760.476.3419
info@chilldyne.com
www.chilldyne.com

5900 Sea Lion Place
Suite 150
Carlsbad, CA 92010

Data Center Cooling FAQs

Contents

Introduction

This is designed to help people understand liquid cooling in general and Chilldyne's solution in particular.

What is most important for data center liquid cooling?

1. Uptime
2. Uptime
3. Uptime
4. No leaks
5. Server-side cost
6. Infrastructure cost
7. Heat capture ratio (percent of heat into the liquid cooling system)
8. Cold plate thermal performance
9. Amount of maintenance required
10. Installation cost
11. Appearance

Which type of data center liquid cooling is best?

A. Direct to the chip liquid cooling (positive or negative pressure):

"*In direct to chip cooling, a liquid coolant is brought via tubes directly to the chip, where it absorbs heat and removes it from the data hall. Because this system cools processors directly, it's one of the most effective and efficient forms of data center heat removal."* [1]

With direct to the chip cooling, not all the server heat is captured. The liquid cooling system can capture about 80% of the heat in servers with 205-watt CPUs. As the CPU power goes up in the future, the percentage of heat captured will go up to 90% or more. If the facility cooling is based on cooling towers, and the data center is allowed to run warm air temperatures, 95-100% of the heat can be captured.

Each server requires a fluid connector and each rack requires a fluid manifold. The data center layout does not change as compared to an air-cooled data center. The same racks and hardware can be used.

B. Direct to the chip liquid cooling with negative pressure CDUs

With negative pressure liquid cooling a leak does not damage the server, it just lets air into the system. Connectors, plumbing and parts can be low cost and generic. Any data center liquid cooling system will have thousands of connections, and there are many ways that they can leak. The negative pressure approach is failure tolerant, so that a system can be deployed quickly with unskilled labor. There isn't a worry about getting expensive servers wet. Negative pressure systems are limited to less than 1 atmosphere delta pressure.

---

[1] The Benefits of Direct to Chip Liquid Cooling in Your Data Center: https://www.rittal.us/contents/what-is-direct-to-chip-liquid-cooling/

C. Direct to the chip Liquid Cooling with positive pressure CDUs

Positive pressure liquid cooling is very common, for instance, gamer cooling systems are all positive pressure. Positive pressure CDUs are simple. Connections to the racks require sweated copper plumbing, installed and leak tested by professional plumbers. Positive pressure may be needed in legacy systems with high pressure drop connectors.

A water leak can ruin an entire rack of servers. Leak detection must be used and air must be purged from the system.

D. Immersion

Immersion cooling removes nearly all the heat from the server. The data center layout is completely different. Optical connections must be completed outside of the liquid. Spinning hard drives must be sealed. Fans must be removed and fan control software updated. Thermal grease must be replaced with Indium or phase change material. Special power and interface cords must be used. The value of the servers in the secondhand market is limited.

E. Oil Immersion

Oil is widely used to cool transformers and high voltage electrical devices. It is used primarily for high voltage utility systems or avionics cooling, where arc resistance and operation down to -40^{o}C is needed. It can cause embrittlement of plastics when used to cool servers. Keep in mind that with immersion cooled systems, fixing them is like fixing an automatic transmission. It will be a messy, expensive process. Maximum power per chip is limited by the heat capacity of the oil.

Vienna HPC gave a paper on their experiences with oil immersion at SC19, they had some plastic power plugs shatter due to the oil, see below.

Figure 1: Oil Immersion image (Power cable after wicking oil from oil bath.) [2]

[2] https://eehpcwg.llnl.gov/assets/sc19_05_1150_1230_impressions_from_5_years_of_oil-based_liquid_immersion_cooling_hofinger.pdf

F. Two phase immersion cooling

The Wright brothers used boiling water to cool the engine in their 1903 flyer. The Cray-2 used boiling fluorocarbon. Today, many Bitcoin miners are cooled by fluorocarbon fluids. Fluids like Novec and Fluorinert leach out plasticizer and cause plastics to crack They also leak easily, evaporate quickly, and are very expensive. The servers must be contained in a sealed box. All power and communications cables must be sealed. For the Bitcoin operation, any chip that fails is abandoned in place, so they don't need to open up the sealed box. For Intel or AMD servers, the server is too expensive to abandon, so it must be removed for repairs. This will result in the loss of expensive fluid.

