Negative Pressure Liquid Cooling

Abstract/Background

A variety of technology, business and market factors are driving increased demand for liquid cooling solutions. This demand was initially focused on supporting the cooling of supercomputers and high-performance applications, but it is expanding to a far broader range of hardware throughout data centers, cloud computing/co-location facilities and a wide variety of other use cases.

The single biggest barrier to widespread adoption is the perceived (and real) fear that a leak will cause severe damage to expensive hardware and cause extended downtime. Data center managers and IT executives are increasingly risk adverse and this has limited the initial receptivity for these solutions.

This short paper addresses the value proposition associated with the use of a negative pressure liquid cooling solution and how it differs from legacy data center cooling approaches

Author(s):

Chilldyne and TechVision Research

Table of Contents

Executive Summary

The standard way of keeping data centers cool is to use expensive and relatively inefficient vapor-compression refrigeration systems at least part of the time. These conventional cooling or "air conditioning" systems often use more power that the computers themselves, all of which is discharged to the environment as waste heat. Water has 3500 times more heat capacity than air by volume, so water is far better for direct heat transfer from heat generating components.

This has led to liquid cooling becoming a necessary prerequisite for all of the top supercomputers and will be required for many-to-most large-scale computing systems as server and rack density increases and chips within the servers increase in component density and power. Add to this the exponential growth in demand caused by increasing remote work, our thirst for data, and traditional organizations becoming increasingly digital, and it's easy to see why conventional cooling approaches can no longer support these systems efficiently and effectively. As the challenges increase, the shift to liquid cooling becomes the obvious solution. But not all liquid cooling is the same. There is an approach using negative pressure (virtually eliminating leaks) that is potentially disruptive in the liquid cooling space and that's the focus of this paper. Some of the key differentiators of liquid cooling systems using negative pressure (vs. conventional liquid cooling) include:

- No leaks (the single biggest concern with liquid cooling)
- Lower cost components; cold plates, plumbing and connectors generally less expensive in that they don't need to work with 3x maximum working pressure (300 psi) as with positive pressure. Positive pressure requires more expensive fittings, and expensive dripless quick disconnects.
- Easier, faster, less complicated installation and maintenance
- Increase of heat transfer by 3X; with higher vacuum levels, cooling system can achieve nucleate boiling at the cold plate. This means that cooling a 30 x 50 mm chip that puts out 10 KW is possible.

A negative pressure cooling approach has several advantages over positive pressure cooling solutions. For a data center operator, the Chilldyne liquid cooling system provides risk mitigation and efficiency through a no leak design, automatic fill and drain capabilities, and automatic coolant anti-corrosion monitoring and control. This makes for dramatically improved system uptime while using lower cost components. For example, the cluster that Chilldyne did for Sandia has been up since October 2019 with ***zero*** downtime due to cooling issues. Specific advantages include:

- Nodes are more accessible; system administrators are able to remove nodes without fear of liquid interface.
- Leak detection is relatively simple with air bubbles forming on the return line of the individual nodes.
- The power draw for operations is lower than most positive pressure Cooling Distribution Units (CDU)s.
- The flexibility in the complete design of the Chilldyne CDU and computer system is a unique approach in its ability to be either air-cooled or liquid cooled with failover valves and built in fans to automatically switch over based on temperature and flow. This means that if major plumbing repairs or modifications are needed, not only will there be no leak, the system can continue operation by switching to air cooling.

The vision is that the pervasive use of liquid cooled clusters, which are more reliable and less expensive to buy and operate than the air-cooled units, will result in reduced power consumption and a smaller carbon footprint. Liquid cooling uses less power, reduces data center size, and requires less cooling equipment. 53 million tons of carbon could be saved if just half the world's data centers used liquid cooling

> Liquid cooling uses less power, reduces data center size, and requires less cooling equipment.

Chilldyne and TechVision Research believe that the negative pressure system will prevail in the long run because of the major TCO benefit at the server level, the ability to scale, the ease of installation/maintenance (without the need for experts), and the high degree of automation for setup and operation with filling, draining and coolant additive controls automated and regulated. We'll now offer further details on negative pressure vs. positive pressure liquid cooling with examples of how the Chilldyne system works. We'll start with the most critical difference; leak prevention.

Leak Prevention

Any data center liquid cooling system will have thousands of connections, and there are many ways that they can leak. The single biggest benefit of negative pressure liquid cooling is that a leak does not damage the server; it simply lets air into the system, which is quickly pumped out. Connectors, plumbing and parts can be low cost and generic.

> The biggest benefit of negative pressure liquid cooling is that a leak does not damage the server; it simply lets air into the system, which is quickly pumped out.

