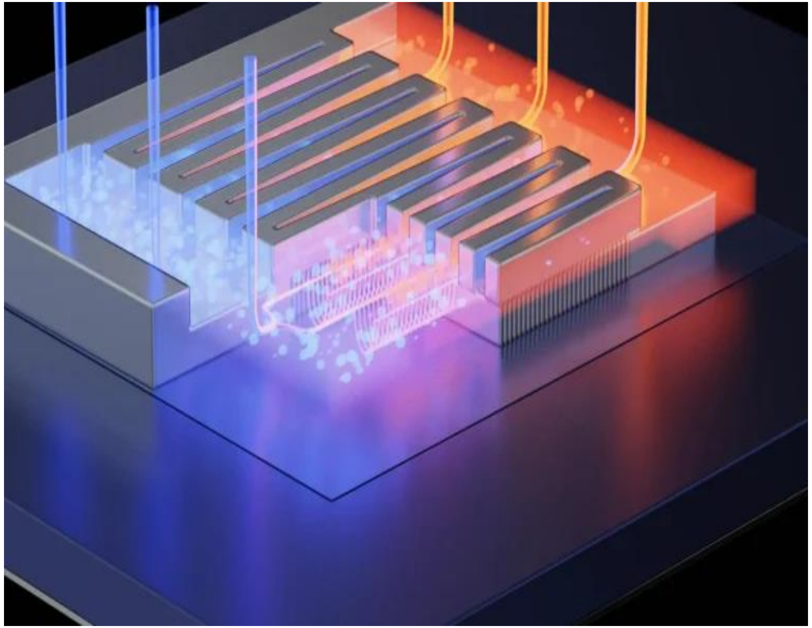

Microfluid channels in a semiconductor chip

# Direct-to-Chip (D2C) - Pros and Cons

Pros:

Very efficient cooling method, Scalable, enables higher clock speeds and overclocking potential, support for high density rack configurations 15 KW – 100 KW.

Cons:

complex installation, may require specialized motherboards and CPUs, potential for leaks, special consideration to rack dimensions for the installation of manifolds.

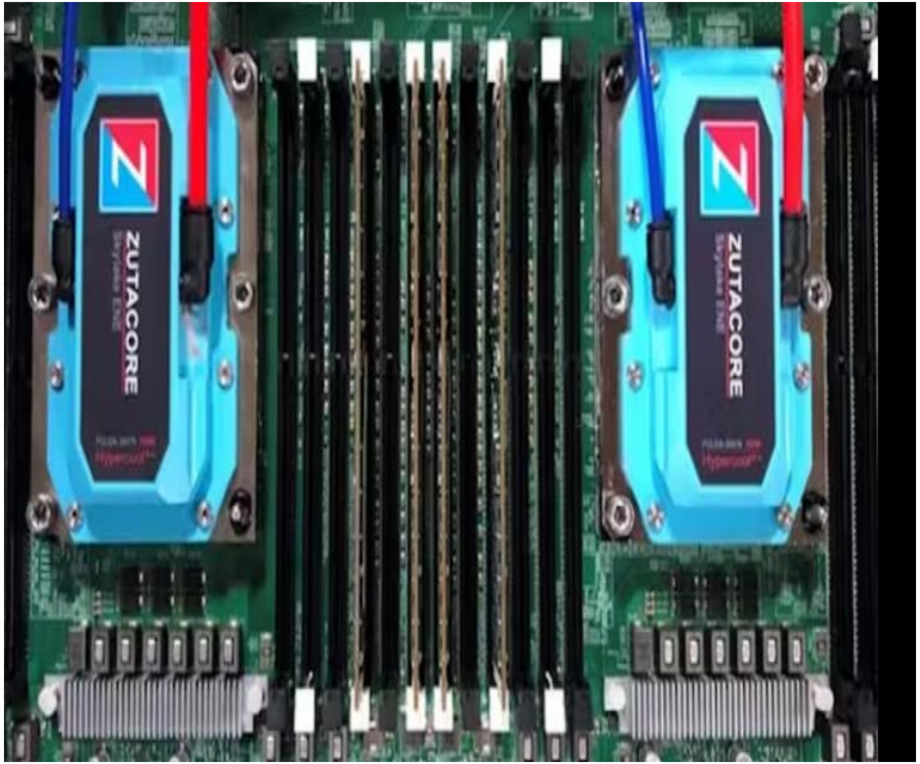

# Immersion Cooling

Submerges the entire server or individual components in a non-conductive dielectric fluid that absorbs heat.

