OXFORD JOURNAL
OF STUDENT
SCHOLARSHIP

Cooling the Cloud - A Study on Innovative Approaches to Cooling AI Data Centers

Adalynn Le
adalynn.k.le@gmail.com

ABSTRACT

Data centers that provide computational power for Artificial Intelligence struggle with resource intensive practices. The cooling systems used for large-scale and commercial AI are mostly liquid-based and utilize large portions of water daily. When paired against a growing demand for water provided by a rising population and inequity crisis, the deficiency causes a substantial issue for general water accessibility. The following paper presents examples and explanations of simple cooling techniques (fan, heat sink, heat pump, and vapor chamber), primarily limited to those used in PCs as well as a comparative literature-based evaluation framework to modify these techniques to be more widely applicable. The information presented here is not a result of a concrete experiment, instead a comparative evaluation framework and process. Although it is true that up-scalling PC-systems for industrial use is novel, the following paper uses pre-existing knowledge to present a process through which such a breakthrough could occur by interpreting and novelizing existing works. The ideal goal of this paper is to present a research plan based on existing data that would yield a product capable of reducing the environmental footprint of AI data centers. While the data in this paper is **not** original (all sources cited), it is a compilation and summation of preexisting data in a novel and condensed format that presents a new use for it.

INTRODUCTION

As AI continues to gain popularity and become more widely used, the data centers that they rely on “may soon consume six times more water than Denmark” (*UNEP*). As population grows and pressure on clean-water strengthens, it becomes overly pressing that something is done to address this issue. In order to prepare a safety net for water scarcity, a reliable plan would be reducing the amount of water used in AI. The root reason as to why AI “guzzles” so much water is because of data centers. These centers are organized so that there is a “cold aisle” where the server draws in cool air, and a “hot aisle” where exhaust is vented” (see Fig. 1) (Barringer, *& The West*). The “heat” is created by an immense reliance on energy and electricity. In order to quickly and efficiently generate this much electricity, computers, such as those found in data centers or PCs, generate heat energy. However, this amount of heat can sabotage safety and performance, meaning that cooling is necessary. In large, industrial areas, the most common form of cooling is liquid cooling, or using water: “Red Canary magazine recently reported that the water demand of nearly 60 centers in Phoenix is about 177 million gallons a day” (Barringer, *& The West).*

Data centers are the backbone for information in AI. Software such as ChatGPT, Copilot, and Gemini are able to generate their answers by quickly sifting through information. The centers contain the bulky and non-portable hardware that computers and AI rely on, including, “the computing infrastructure that IT systems require, such as servers, data storage drives, and network equipment” (*Amazon*). Effectively, AI is heavily reliant on a water-using system. Without this, AI would cease to function, something that would disrupt the flow of modern life as well as disrupting business the economy. Due to the ramifications and inability to remove or lessen the use of data centers nor AI, the main focus for improving the environmental sustainability should cater towards lessening the total environmental footprint of data centers.

CURRENT WORK

Liquid cooling is currently the most common in data centers, though many also adopt hybrid cooling. There is a variety of reasons for this, though the underlying theme is that-in large scale operations-liquid-cooling is simply more convenient. Many smaller centers or personal PC’s use fan-cooling. However, this is not practical for a large-scale industry, which could provide a reason as to why liquid-cooling is seen as a better option. To use fan-cooling would require the establishment of a long array of costly fans that operate on electricity. Although this is not optimal in terms of cost, design, and infrastructure, it may provide environmental benefits. Other types of cooling include semi-fan cooling and non-fan-cooling. Fan cooling is relatively efficient, with a Cooling Efficiency Ratio (CER) typically ranging from 5-15. While this provides a considerably more energy-efficient solution, implementing large scale fans into a data center will be inconvenient. Some data centers also already implement hybrid systems, as many computers have fans, although it is not to a level of environmental efficiency quite yet.

Fanless cooling, despite being less efficient than Fan-cooling, could prove itself to be more effective in this sort of situation. Although a lot of the solutions are currently energy inefficient, they may be more convenient than building an array of fans at a data center. Different types of these exist, and only a few of many are explained here. The first, a heat sink, involves effectively “draining” heat from the CPU and moving it into a thermally conductive metal, such as aluminum or copper. The process for this is slow, and the CER is low at around 1-3, although it may be combined with other types of cooling. The next is heat pipes, which is similar in concept to heat sinks. However, instead of using conductive metals, it uses “a heat transfer liquid, which evaporates by removing latent heat. The heat transfer fluid that becomes vapor carries the heat to the cooler end, thereby releasing heat” (*DNP*). A similar solution, known as a vapor chamber, employs the same technology while using a chamber instead of a pipe. This is the most energy efficient with a CER around 20. The final solution this paper will address is a Peltier (PEL-tee-ay) module. As the name suggests, the Peltier module takes inspiration from the Peltier Effect (also known as the thermoelectric effect). This technique “corresponds to the heat extraction or absorption occurring at the contact between two different conducting media when a direct current (DC) electric current flows through this contact” (Gurevich & Velasquez Perez). Essentially, this means that when connecting two electrical wires in a specified way, you can end up with one “cool” end and one “hot” end to transfer energy. Peltier modules take advantage of this by constructing a panel with two sides. One end, the cool one, is attached to the module that needs to be cooled. The metal parts inside transfer heat to the other end. Peltier modules are small and compact, making them relatively easy to install. That being said, the modules are high maintenance and are inefficient with energy, meaning that only abtout 1/10 of energy input is used for cooling. Due to how small-scale and inefficient Peltier modules are, I would not be inclined to study them in my project. A comparison of PC cooling systems can be found in Figure 2.