An arc in these liquids can create hydrofluoric acid and requires a special cleaning process. LLNL wrote a paper in which they did not recommend two phase cooling and described how it caused tin whiskers shorting out power supplies:

> *"The failures resulted in permanent component damage (restart attempts were unsuccessful) for 58 of the 72 nodes."* [3]

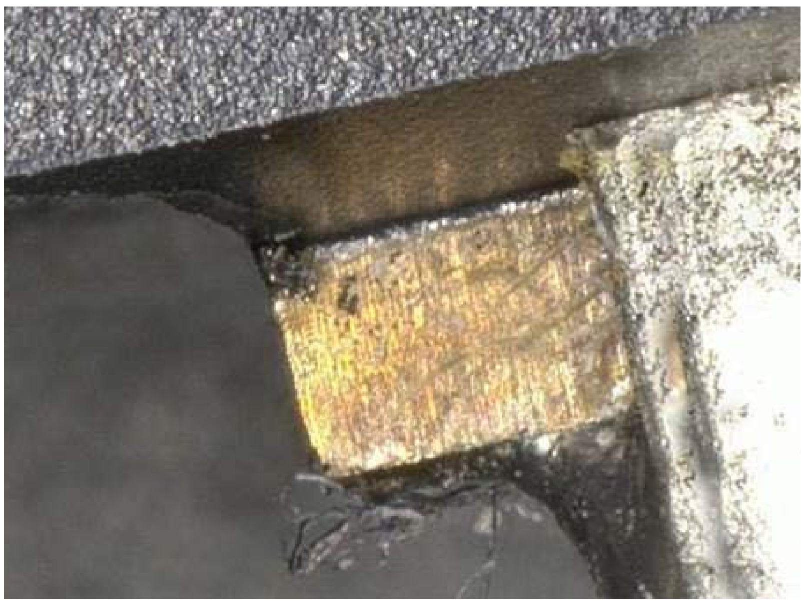

Figure 2: Two phase immersion cooling damage (Tin whisker on power supply)

How does Chilldyne's negative pressure CDU work?

Chilldyne's negative pressure CDU operates under a vacuum which allows for leak-free operation. The chamber system of the CDU pumps the coolant and stores it. The chamber is divided into three smaller chambers: Auxiliary, Reservoir, and Main. Together, these are referred to as the ARM chamber. See Figure 3 for details. 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

[3] https://datacenters.lbl.gov/sites/default/files/ImmersionCooling2016.pdf

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 is pumped 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.

Because the CDU keeps the entire system under vacuum, water cannot leak out. If a line is damaged or a seal fails, air leaks into the system instead. 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. The vacuum also allows servers to be disconnected from a live system without shutting off flow to the rack or the CDU. When a server is disconnected, the water inside is automatically evacuated, leaving the server dry for maintenance. We are continuing to improve the CDU design for greater reliability and lower cost.

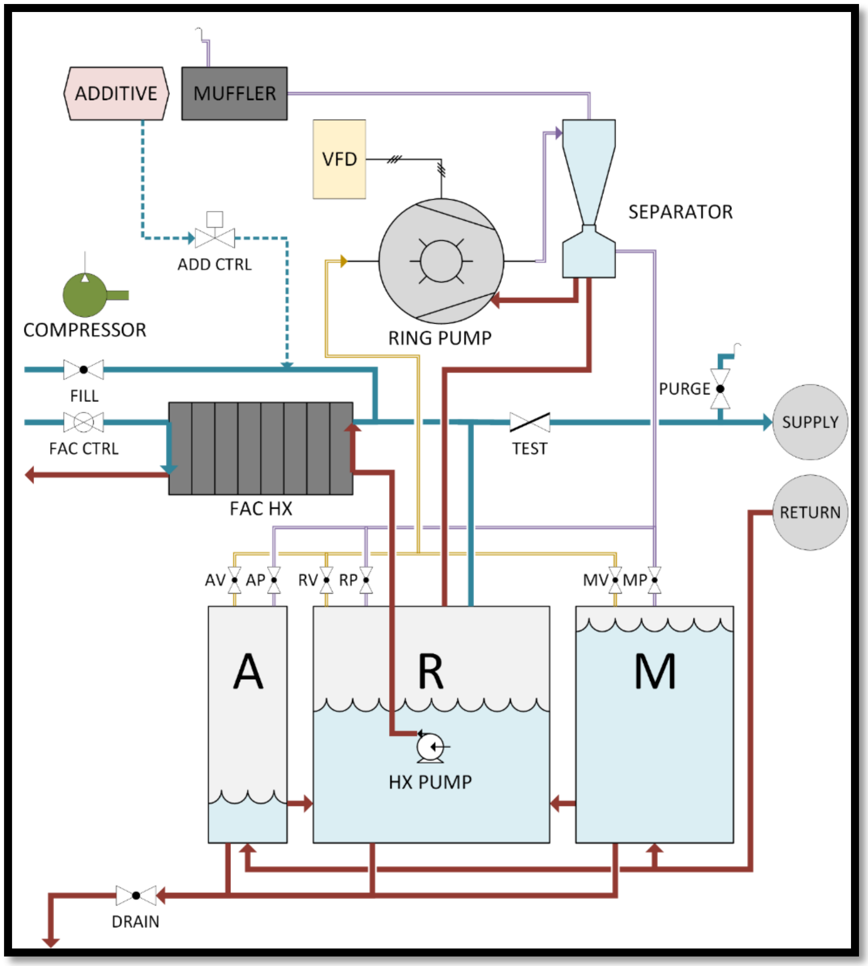

Tag	Name	Description
MUFFLER	Muffler	Prevents droplets from escaping and reduces audible volume of system.
ADDITIVE TANK	Additive Tank	Stores coolant additive solution for periodic distribution.
VFD	Variable Frequency Drive	Provides AC power and speed control for LRP.
RING PUMP	Liquid Ring Pump (LRP)	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.