The Chilldyne approach is failure tolerant, so that a system can be deployed quickly with unskilled labor and there isn't a worry about getting expensive servers wet. Stopping leaks before they start means there is no damage to hardware and no downtime; these are major advantages that negative pressure systems have over positive pressure liquid cooling systems. To understand how leaks are prevented we'll start with a description of the how the Chilldyne CDU's are designed, how they operate, and we'll then describe the specifics of how leaks are prevented.

Chilldyne Negative Pressure CDU Architecture

Chilldyne's negative pressure CDU operates under a vacuum that enables out-of-the-box leak-free operation. The chamber system of the CDU, which Chilldyne calls the "ARM" chamber, (Auxiliary, Reservoir, Main) pumps the coolant and stores it. The ARM chamber is divided into three smaller chambers: Auxiliary, Reservoir, and Main. The pumping action of the CDU is cyclical. In the first stage, the CDU applies vacuum to the Main chamber. Fluid is drawn out of the reservoir and through the servers into the main chamber. When the Main chamber is nearly full, the CDU draws vacuum on the Auxiliary chamber, and the Main chamber is allowed to drain into the Reservoir. When the Auxiliary chamber is nearly full, the cycle repeats. By alternately applying vacuum to the Main and Auxiliary chambers, the CDU creates a steady flow of water out of the Reservoir chamber, through the servers, and back into the CDU.

After the warm fluid returns to the CDU, it passes through two heat exchangers that reject the heat to a source of facility cooling, such as a cooling tower. A coolant additive management system regulates the level of anti-corrosion and biocide additives in the water to further mitigate the risk of extended downtime.

> The system can continue to operate even with minor leaks present

The basic principle is that since the CDU keeps the entire system under vacuum, water cannot leak out. If a line is damaged or a seal fails, air will leak into the system instead of water leaking out. The air is evacuated from the system via the liquid ring vacuum pump and a fluid separator, so the system can continue to operate even with minor leaks present. This recorded demo provides a visual as to what happens when a line is cut: https://www.youtube.com/watch?v=552tzND2Xx0

The vacuum also allows servers to be disconnected from a live system without worrying about leaks or shutting off flow to the rack. This allows a level of uptime, ease of maintenance and flexibility rarely available within data centers. When a server is disconnected, the water inside is automatically evacuated, leaving the server dry for maintenance. The following figure and subsequent descriptions show the basic components and flows within the negative pressure liquid cooling system:

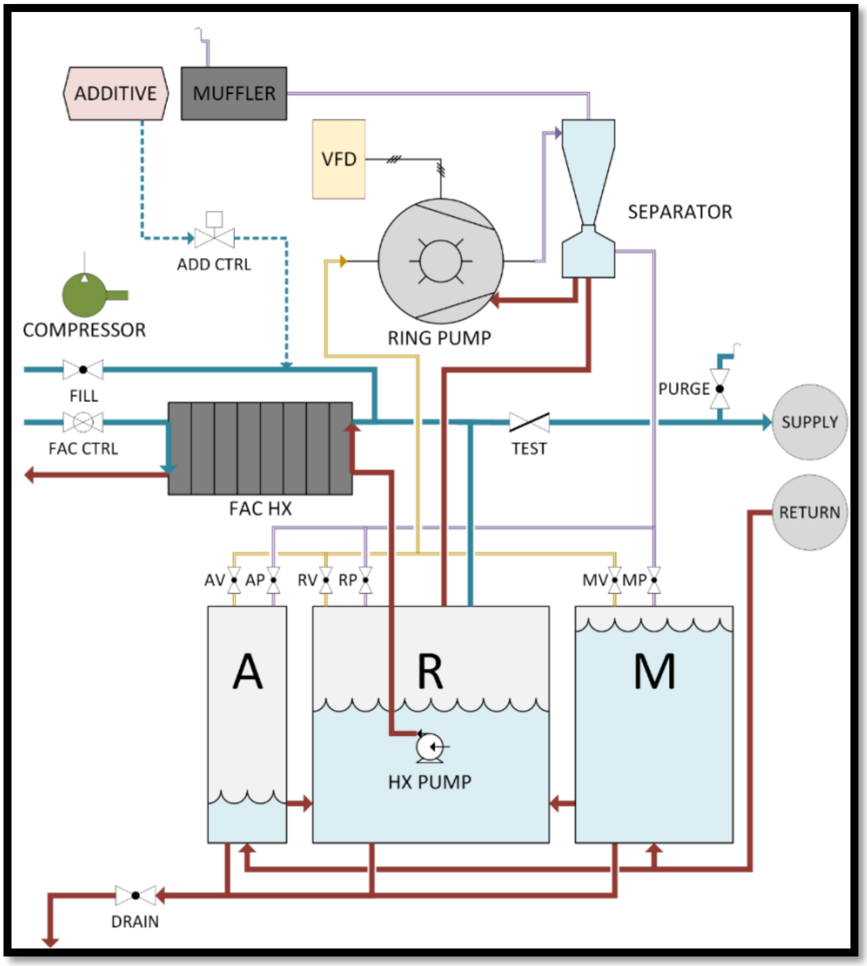

Figure 1: CDU Hardware Layout of Chilldyne's CDU

This table describes the components referenced in figure 1 for ease of understanding.

