

Data Center Cooling FAQs

Contents

Data Center Cooling FAQs.....	1
What is most important for data center liquid cooling?	3
Which type of Data center liquid cooling is best?	3
How does the Chillydne system work with a leak?	5
How does the Chillydne system drain the servers and racks when they are disconnected during operation?	5
How does the Chillydne system work with only 1 atmosphere of available delta pressure?	6
Does the Chillydne system at higher altitudes with lower ambient pressures?	7
Does the Chillydne system work with overhead and raised floor connections?	7
How does the system still work with a leak?	8
Since liquid cooling has been around for years, it must be mature technology, why shouldn't I buy the cheapest system?.....	9
Why is your technology better than other companies?.....	9
What are the typical failure modes of a liquid cooling system?	10
Why is the Chillydne system better?.....	10
Should I implement Heat recovery with my liquid cooling system?	11
How efficient is the liquid cooling system?	11
Condensation prevention.....	12
What kind of leak detection system is needed?.....	12
How to service a liquid cooled server:	12
How to figure Total Cost of Ownership (TCO) of a liquid cooling system.....	12
What type of coolant to use?	13
How to reject the heat	13
What type of connectors to use?.....	13
What is the recommended flow rate for a given power at what temperature?	13
What pressure is used and what kind of water?	13
How are leaks detected?	14
How to service the liquid cooling system.	14
How does an automatic fail over valve prevent downtime?	14
How should I prepare the site for Chillydne liquid cooling?	14

What type of sensors are used to measure the performance of the liquid cooling system.....	14
What about leak detection?	14
How to prevent corrosion of metal parts or chemical reactions of the plastic with the coolant additive?.....	15
How to drain the system?.....	15
How much heat is captured into the air?	15
Is the Chillyne CDU certified:.....	15
What will my PUE be with liquid cooling?.....	15
Will my hardware last longer with liquid cooling?	16
Will my CPUs speed up?.....	16
How does the CDU work?	16

What is most important for data center liquid cooling?

1. Uptime
2. Uptime
3. Uptime
4. No leaks
5. Server-side cost (servers get replaced every 4 years or so)
6. Infrastructure cost (stays in the data center for number of server refresh cycles)
7. Heat capture ratio (percent of heat into the liquid cooling system, depends on the server and data center design)
8. Cold plate thermal performance (needs to be good enough, but not cost too much, nor have reliability issues)
9. Appearance (people like shiny things, but only if it doesn't cost more)

Which type of Data center liquid cooling is best?

Liquid Cooling with negative pressure CDUs

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. Our approach is failure tolerant, so that you can deploy systems quickly with unskilled labor and not worry about getting expensive servers wet

Direct to the chip cooling:

Not all the heat is captured. With 205-watt CPUs, about 80% of the heat is captured by the liquid cooling system. 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 95-100% of the heat can be captured. Rear doors can

Each server requires a fluid connector and each rack requires a fluid manifold.

The data center layout does not change as compared to a liquid cooled data center. The same racks and hardware can be used.

Liquid Cooling with positive pressure CDUs

Positive pressure CDUs are simple. Connections to the racks require sweated copper plumbing, installed and pressure tested by professional plumbers.

A water leak can ruin an entire rack of servers. Leak detection must be used.

Immersion

The good thing about immersion cooling is that all the heat is captured from the server. The data center layout is completely different. Optical connections must be completed outside of the oil. Spinning hard drives must be sealed. Fans must be removed and fan control software updated. Thermal grease must be replaced with something that won't dissolve in the liquid. Special power and interface cords must be used. The value of the servers in the secondhand market is limited.

Oil Immersion

Widely used to cool transformers and high voltage electrical devices. Can cause embrittlement of plastics when used to cool servers. Immersion is great when you need it, primarily for high voltage utility systems (like a power transformer) or avionics cooling, where the arc resistance at altitude and ability to work down to -40C is needed. Keep in mind that with immersion cooled systems, fixing them is like fixing an automatic transmission. It will be a messy expensive process.

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



Figure 1: Power cable after wicking oil from oil bath. ¹

Two phase immersion cooling

The Wright brothers used boiling water to cool the engine in their 1903 flyer. The Cray-2 used it. Today, many Bitcoin miners are cooled by fluorocarbon fluids. Novec and Fluorinert leach out plasticizer, 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 server board that fails is left in place, so they don't need to open up the sealed box to repair the server.