Figure 3: Hardware Layout that depicts the major components of Chilldyne's CDU

How does the Chilldyne system work with an air leak?

The Chilldyne system works with a server completely open to air. 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 liters per minute(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°C, but the system still works to liquid cool all of the servers except the one with the leak, and it can use the back-up 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)

How does the Chilldyne system drain the servers and racks when they are disconnected during operation? How does the CDU keep from overfilling when racks are evacuated of coolant?

The CDU contains about 50 liters of coolant. In the event that more than about 15 liters is removed from the racks and servers and gets sucked into the CDU, the CDU has a drain connection which typically connects to the sewer, and the excess coolant is pumped down the drain while it is running, The CDU also has a water fill connection, which is typically connected to a reverse osmosis water supply, so that when the racks are added to the system, the CDU refills and continues to cool the existing racks and added racks with no downtime. It also automatically adds more coolant additive if required.

When servers are disconnected from the cooling system, the supply line is disconnected first, and then the return line sucks out all the coolant from the server. The air flow into the rack manifold and the CDU is limited by the Venturi. The CDU vacuum pump has a capacity of 1,200 lpm so the air leak does not reduce the system performance significantly. (It doesn't lose suction, just like the famous vacuum cleaner)

How does the Chilldyne system work with only 1 atmosphere of available delta pressure?

Less than 1 atmosphere of pressure works because the system uses slightly larger ID tubing, (pressure drop goes as the inverse fifth power of pipe diameter) and no quick disconnects are needed except at the servers. The positive pressure systems use plumbing similar to tap water plumbing, but Chilldyne takes advantage of the negative pressure by using wire reinforced flexible PVC tubing which is much easier to install and has lower pressure drop than metal pipe plumbing with sharp elbows and restrictive quick disconnects. Below is a Chilldyne installation at Sandia National Laboratory. The pressure drop accounting is shown in *Table 1: Pressure Drop Accounting*. The low pressure drop has the added advantage of reducing the power required for the negative pressure liquid cooling system as compared to positive pressure systems. The system uses about 3.8-inch Hg (4 ft of water) of pressure drop in the servers so that the racks clear air bubbles automatically, and the flow does not short circuit through the lower servers in a rack.

Figure 4: Chilldyne Installation at Sandia National Laboratory

Table 1: Pressure Drop Accounting

Pressure Drop Accounting	**Flow lpm**	**ΔP in Hg**
Server	0.4	3.8
Venturi Flow limiter	0.4	1
Leaky Check Valve	0.4	0.1
Rack Manifold	24.8	0.5
1-inch Tee	24.8	0.1
10 ft 1-inch Tubing	24.8	0.3
1.25 Tee	49.6	0.2
10 ft 1.25-inch tubing	49.6	0.4
Fail over valve	198.4	0.1
40 Ft of 2 Inch tubing	198.4	2
Rack (subtotal)	1.6	
CDU-Rack (subtotal)	3.1	
Total	**8.5**	
Available DP @ 200 lpm	**11 avg, 17 max**	

In the Chilldyne system a separate pump is used to pump water from the reservoir through the heat exchangers and back to the reservoir. (See the image in *Figure 5: Liquid Cooling Schematic*) Thus, the Heat exchanger pressure drop does not add to the server loop pressure drop and the approach does not depend on the server flow rate. (The approach is 1-2 ^{o}C at 200 kW)

Does the Chilldyne system work at higher altitudes with lower ambient pressures?

The maximum available pressure is the difference between ambient pressure and the vapor pressure of water at the highest temperature in the server loop. At higher altitudes, the outside air temperature is lower so the available pressure is similar to sea level. If we assume that we are using dry air coolers, the cooling water temperature is 10 ^{o}C warmer than the air temperature. If we further assume that the temperature rise on the server loop is 10^{o}C, then the highest water temperature is 20^{o}C warmer than the ambient temperature.

For example,

At Los Alamos, at 7000 feet, the highest ambient temperature is about 35^{o}C, so the cooling water is 45^{o}C and the return water is 55^{o}C. At 55^{o}C, the vapor pressure is 4.4 In Hg, and the absolute pressure is 23.1-inch Hg, so the available pressure is 18.7-inch Hg. The system works with cooling water at 45^{o}C.

In Tucson, near sea level, the highest ambient temperature is about 47^{o}C. In this case we would have return water at 67^{o}C. At 67^{o}C the vapor pressure of water is 7.7-inch Hg, the total available ΔP is 22.2-inch Hg. In this case the return water is hot enough to scald a technician and the negative pressure provides safety since a leak of hot water is not a concern.

Does the Chilldyne system work with overhead and raised floor connections?