All solutions addressed above are common in PC's or personal units. None have been adapted for a large-scale use, such as the kind used in data centers, although the purpose of this experiment and process is to adapt such technologies to large-scale industry based on efficiency.

SOLUTION

After analyzing and evaluating the sources listen and research, the conclusion I have reached is that a potentially feasible solution process could be as follows: A research-process that uses pre-existing data and achievable research to synthesize a process for cooling. Such a process would be divided into 3 sections:

1) Testing the different types of cooling (fan, heat sink, heat pipe, vapor chamber) on PC's of different scale and recording major variables being:
    a) Energy Efficiency (measured in CER).
    b) Cooling Power (measured in Cubic Feet per Minute for fans or heat removed in general).
    c) Environmental Consequence
        i) Since this data is qualitative, the most efficient way to calculate this would be through further research of the techniques. Due to the present problem, it is widely known and accepted that techniques that rely on water-use would have very low environmental efficiency, whereas air-based and energy efficient techniques would be more environmentally-safe.
    d) Costliness (costs will be noted and compared)
    The suggested testing process would involve the implementation of cooling techniques to PCs of varying size and make in order to simulate the effects on a real-world data center. In order to optimize efficiency, the best course of action would include gathering a sample of about 3-5 PCs and applying the cooling techniques while noting the metrics listed above. Each technique will be applied to the same PCs (in order to ensure a repeatable, fair, and standardized experiment). In order to prevent bias, different trials on different types of PCs will be used as well.
2) Researching and analyzing these cooling systems to see how they work and devising a solution to get the top contenders in step 1 to be accessible for a large-scale data centers
3) (Optional, but ideal) Reaching out to business professionals and leaders about implementing environmentally-clean practices into their AI data centers

PLAN AND APPROACH

The suggested solution and experiment for the presented issue can be divided into three steps, with the final one being an optional venture towards implementation. The first of the three includes rigorous testing. Four cooling techniques used on PCs (fan, heat sink, heat pump, and vapor chamber) will be measured in terms of energy efficiency (how much energy they "waste") and cooling power (how much they are able to bring down temperatures). The techniques will be used on different PCs (make and model kept constant in order to ensure a repeatable and fair test) with sensors embedded to track these details. PC cooling efficiency can be measured with different metrics. In the case of air-cooling, the most popular and standardized unit is Cubic Feet per Minute (CFM), a measure of how much air is moved per minute. This is important as higher circulation yields more cooling. This metric, or another such measurement depending oni the cooling type, can be divided by the energy consumed to find Cooling Efficiency Ratio (CER). The latter unit is the most applicable for this study, as a large-scale implementation would prioritize efficiency and scarce use of energy. Feasibility, and environmental efficiency, the most notable qualitative data assessments in this experiment, will be assessed through research. Costliness will also be researched and noted on a scale. After concluding which are the best suited for both PCs and large-scale systems, we will move on to our more innovative stage. During stage 2, research and prototyping would be done to conceptualize a large-scale application for the chosen techniques. This is the most tentative and far-off stage. Research should be documented in a journal. If a solution is found, I would also seek to publicize the findings and attempt to implement them into a real industry in order to have a measurable impact.

RISKS

1) Insufficient funding to support acquisition of PCs and cooling

Request donations of old PCs and computers, focus on CPU instead of PCs, try to reuse parts instead of buying multiple

2) PCs and Data Centers are different

This is an issue I am well aware of. However, my solution emphasizes making a large-scale change out of very little. Thus, the main premise is to use PC cooling system so that the proposed solution can be cheap, energy efficient, safe, and simple.