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."

¹ https://eehpcwg.llnl.gov/assets/sc19_05_1150_1230_impressions_from_5_years_of_oil-based_liquid_immersion_cooling_hofinger.pdf




Figure 2: Tin whisker on power supply²

How does the Chillydyne system work with a leak?

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 .1 in 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 manifold to 2 lpm on the supply side. 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 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)

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

How does the Chillydyne system drain the servers and racks when they are disconnected during operation?

How the CDU keeps from overflowing 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, 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

² <https://datacenters.lbl.gov/sites/default/files/ImmersionCooling2016.pdf>

racks are added to the system, the CDU refills and continues to cool the existing 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, and the air flow into the rack manifold and the CDU is limited by the Venturi. The CDU vacuum pump has a capacity of 1200 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 Chillydyne system work with only 1 atmosphere of available delta pressure?

The less than 1 atmosphere of pressure works because we use 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 we take advantage of the negative pressure to use wire reinforced flexible PVC tubing which is much easier to install and has lower pressure drop than copper tubing with lots of 90 degree elbows. You can see the installation at Sandia below. The pressure drop accounting is in figure XX. The low pressure drop has the added advantage of reducing the power requirements for the liquid cooling system so that they are less than competitor systems. We use about 3.8 In 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.



Pressure Drop Accounting	Flow lpm	DeltaP 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 our system we use a separate heat exchanger pump to pump water from the reservoir through the heat exchangers and back to the reservoir. (You can see a schematic on figure X) This way the HX 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 degrees C at 200 Kw per CDU.)

Does the Chillyne system at higher altitudes with lower ambient pressures?

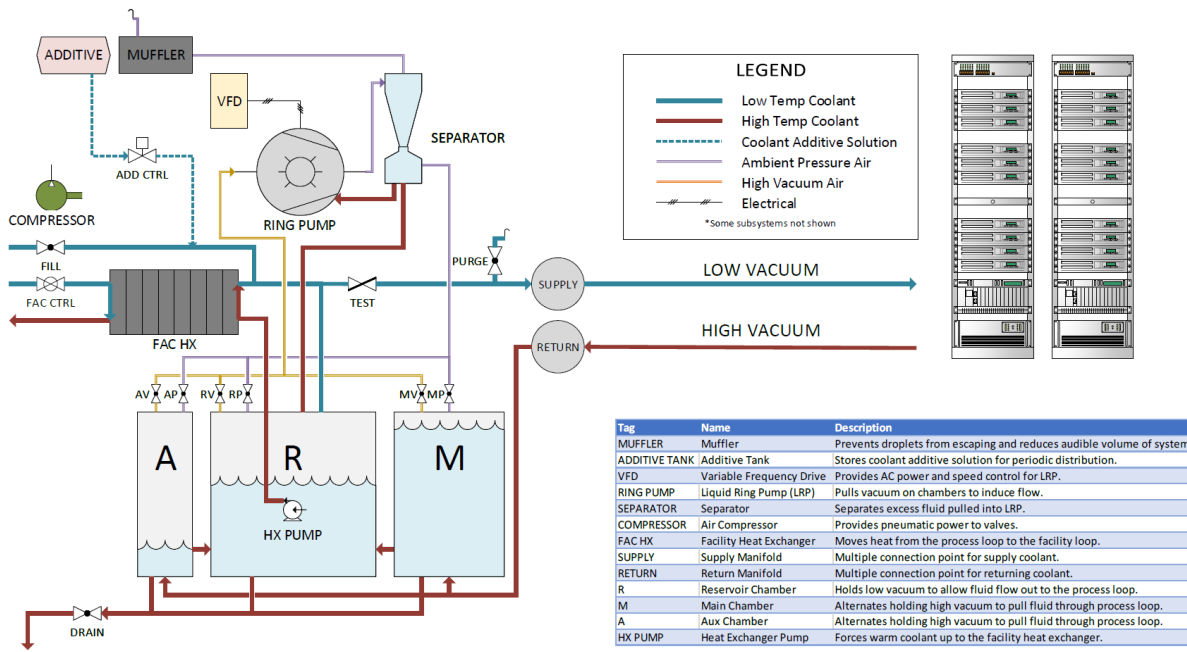
At higher altitudes, the outside air temperature is lower so the available pressure is similar to sea level.