As with the power and network cables, installation is easier with a raised floor, but overhead works as well. In a recent installation, there are 90 nodes per OCP rack, dual 165-watt Xeons with an overhead connection. The total power is 50 kW per rack.

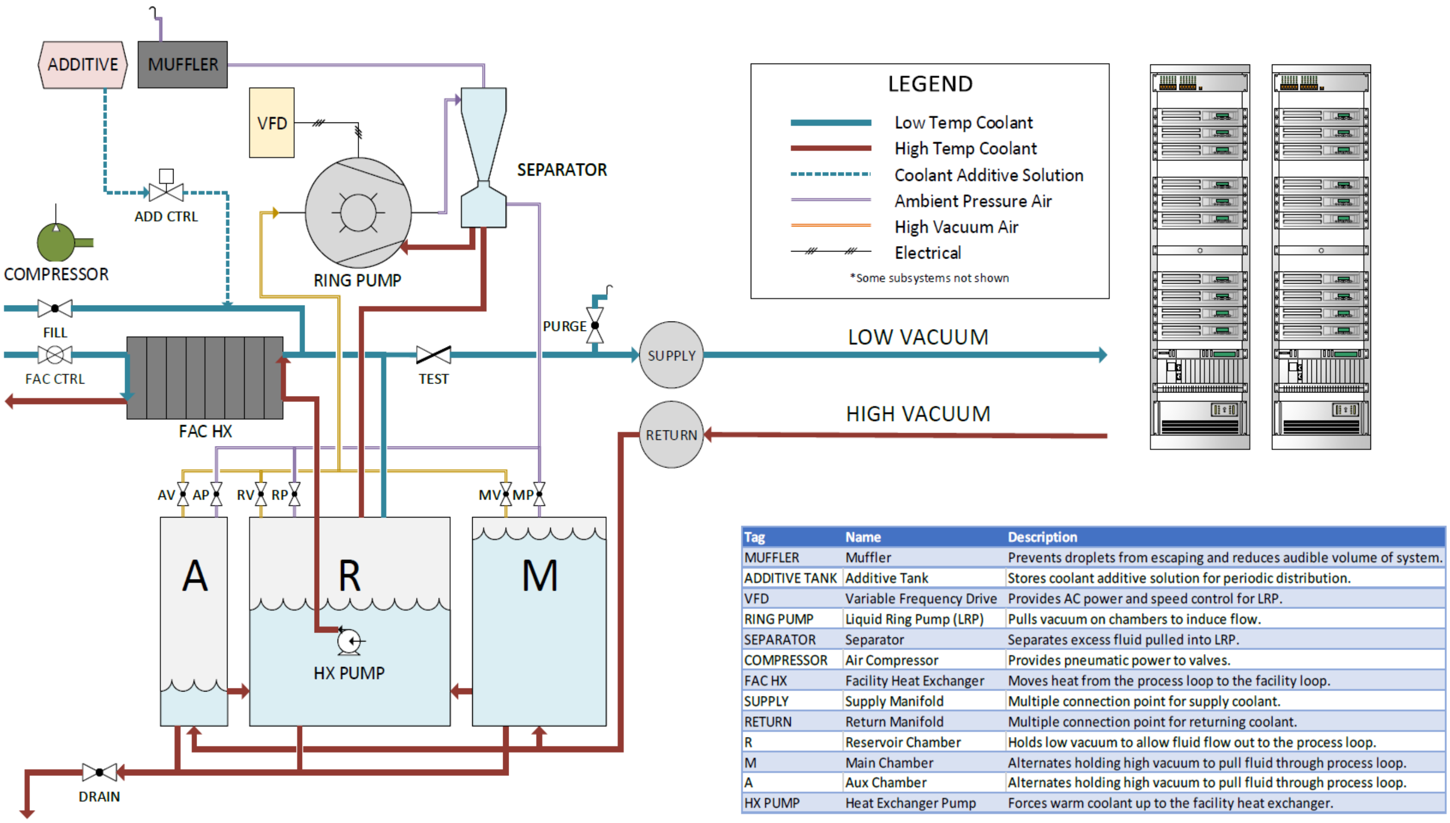

Tag	Name	Description
MUFFLER	Muffler	Prevents droplets from escaping and reduces audible volume of system.
ADDITIVE TANK	Additive Tank	Stores coolant additive solution for periodic distribution.
VFD	Variable Frequency Drive	Provides AC power and speed control for LRP.
RING PUMP	Liquid Ring Pump (LRP)	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.

Figure 5: Liquid Cooling Schematic (Heat Exchanger pump is R portion of chamber)

Since liquid cooling has been around for years, it must be a mature technology. Why shouldn't I buy the cheapest CDU with the cheapest cold plates, manifolds and quick connects?

Data center liquid cooling is not a commodity. Each system is still being improved upon and the sweet spot for price, quality and reliability has not been determined. The cheapest solution may cost a lot of money for downtime. Those with a large budget may want to use the Cray or IBM style of using aerospace fittings and connections, which are designed for aircraft with 30-year lifespans but are too expensive for a computer that goes obsolete in 4 years. With the Chilldyne system, high reliability does not require aerospace quality fittings.