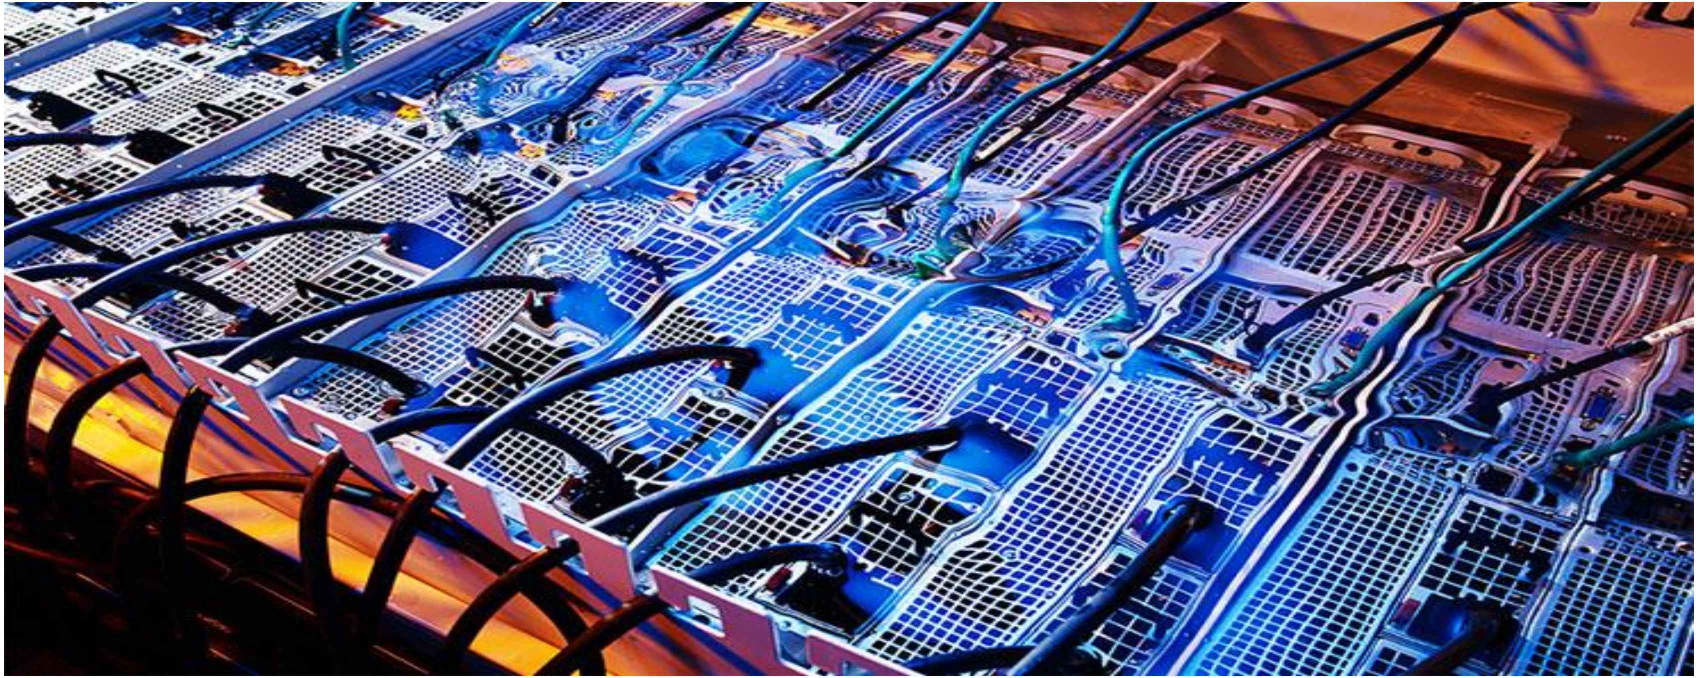

# Immersion Cooling – Pros and Cons

**Pros:** Extremely efficient cooling, silent operation, high power density potential, reduced hardware maintenance, could support 20 – 200 KW per rack.