For example, at Los Alamos, at 7000 ft with 55 C return water 55 C, the vapor pressure is 4.4 In Hg, and the absolute pressure is 23.1 In Hg, so the available pressure is 18.7 in Hg. The highest temperature there is about 32C, so the system works with cooling water at 45 C.

In Tucson, the highest temperature there is about 40C we would have return water at 65 C, at 65 C, the vapor pressure of water is 7.4 in Hg total available DeltaP is 22.5-inch Hg.

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

As with the other power and network cables, installation is easier with a raised floor. But overhead works as well. In this installation, there are 90 nodes per OCP rack, dual 165-watt Xeons. Overhead plumbing at 50 kw per rack.



The goal of the Chillydyne system is to make the cost of a liquid cooled data center and server less than the air cooled one, and 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 the air.

We believe 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 level. In addition, it scales easier because experts are not required for setup, and fill, drain and coolant additive control are all automated.

How does the system still work with a leak?

Server flow limiting:

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 .1 in 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 manifold to 2 lpm on the supply side. 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 2 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. (Sonic Nozzle)

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>

Since liquid cooling has been around for years, it must be mature technology, why shouldn't I buy the cheapest system?

Data center liquid cooling is not a commodity. If you pick the cheapest solution, you may spend a lot of money on downtime. It is too early to standardize, unless you want to use the Cray or IBM style of using aerospace fittings and connections, which are too expensive.

We are changing that calculation with low cost, 100% uptime liquid cooling.

No data center operator wants to be the person who ordered the system that caused downtime.

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

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>

We have a solution to the leak detection problem, the CDU measures flow of air out of the system and alerts the user if more than 10 lpm or of air is coming out.

Why is your technology better than other companies?

We have the best team for the technology challenge.

We helped develop the liquid cooling system for the Hunter UAV for Northrop Grumman:

<http://www.flometrics.com/project/aircraft-cooling-system-design/>

We also developed cooling systems for rocket engines, lasers, and medical devices. Parker, Ametek and Lytron are involved in avionics cooling. Parker was at SC19 selling connectors. It seems to me that the major aerospace companies have internal engineering teams that develop cooling systems, and they have outside vendors supply components. 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.

The automotive companies have internal teams working on battery and engine cooling. There are some external consulting groups like, but they don't show any inclination to become a manufacturer. For many of these consultant groups, if they have already consulted on data center liquid cooling under proprietary contracts, they are constrained in developing their own products.

Our advantage is that we have 25+ years of thermal and fluids design experience, plus 9 years of data center liquid cooling experience. We recently developed an avionics cooling system for a defense contractor, and they did not have any other decent bidders on the job.

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, we have flow limiting valves see (Why is your technology better than other companies?) This means that any single server failure will not result in reduced cooling for the other servers. And the server can fail over to air with fins on the cold plates.

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 checks to see if the level, temperature and pressure sensors are reading correctly and if not, it issues a warning. We also check for air flow out of the CDU to see if there is a leak of air into the system, and we issue a warning. In some cases, the CDU can continue to work with a defective sensor

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 XX 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.

Our current practice is to replace any problematic component with ones of higher reliability. If we see any components wear out, we will replace them on a maintenance schedule to prevent downtime. We have an hour meter on the CDU so that we can determine appropriate replacement intervals for components subject to wear.

Why is the Chillydyne system better?

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

- 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
- The flexibility in the complete design of the Chilldyne CDU and computer system, to my knowledge seems to be unique in the approach of having the ability to be either air cooled or liquid cooled with failover valving and built in fans to automatically switch over based on heat capacity
- Data on the system operation can be accessed using Modbus, SNMP, ftp, or the web page.

Should I implement Heat recovery with my liquid cooling system?

When looking at heat recovery, you have to balance 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 and a typical heat pump has a COP of 4, so that It adds 4 KW of heat to the building while using 1 KW of electricity.

This means that if you run your 100 watt processor hotter so it uses 105 watts to get usable heat, and the heat you are replacing requires 25 watts of power to produce 100 watts of heat, the 20 watt power savings for a 100 watts of heat may not justify the expense of hooking the computer to the heating system.

Also, once you hook up a heat source that you can't control (A computer) to a heating system, you need to have a backup source of heat.