Why is the Chilldyne system better?

The negative pressure cooling approach has a few more added advantages than a positive pressure cooling solution.

The whole idea of the liquid cooling system is that there are no worries for the data center operator. No liquid leaks, automatic fill and drain, automatic coolant anti corrosion monitoring and control are why Chilldyne is better. This makes for good system uptime. The Attaway cluster for Sandia has been in operation since October 2019 with no downtime due to cooling issues.

- 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 CDU's.
- Experts understand that the flexibility in the design of the Chilldyne CDU and cooling system. The ability to be either air cooled or liquid cooled with fail over valving and built in fans to automatically switch over to maintain uptime, regardless of breakdowns is unique.
- Data on the system operation can be accessed using Modbus, SNMP, ftp, or the web page.

Why is your team best at data center liquid cooling engineering?

The Chilldyne team has over 100 combined years of thermal and fluids design experience, plus 9 years of data center liquid cooling experience, so we are well equipped for the technology challenge. We recently developed a liquid cooling system for a defense contractor, and we were by far the lowest price and fastest delivery, and they never had to worry about leaks while testing their new product.

We also helped develop the liquid cooling system for the Hunter UAV for Northrop Grumman: http://www.flometrics.com/project/aircraft-cooling-system-design/

We have developed cooling systems for rocket engines, lasers, and medical devices. Applying aerospace quality systems to data center liquid cooling results in systems like what is sold by IBM, Cray, SGI, i.e. failure proof, expensive systems. We have experience with disposable medical fluid systems which must be low cost *and* high reliability. We have applied this know-how to the data center liquid cooling problem.

With Chilldyne, we have rocket scientists designing cooling systems. If you think cooling a high-power processor is hard, try cooling a rocket engine.

What are the typical failure modes of a liquid cooling system?

Leak in server, i.e. hose not connected or broken Quick Disconnect

For each server, there are flow limiting valves, this means that any single server failure will not result in reduced cooling for the other servers. The server can fail over to air with fins on the cold plates. In most cases, the processor will not throttle.

CDU failure

Any CDU failure will result in the activation of a fail over valve to switch to a backup CDU. (This is our current recommended practice, N+1 CDUs). At the CDU level, the software constantly checks to see if the water level, temperature and pressure sensors are reading correctly and if not, it issues a warning. There is also a check for air flow out of the CDU to see if there is a leak of air into the system. If there is air leaking, a warning is issued. In some cases, the CDU can continue to work with a defective sensor, like a modern car with a *limp home mode* in case of a sensor failure.

The CDU is designed to never apply positive pressure to the servers, including in the event of a liquid cooling system power outage with the servers still running at full power. The valves in the CDU are designed to have power off valve positions such that the system is vented under power off conditions. Referring to the schematic in *Figure 5: Liquid Cooling Schematic* the test valve is normally closed and the purge valve is normally open. This means that if the power shuts off, the vacuum in the pump chambers will suck some of the coolant into the pump chambers, while the purge valve will let air into the system and the test valve will prevent more coolant from entering the system.

Component Failure

Our current practice is to replace any problematic component with ones of higher reliability. Before any components wear out, they will be replaced on a maintenance schedule to prevent downtime. The CDU has an hour meter on it which is used to recommend appropriate replacement intervals for components subject to wear.

Should I implement heat recovery with my liquid cooling system?

When looking at heat recovery, consider balancing the extra power required at higher CPU temperature with the cost of efficient heat pumps (a one to one comparison of watts of heat to watts of computer power is not appropriate). A hotter running processor may require 5% more power. A typical heat pump has a coefficient of performance (COP) of 4 so that is provides 100 watts of heat to the building while using 25 watts of electricity.

This means that a nominal 100-watt processor running hot enough to generate useful heat will use 5 extra watts of power. 100 watts of building heat normally requires 25 watts of heat pump power. This means about 20 watts of power savings per 100 watts of heat. This may not justify the expense of attaching the computer to the heating system.

Also, once a building uses a computer for heating there needs to be a backup source of heat since the computer may not operate all the time

Depending on the cost of energy in your region and required heating, you could save on the cost of heating with the waste heat from liquid cooling. This will work best in colder areas when the computers are on all the time and electricity is expensive.

How efficient is the liquid cooling system?

If the data center is air cooled, the air cooling will use 10-30% of the server power to run fans (server fans and data center fans), air conditioners, etc. Liquid cooling reduces that number to 2%, as long as chillers are not used to provide cooling water. The liquid cooling uses 5 to 15x less power than air cooling.

The difference between and efficient and inefficient liquid cooling system is very small. Heat capture ratio is more important:

Let's suppose that the liquid cooling system has a pumping efficiency of 25% and that it runs at 1 kW per 1 lpm which implies a 14^{o}C rise for the cooling water. The ΔP is assumed to be 8 psi for a positive pressure system and 20 psi for a negative pressure system.