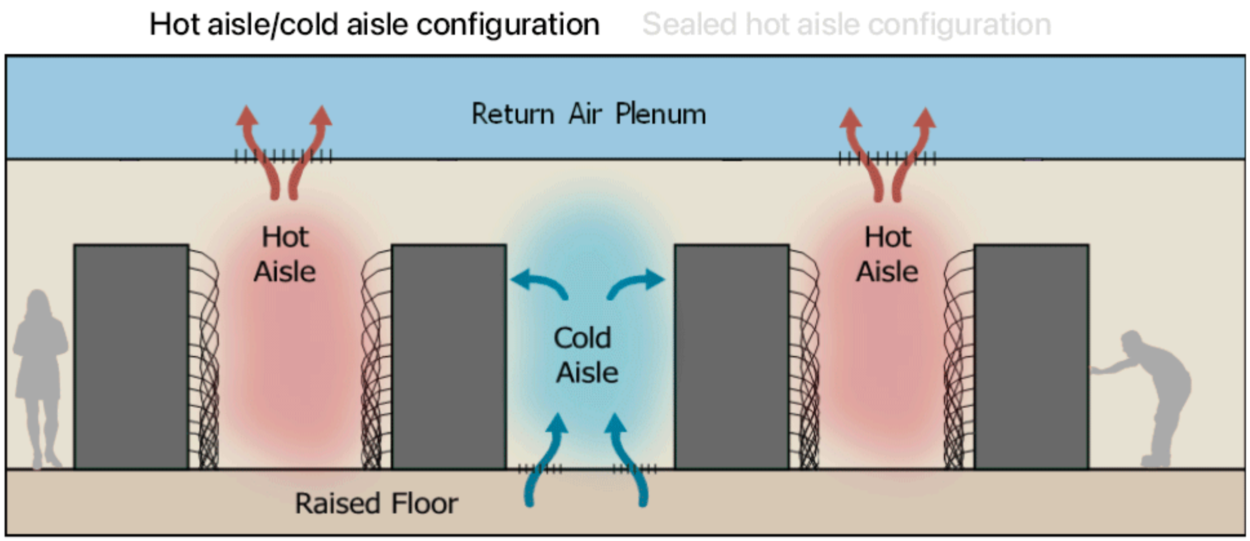

Figure 1 (*& the West)*

	Cooling Method		
	Option 1: Air Cooling	**Option 2: Water Cooling**	**Option 3: Refrigeration (R134a) Cooling**
CPU Cooling	Copper-fin heat sinks with embedded heat pipes	Option 2a: Series-parallel copper cold plates Option 2b: Parallel copper cold plates	Series-parallel copper cold plates
Heat Exchanger Location	Air-to-liquid on rear door of cabinet	Liquid-to-liquid in Coolant Distribution Unit	Option 3a: Liquid-to-liquid in Refrigerant Distribution Unit Option 3b: Liquid-to-air in Rooftop Condenser
Pump Location	Rooftop Chiller	Coolant Distribution Unit	Cabinet Level Pumps
Rooftop Cooling	Compressor-Enabled Chiller	Water-to-Air Dry Cooler	Option 3a: Water-to-Air Dry Cooler Option 3b: R134a Vapor-to-Air Condenser

Note: External Ambient air temperature = 25 °C

Figure 2 *(Electronics and Cooling)*

REFERENCES

"AI Has an Environmental Problem. Here's What the World Can Do about That." *UNEP*, 2025, www.unep.org/news-and-stories/story/ai-has-environmental-problem-heres-what-world-can-do-about.

Barringer, Felicity. "Thirsty for Power and Water, AI-Crunching Data Centers Sprout across the West." *& the West*, 9 Apr. 2025, andthewest.stanford.edu/2025/thirsty-for-power-and-water-ai-crunching-data-centers-sprout-across-the-west/.

Facts and Statistics | Wateraid, Leading Water Charity, Est. 1981. 25.9 Million People Reached with Clean Water, www.wateraid.org/us/the-water-crisis-facts-and-statistics. Accessed 17 Dec. 2025.

Group, DNP. "The Basics of Fanless Cooling Technology: Colomn: Solutions/Products/Services: DNP Dai Nippon Printing." Dai Nippon Printing Co., Ltd., DNP, www.global.dnp/biz/column/detail/10162626_4117.html. Accessed 17 Dec. 2025.

Gurevich, Yuri G., and Jesus Enrique Velazquez-Perez. "Peltier Effect in Semiconductors." *ResearchGate*, www.researchgate.net/publication/262179901_Peltier_Effect_in_Semiconductors. Accessed 18 Dec. 2025.

If you want your PC performing at its best, you're going to want to learn more about PC cooling.1 2 If you want your PC performing at its best. "PC Cooling: The Importance of Keeping Your PC Cool." *Intel*, www.intel.com/content/www/us/en/gaming/resources/pc-cooling-the-importance-of-keeping-your-pc-cool.html. Accessed 17 Dec. 2025.

"What Is a Data Center? - Cloud Data Center Explained - AWS." *What Is a Data Center?*, Amazon, aws.amazon.com/what-is/data-center/. Accessed 18 Dec. 2025.

"World Population Projected to Reach 9.8 Billion in 2050, and 11.2 Billion in 2100." *United Nations*, United Nations, www.un.org/en/desa/world-population-projected-reach-98-billion-2050-and-112-billion-2100. Accessed 17 Dec. 2025.

Yatskov, Alexander. "Comparison of HPC/Telecom Data Center Cooling Methods by Operating and Capital Expense." *Electronics Cooling*, Articles, Coolers, Data Centers, Featured, Free Air Cooling, Liquid Cooling, Refrigeration, 21 June 2018, www.electronics-cooling.com/2018/06/comparison-of-hpc-telecom-data-center-cooling-methods-by-operating-and-capital-expense/.