I don't think it makes sense for San Diego, but it might for a place where heat is expensive and the computers are running at full load all of the time.

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 you don't use chillers to provide cooling water.

The difference between an efficient and inefficient liquid cooling system is very small.

Heat capture ratio is more important:

Let's suppose that the air conditioning system has a COP of 3.3.

Let's suppose that the liquid cooling system has an efficiency of 25% in terms of coolant pumped and that it runs at 1 kw per 1 lpm. (14 C rise) The deltaP would be 20 psi for positive pressure systems and 8 psi for negative pressure systems).

The flow work for 1 lpm at 20 psi is 1-2.3 watts

So, the liquid cooling system power is 9.2 watts (or less) and the liquid cooling system uses .9% of the server power.

Suppose that the heat capture is 80%, then the HVAC power required is 6% of the server power if we capture 85% of the heat, then the HVAC power is 4.5 % of the server power.

So, the power for pumping the liquid cooling system does not matter compared to the power to remove the heat for small amount of heat not captured.

There is also a processor temperature effect, the processor might use 5% less power liquid cooled, this means that more liquid cooling system flow (and 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.

Condensation prevention

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

What kind of leak detection system is needed?

Each rack will have a temperature sensor on the rack manifold to measure the temperature into and out of the rack and a flow sensor to measure the liquid flow rate through the rack. The Chillydyne negative pressure system never leaks onto a server so there is no need for the leak detection system. The Chillydyne system detects leaks 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 to service a liquid cooled server:

Each server can be serviced by removing the blue inlet connector first and then the red outlet connector after a few seconds. This leaves the server basically dry inside so it then can be shipped back to the manufacturer if necessary or component level disassembly and repair can be accomplished with no worries about spilling coolant onto the motherboard.

How to figure Total Cost of Ownership (TCO) of a liquid cooling system

1. Hardware cost
2. Installation cost
3. Facilities costs associated with installation, such as any 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
5. Annual cooling energy consumption (including chillers, pumps, air handlers, etc.)
6. Annual maintenance of the Servers and supporting mechanical facilities, including annual costs to maintain and support any cooling distribution system required for liquid-cooled SUs.

Liquid cooling systems will reduce air conditioning loads and will reduce fan vibration and server temperatures. The fans will never wear out, in fact less expensive sleeve bearing fans can be substituted for ball bearing fans due to lower temperature and lower fan speeds.

2) The Chilldyne liquid cooling system uses less energy than other liquid cooling systems Sandia measured 1.6% on Attaway.

3) 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.

4) Fill, vacuum test and drain are all automated, so there is no need for experts for any liquid cooling installation or maintenance.

5) The coolant can be replaced while the system is running, so no need for downtime for this activity.

What type of coolant to use?

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 2 x lower viscosity than glycol. This results in 15% more flow with 25% Glycol for the same cooling capacity. 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, better coolant is more important.

How to reject the heat

Cooling towers are best, unless water is scarce, in which case dry coolers can be used.

What type of connectors to use?

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 use approximately one liter per minute of flow to cool one kilowatt of server power. This leads to a temperature rise of approximately 12°C. The use can adjust the temperature rise. 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.

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. We can use tap water although we

recommend using reverse osmosis water in the server along with our additive. Our system includes plastic, stainless steel, copper and brass in contact with the liquid.

How are leaks detected?

For leak detection we utilize an air flow sensor that measures the air flow out of the coolant distribution unit to determine if there is a leak in the system and give a warning to the operator. 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 have already been raised due to the CDU failure

How to service the liquid cooling system.

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.

How does an automatic fail over valve prevent downtime?

We recommend our automatic fail over valve which 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.

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 probably already available in the data center.. The CDU can be fitted with a pump to pump the drained fluid out overhead if necessary.

The tubing to connect the servers and racks to the CDU is all under negative pressure so we use wire reinforced flexible PVC tubing. 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 measure 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. The flow is measured 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 and calculates the heat dissipated by the load on the CDU. The facility input and output temperature is measured and the data is used to determine the facility water flow rate. All the data is available via SNMP or Modbus.

What about leak detection?