The flow work for 1 lpm at 20 psi is 1 watt for a negative pressure system and 2.3 watts for a positive pressure system. This is only 0.2% of the server power. Since the pump is 25% efficient, the pump power is 2.3/.25=9.2watt

So, the liquid cooling system power is 9.2 watts (or less) and the liquid cooling system uses 0.9% of the server power.

Suppose that the air conditioning system has a COP of 3.3 and that the heat capture is 80%, then the HVAC power required to remove the 20% of the heat is 20%/3.3 or 6% of the server power. If 85% of the heat is captured, then the HVAC power is 4.5 % of the server power.

So, the power for pumping the liquid cooling system is much less than the power for the air conditioning to remove the heat for the small amount of heat not captured. If the remainder of the heat is removed by fans instead of air conditioners, this is still true.

The most important factor in efficiency is the heat capture ratio, and it depends more on the server design than the liquid cooling system design.

How efficient is the liquid cooling system at cooler and hotter temperature settings?

There is a processor temperature effect, which means it is more efficient at colder temperatures. The processor might use 5% less power liquid cooled; this means that more liquid cooling system flow (which takes a small amount of power) can result in less overall data center power.

So, the benefit for running the processors hotter and saving pumping power must be balanced against the extra power for the processors when they run hot, and the extra cooling required for keeping the data center air temperature cool enough for people to work.

How much heat is captured into the air?

The most important factors for heat capture are 1) the CPU power and 2) the difference between the cooling water temperature and the data center air temperature. Using cooling tower water and a warm data center will generally be the most efficient. As processor power continues to increase, the percentage of heat captured by a direct to chip liquid cooling system will also increase. At 200 watts CPU power, approximately 80% of the heat goes into the CPU. As the power increases the percentage will also increase. If necessary, recirculating air inside the server can be used to capture even more of the server heat into the liquid cooling system. In lab tests 90% of heat has been captured using air recirculation.

How does the Chilldyne system prevent condensation?

The Chilldyne system uses an onboard humidity sensor to measure the dewpoint and controls the coolant output temperature with redundant systems. A liquid cooling system must never provide the servers with water that is below the dewpoint, even if the servers are off.

What kind of server air leak detection system is needed?

The Chilldyne negative pressure system never leaks fluid onto a server so there is no need for the coolant leak detection system in the server racks. The Chilldyne system detects air leaks into the system by measuring the air flow in the return line to the CDU. Once a leak is detected it can be located by following the bubbles in the coolant return line to the source of the leak.

How are system wide air leaks detected?

For system leak detection an air flow sensor measures the air flow out of the CDU. If the air flow out of the system is above a threshold, a warning is issued.

There is no need for rack-based moisture detection systems. The only way for a leak to occur is in the event of a leak on the server side and a CDU failure. In this rare double failure scenario, the amount of the leak will be limited to the fluid volume in the rack, and a moisture detection system would cause an alarm that had already been raised due to the CDU failure.

What action is required in case of an air leak?

No action is required in the event of an air leak other than a routine maintenance notification. Any server level air leak can be repaired during the next business day or whenever it is convenient. In the event of this air leak, the Chilldyne system will cool the other servers in the rack with a 1-2°C increase in CPU temperature compared to the CPU temperature before the leak. The air leak can be up to 10 lpm of air, which would be an air leak due to a server or fluid connector fully open to air. The Chilldyne system includes flow limiting valves to ensure that any leak in a server or due to a broken quick disconnect will not result in downtime of any server besides the one with the air leak. The server with the air leak will have backup air cooling via fins on the cold plate so that it can continue to operate at a reduced clock speed and power dissipation. In the event of a rack level air leak, the system may stop cooling.

How is the liquid cooling system serviced?

The Chilldyne liquid cooling system is designed for continuous operation with no shut down required for changing filters, coolant additive or coolant. Furthermore, each server can be easily removed and replaced with only a few seconds of time required to disconnect the fluid couplings. See *How is the server serviced?*

How is the server serviced?

Each server can be serviced by removing the blue supply connector first and then the red return connector after a few seconds. This leaves the server basically dry inside so component level disassembly and repair can be accomplished without the worry of spilling coolant onto the motherboard. If necessary, it then can be shipped back to the manufacturer for service. We also have an avialable no drip hot swap connector that drains the server on disconnection.

What is the breakdown of the Total Cost of Ownership (TCO) of a liquid cooling system?

1. *Hardware cost. These are less expensive with a negative pressure system, as the requirements are less.*
   - Cold plates, internal tubing, quick connects (replaced every server refresh)
   - Manifolds and tubing from server to racks (may be replaced or reused with each server refresh)
   - Tubing from CDUs to manifolds, (may be left in place)
   - CDUs and fail over valves, left in place, refurbished every server refresh cycle.

2. *Installation cost. These are less expensive with negative pressure, as tasks don't require experts.*
   Fill, purge, vacuum test and drain are all automated, so there is no need for experts for any Chilldyne liquid cooling installation or maintenance.