**Cons:** High upfront cost, requires specialized tanks and fluids, potential for leaks, may require specific hardware and cables compatibility, Operational complexity.

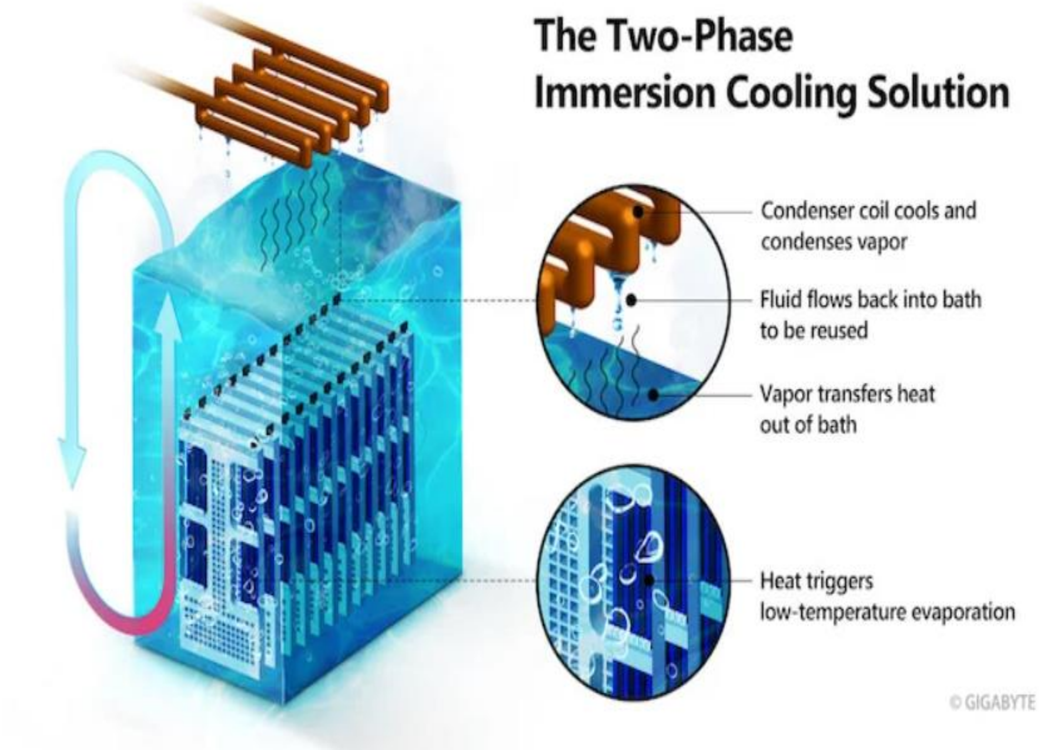

# Rear-Door Heat Exchanger (RDHX) Cooling

Attaches to the rear of server racks and uses coolant to absorb heat from the exhaust air.