No external leak detection system is required. No action is required in the event of a leak, other than a routine maintenance notification. Any server level leak can be repaired during the next business day or whenever is convenient. In the event of a leak, the Chilldyne system will cool the servers with a 1-2°C increase in CPU temperature compared to the CPU temperature before the

leak. The leak can be up to 10 lpm of air, which is the 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 leak. The server with the 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 even of a rack level leak, the system may stop cooling.

How to 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 and additive in the system have 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 .25 mm and is Corrosion tolerant up to 150 micron. 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 to 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.

How much heat is captured into the air?

The most important factors for heat capture are the CPU power and the difference between the cooling water temperature and the data center air temperature. Using cooling tower water and a warm data center will work well. 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 and as the power increases that percentage will also increase. If necessary, we can recirculate the air inside the server to capture even more of the server heat into the liquid cooling system. We have captured up to 90% of heat in lab tests.

Is the Chilldyne CDU certified:

It meets UL and FCC requirements for emissions and safety.

What will my PUE be with liquid cooling?

It will be less than 1.1. However, PUE is not that good of a metric for liquid cooling efficiency. Liquid cooled computers need very low fan power, and the CPU will use less power, so the liquid cooling saves power on the server side, which increases PUE. In general, a cluster cooled by air conditioning with good 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 PUE of 1.1, but it will still use 5-10% less power at the server.

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.

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.

How does the CDU work?

Chillydyne's negative pressure CDU operates under a vacuum which allows for leak-free operation. The chamber system of the CDU, which Chillydyne 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 the Thermosyphon developed by Johnson Controls (more detail provided in Section 1.3). 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

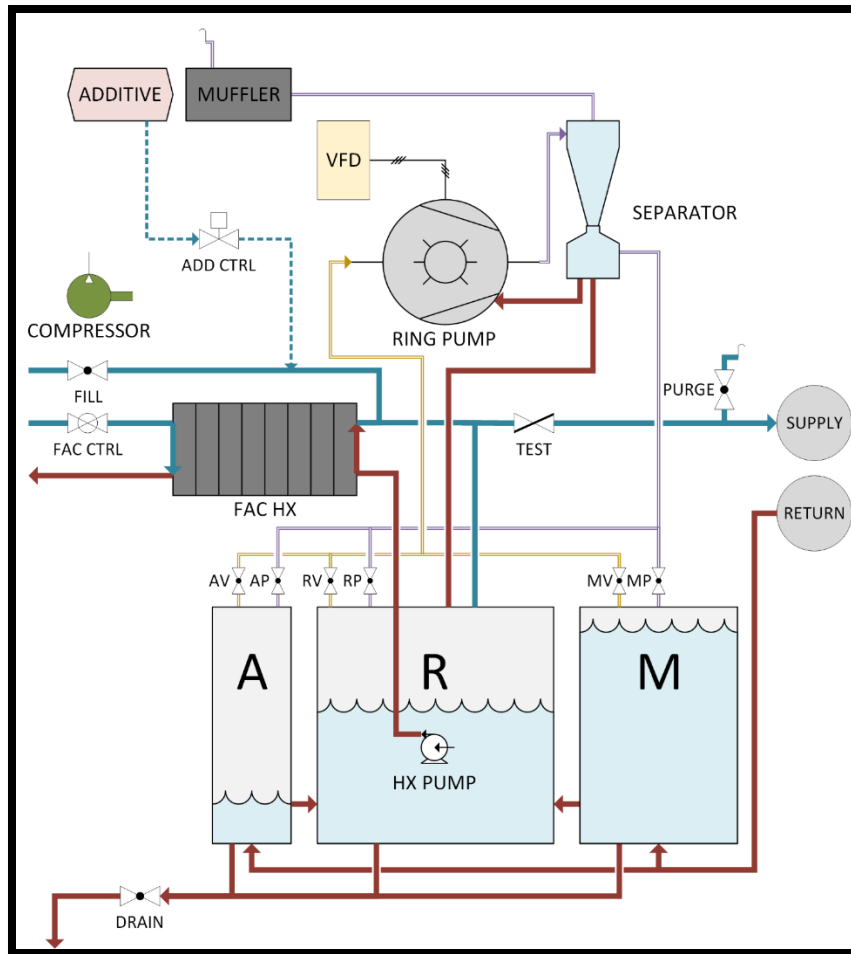


Figure 3: CDU Hardware Layout that depicts the major components of Chilldyne's CDU

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.