3. *Facilities costs associated with installation: only the CDU needs plumber installed connections, the racks can be hooked up by non-experts*

   This includes mechanical upgrades needed to support new air-cooling loads or any liquid-cooling infrastructure needed to install a liquid-cooled system.

4. *Annual IT energy consumption (less energy needed as a side effect of lower delta pressure required)*

   The Chilldyne liquid cooling system uses less energy than other liquid cooling systems Sandia measured 1.6% of the server power on Attaway. Typical efficient air-cooled data centers use 30% of the server power to cool the servers.

5. *Annual cooling energy consumption is lower.*

   This includes chillers, pumps, air handlers, etc. Liquid cooling systems will reduce air conditioning loads and will reduce fan induced vibration and server temperatures. The server fans will never wear out. Less expensive sleeve bearing fans can be substituted for ball bearing fans due to lower temperature and lower fan speeds.

6. *Annual maintenance of the servers and supporting mechanical facilities is reduced. Much of the maintenance has been automated at the CDU.*

   This includes annual costs to maintain and support any cooling distribution system required for liquid-cooled systems. The coolant can be replaced while the system is running, so no need for downtime for this activity.

What type of coolant needs to be used?

The Chilldyne system uses water with anti-corrosion and anti-bacterial additives as coolant. This is better than glycol mixtures because pure water has 4% higher heat capacity and 2x lower viscosity than glycol. This results in 15% more flow required with 25% glycol for the same cooling capacity using water. The only reason to use glycol is for systems that must be shipped full of coolant in potential freezing environments. The Chilldyne system ships dry so there is no risk of freezing. As processor power goes up, higher performance coolant is more important.

The system includes on board coolant additive to reduce the need for periodic water testing to be done by facility employees. The standard coolant replacement schedule is once every 18 months. The Chilldyne coolant additive is composed primarily of household chemicals such as lye, fertilizer and borax and can be poured down the drain in most jurisdictions. This reduces the total cost of ownership of the Chilldyne system significantly compared to other systems.

What is the recommendation to reject the heat?

Cooling towers are best, unless water is scarce, in which case dry coolers can be used. Misters can be combined with the dry coolers for extra cooling on hot days.

What type of connectors are used?

The Chilldyne negative pressure system uses low cost generic liquid connectors because a leak is not a high priority issue. The negative pressure system does not require aerospace quality connections.

What is the recommended flow rate for a given power at what temperature?

We recommend one liter per minute of flow to cool 1kW of server power. This leads to a temperature rise of approximately 12°C. The cooling flow can be customized based on CPU requirements and available cooling. For example, if the CPU maximum temperature is low, and the local climate is hot, then more flow per kW would be needed.

We can use water from 0 to 50°C as the input although we do not recommend using hot water as this represents a burn hazard for people servicing the servers. Colder water results in more heat capture at the server level, lower server power and greater reliability for the semiconductors. We recommend the lowest temperature coolant that can be achieved without vapor compression chillers.

What pressure is used and what kind of water?

Water pressure is below atmosphere, approximately 4 inches of mercury vacuum on the supply side, and 18 inches of mercury vacuum on the return side. Reverse osmosis (RO) water in the server along with the Chilldyne coolant additive to reduce water chemistry issues is recommended, but tap water can be used as well. The system includes plastic, stainless steel, copper and brass in contact with the liquid.

How does an automatic fail over valve prevent downtime?

The automatic fail over valve switches a set of racks from a main CDU to a backup CDU when the flow is too low or the return temperature is too high. In this way CDU downtime does not cause cluster downtime. This is similar to air cooling with N+1 HVAC systems. When the failover valve activates, is splits the load with the remaining CDUs.

How should I prepare the site for Chilldyne liquid cooling?

The CDUs need facility cooling water, RO or tap water and drain pipes. All these are also needed by a typical CRAH, so they are generally available in the data center. The CDU can be fitted with a pump to pump the fluid drained out overhead if necessary.

The tubing to connect the servers and racks to the CDU is all under negative pressure so wire reinforced flexible PVC tubing is used. This tubing can be installed by data center technicians very quickly with no plumbers required. If the technician makes a mistake installing the tubing, the mistake is easily detected because the transparent tubing allows the operator to see bubbles inside the tubing which can be traced back to the leak after the CDU is turned on.

What type of sensors are used to measures the performance of the liquid cooling system?

The Chilldyne rack manifolds include accurate thermistor-based temperature sensors (0.2°C accuracy) to measure the input and output temperatures for each rack.

How is the flow measured?

The flow is measured at the rack level by a vortex flowmeter with accuracy 1.5% full scale (1.-26.4 GPM) In addition, the CDU measures flow rate and input and output temperature then calculates the heat dissipated by the load on the CDU. The facility input and output temperature are measured and the data is used to determine the facility water flow rate. All the data is available via SNMP or Modbus.

How do you prevent corrosion of metal parts or chemical reactions of the plastic with the coolant additive?