Tag	Name	Description
MUFFLER	Muffler	Prevents droplets from escaping and reduces audible volume of system.
ADDITIVE TANK	Additive Tank	Stores coolant additive solution for periodic distribution.
VFD	Variable Frequency Drive	Provides AC power and speed control for LRP.
RING PUMP	Liquid Ring Pump (LRP)	Pulls vacuum on chambers to induce flow.
SEPARATOR	Separator	Separates excess fluid pulled into LRP.
COMPRESSOR	Air Compressor	Provides pneumatic power to valves.
FAC HX	Facility Heat Exchanger	Moves heat from the process loop to the facility loop.
SUPPLY	Supply Manifold	Multiple connection point for supply coolant.
RETURN	Return Manifold	Multiple connection point for returning coolant.
R	Reservoir Chamber	Holds low vacuum to allow fluid flow out to the process loop.
M	Main Chamber	Alternates holding high vacuum to pull fluid through process loop.
A	Aux Chamber	Alternates holding high vacuum to pull fluid through process loop.
HX PUMP	Heat Exchanger Pump	Forces warm coolant up to the facility heat exchanger.

Table 1: Descriptions of key components

How does the Chilldyne system work with a leak?

This provides a bit more technical detail on the Chilldyne solution. First the Chilldyne system works with a server that is completely open to air even as air flows into the server fluid connections. The connections to the server include a check valve with a controlled leak in the reverse direction on the supply side, and a sonic nozzle venturi on the return side. These valves limit the flow of air into the rack manifold in the event that there is a major air leak into the server.

Under normal operation, the flow resistance of the check valve on the supply side is about 0.1-inch Hg and the flow resistance of the venturi is about 1-inch Hg. Under a leak condition, the controlled leak in the check valve limits the air flow into the supply manifold to 2 lpm. This results in some bubbles in the coolant for the servers downstream of the leak, but the bulk density of the coolant is lower so the volume flow rate is higher, and the net result is that the downstream server temperature may go up by 1 to 3^{o}C, but the system still works to liquid cool all of the servers except the one with the leak, and it can use air cooling. On the return side, the venturi limits the air flow to about 10 lpm as the air flow cannot exceed the speed of sound in the narrow part of the venturi. (Sonic Nozzle).

Basically, this means that the negative pressure is sufficient to both eliminate the leak and maintain operations during the leak condition. No other system offers this level of failure tolerance.

Ease of Installation/Maintenance

As Dell and others consider scaling the negative pressure approach, it is important to consider the ease of installation and maintenance. There are inherent advantages to the negative pressure liquid cooling approach in general, but also some specific advantages of the Chilldyne system. We'll start with a general look at positive vs negative pressure with a real use case.

Positive pressure government use case: It took over 2 weeks using a professional plumber to install plumbing from the CDU to support 25 racks with positive pressure cooling. This includes cutting pipe, soldering, threading and expensive pressure-rated materials. Scalability is constrained by the local availability of skilled labor.

Negative pressure government use case with Chilldyne: It took 1 week to install; there was no mess and no plumbing professionals were required. The Chilldyne system was installed by unskilled labor using flexible plastic tubing and a knife. Note: the CDUs do need to be within 100 feet of the racks and on the same level for a negative pressure solution to work correctly so this may not be applicable in all situations.

Beyond installation, the on-going operation in a data center will be more cost effective and efficient with a correctly architected negative pressure solution. Ongoing benefits include reduced costs and higher uptime when replacing and removing servers/components. The data center operator can also simply shut down a rack (not the entire system) if a leak occurs until the source is detected…and the leak doesn't create an emergency condition.

Other benefits include no downtime associated with standard maintenance/filter replacement. Lower cost plumbing materials (plastic vs copper) can be used. With liquid cooling, most solutions have a requirement to the replace coolant every few years. The Chilldyne system drains/adds coolant and additive automatically so there is not downtime during regular maintenance.

The bottom line is that negative pressure is lower cost, easier to install, and easier to maintain. It is engineered to deliver substantially higher availability and mitigates the risk of a leak causing major hardware damage and significant downtime.

Summary

Liquid cooling is the future of computing; the combination of increasingly powerful and higher density chips, and the associated heat they generate make liquid cooling a future-state necessity. Within the liquid cooling space, negative pressure is emerging as the strongest model as it addresses the single biggest issue with liquid cooling—elimination of leaks. Reducing that threat, coupled with easier and less expensive installation, lower cost components, better reliability and lower total cost of ownership make this a leading technology to consider in building next generation data centers, cloud computing facilities and supercomputers.