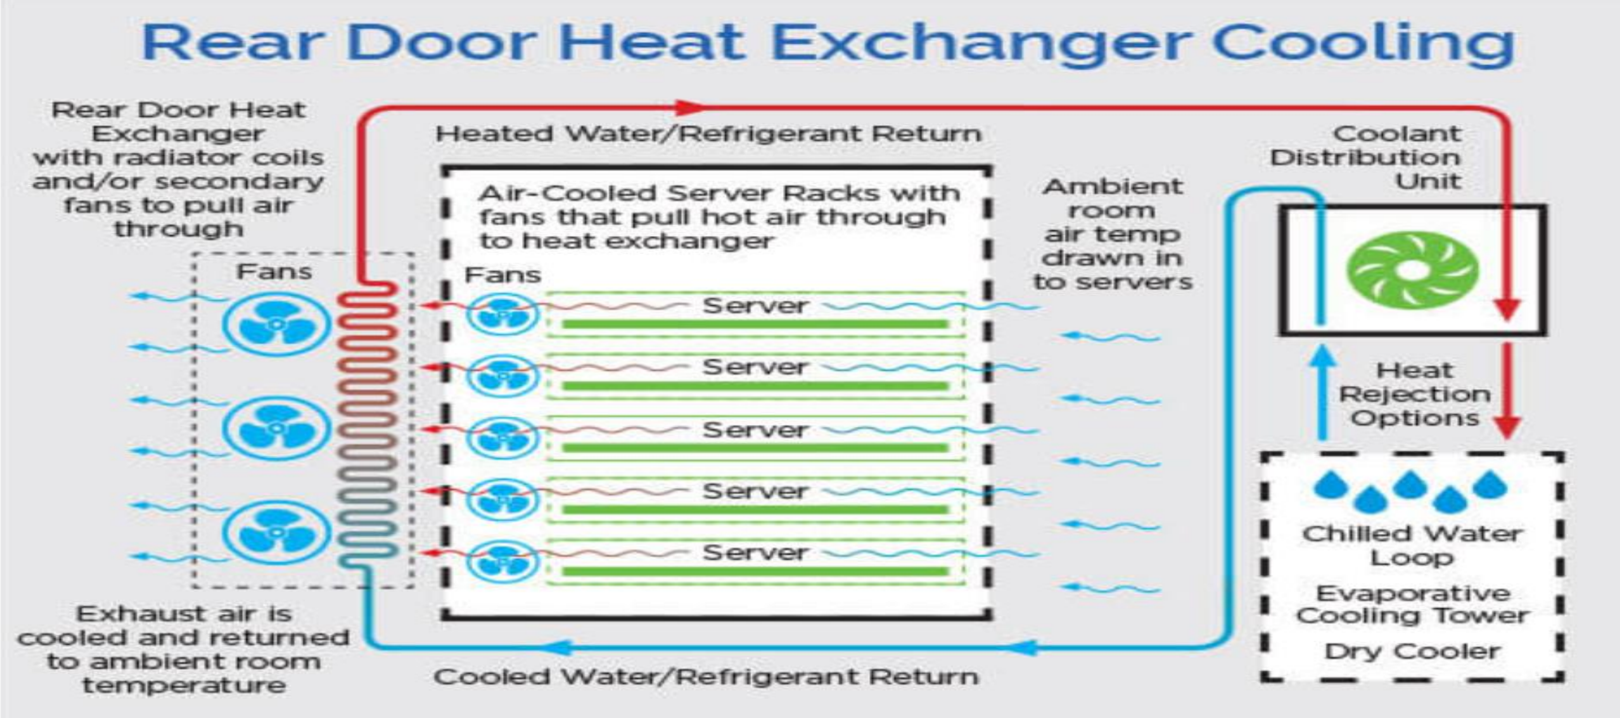

# Rear-Door Heat Exchanger Cooling – Pros and Cons

**Pros:** Easier to retrofit into existing data centers as a replacement of rear rack doors, less invasive than D2C or immersion cooling.

**Cons:** Less efficient than D2C or immersion cooling, limited cooling capacity for high-density racks. 15-25 KW cooling capacity.

*Note: Cost to deploy 10 racks to support 25KW per rack with RDHX is approx. USD 1.2M in GCC-CRC with no existing chilled water infrastructure in place.*

# Liquid Cooling Comparison

| Parameter | Rear-Door Heat Exchanger (RDHX) | Direct-to-Chip (D2C) | Immersion Cooling |
|---|---|---|---|
| Heat Load per Rack (kW) | 15-20 kW | 15-100 kW | 20-100+ kW |
| Cost | Low | Moderate | High |
| Installation Complexity | Low | Moderate | High |
| Scalability | Limited | Moderate | High |
| Efficiency | Low | High | Highest |
| Sustainability | Moderate | High | Highest |
| Operational Ease | High | Moderate | Low |

# Key Steps

- Need Assessment:  CPU vs GPU ? Hardware Strategy  – in Progress
- Cost/Budget
- Choosing the right Liquid Cooling Technology
    - Evaluate coolant options
- Space/Infrastructure requirements
- Research and vendor selection
- Pilot Test
- Plan for Implementation
- Monitor and Optimize

# Questions ?