The cold plates have all metal fluid passages to eliminate the possibility of leaks due to thermal expansion mismatch between plastic fluid passages and copper heat transfer surfaces. Although the Chilldyne system uses negative pressure, the server-side assembly has been tested to 200 psi. The coolant additive in the system has been tested long term with the Buna Rubber, PVC, CPVC Silicone, urethane and Loctite sealant. The Chilldyne cold plates use turbulators in drilled passages for resistance to corrosion and contamination. The cold plate passes contaminants up to 0.25 mm and is corrosion tolerant up to 150 microns. The CDU has a 5-micron filter which filters 5% of the coolant flow. The filter can be changed without shutting down the CDU. Strainers are used in the tubing to the racks to prevent any problems due to debris getting in the system.

How do you drain the system?

The CDU drains automatically by pressing the drain button on the control panel. The coolant with the standard concentration of coolant additive can be flushed down the drain, as it contains a diluted solution of lye (drain cleaner), borax (laundry detergent), sodium nitrite (preservative) and sodium molybdate (fertilizer). In the event that the local wastewater district is not comfortable with these chemicals being flushed down the drain occasionally, it can be drained into a container and disposed of in accordance with local regulations.

Is the Chilldyne CDU certified:

It meets UL and FCC requirements for radio emissions and safety.

What will my PUE be with liquid cooling?

It will be less than 1.1. However, the reduction in power at the server must be included, as PUE does not take this into consideration. See *How does a need for lower CPU fan power help improve efficiency?*

How does a need for lower CPU fan power help improve efficiency?

Liquid cooled computers only need a few watts of fan power, and the CPU will use about 5% less power due to lower leakage current. So, the liquid cooling saves power on the server side, which increases PUE. In general, a cluster cooled by air conditioning with aisle containment will have a PUE of 1.3. one cooled by fans and misters will be at 1.15, and a liquid cooled one will have a less PUE of less than 1.1. The liquid cooled data center will use ~15-40% less power overall due to lower server power and reduced air-cooling power.

Will my hardware last longer with liquid cooling?

Heat and vibration reduce hardware life. With liquid cooling and idling fans, the hardware will last much longer. Spinning hard drives work better in a low vibration environment, according to an ASHRAE white paper, *Hard Disk Drive Performance Degradation Susceptibility to Acoustics*

Will my CPUs speed up?

They can if they are unlocked. Most server CPUs are locked to a specific speed. However, a few researchers have found that the slowest CPU in a cluster runs faster when liquid cooled. So, if the program requires all the nodes to be done with one step before the next step will start, you can expect a speed increase of about 7%. The latest Nvidia A100 GPUs will speed up when liquid cooled.

What is the environmental impact of liquid cooling?

The goal of the Chilldyne system is to make the cost of a liquid cooled data center and server less than that of an air cooled one Also it is to make the liquid cooled cluster more reliable than the air cooled one so that everyone uses liquid cooling, saves electricity and puts less carbon in to the air. Liquid cooling will reduce carbon output as less power is used, the data center is smaller, and less cooling equipment is required. Chilldyne has estimated an annual reduction of 53 million tons of carbon if just half the world's data centers used liquid cooling.

What is the long-term impact of negative pressure liquid cooling?

Chilldyne believes that the negative pressure system will prevail in the long run because it has a minor impact on CDU cost and it has a major impact on cost reduction at the server, rack and interconnect level. In addition, it scales easier because experts are not required for setup, and because fill, drain and coolant additive control are all automated.

How do I sell Chilldyne Liquid cooling?

Every data center operator knows someone who liquid cooled their own PC and had a leak, so leaks are a major concern. Chilldyne manufactures a portable demo system that cools one server that can be taken into the data center for customers to view and cut the plastic tubing line. Once they see what happens when they cut the line they never mention leaks again. You also explain how most of the maintenance is automated, and air leaks into the system are a minor issue. Your customer does not have to a be a pump expert, airplane mechanic and pool boy rolled into one. And there are no worries about ruined racks of wet servers. The Chilldyne solution costs about 3-5% of the server cost. This is about the same as an air-cooling system, and the liquid cooling cost will drop as compared to air cooling as the power per CPU or GPU goes up.

What density does liquid cooling support?

With water based liquid cooling. density is limited by the power into the rack and the power into the CPU/GPU. With direct to the chip water cooling, the power density can be up to 1 kW/cm^2 at the chip level. (AMDs 280-watt Rome CPU is 67 $watt/cm^2$). Chilldyne's parent company Flometrics has tested a water-cooled rocket engine with 2300 $watt/cm^2$ power density.

Where can I find Videos of the Chilldyne system in operation:

Here is a video of how the flow limiting valves work:
https://youtu.be/weHmijmbL6E

Here is a video of what happens when you cut the cooling line on a Chilldyne server:
https://www.youtube.com/watch?v=552tzND2Xx0

Here's a video on a OCP installation of Chilldyne Liquid Cooling:
https://www.youtube.com/watch?v=MzqIouc1T9w

And here's a video of how the Chilldyne CDU works:
https://www.youtube.com/watch?v=w1Ouz7cHhrk