CHILLDYNE

LIQUID COOLING SOLUTIONS

WHITE PAPER

FAQ GUIDE TO DATA CENTER LIQUID COOLING

Answering 42 Questions on Liquid Cooling Techniques, CDUs, Negative Pressure Technology, and Direct-to-Chip Cooling Economics

Rev X.4

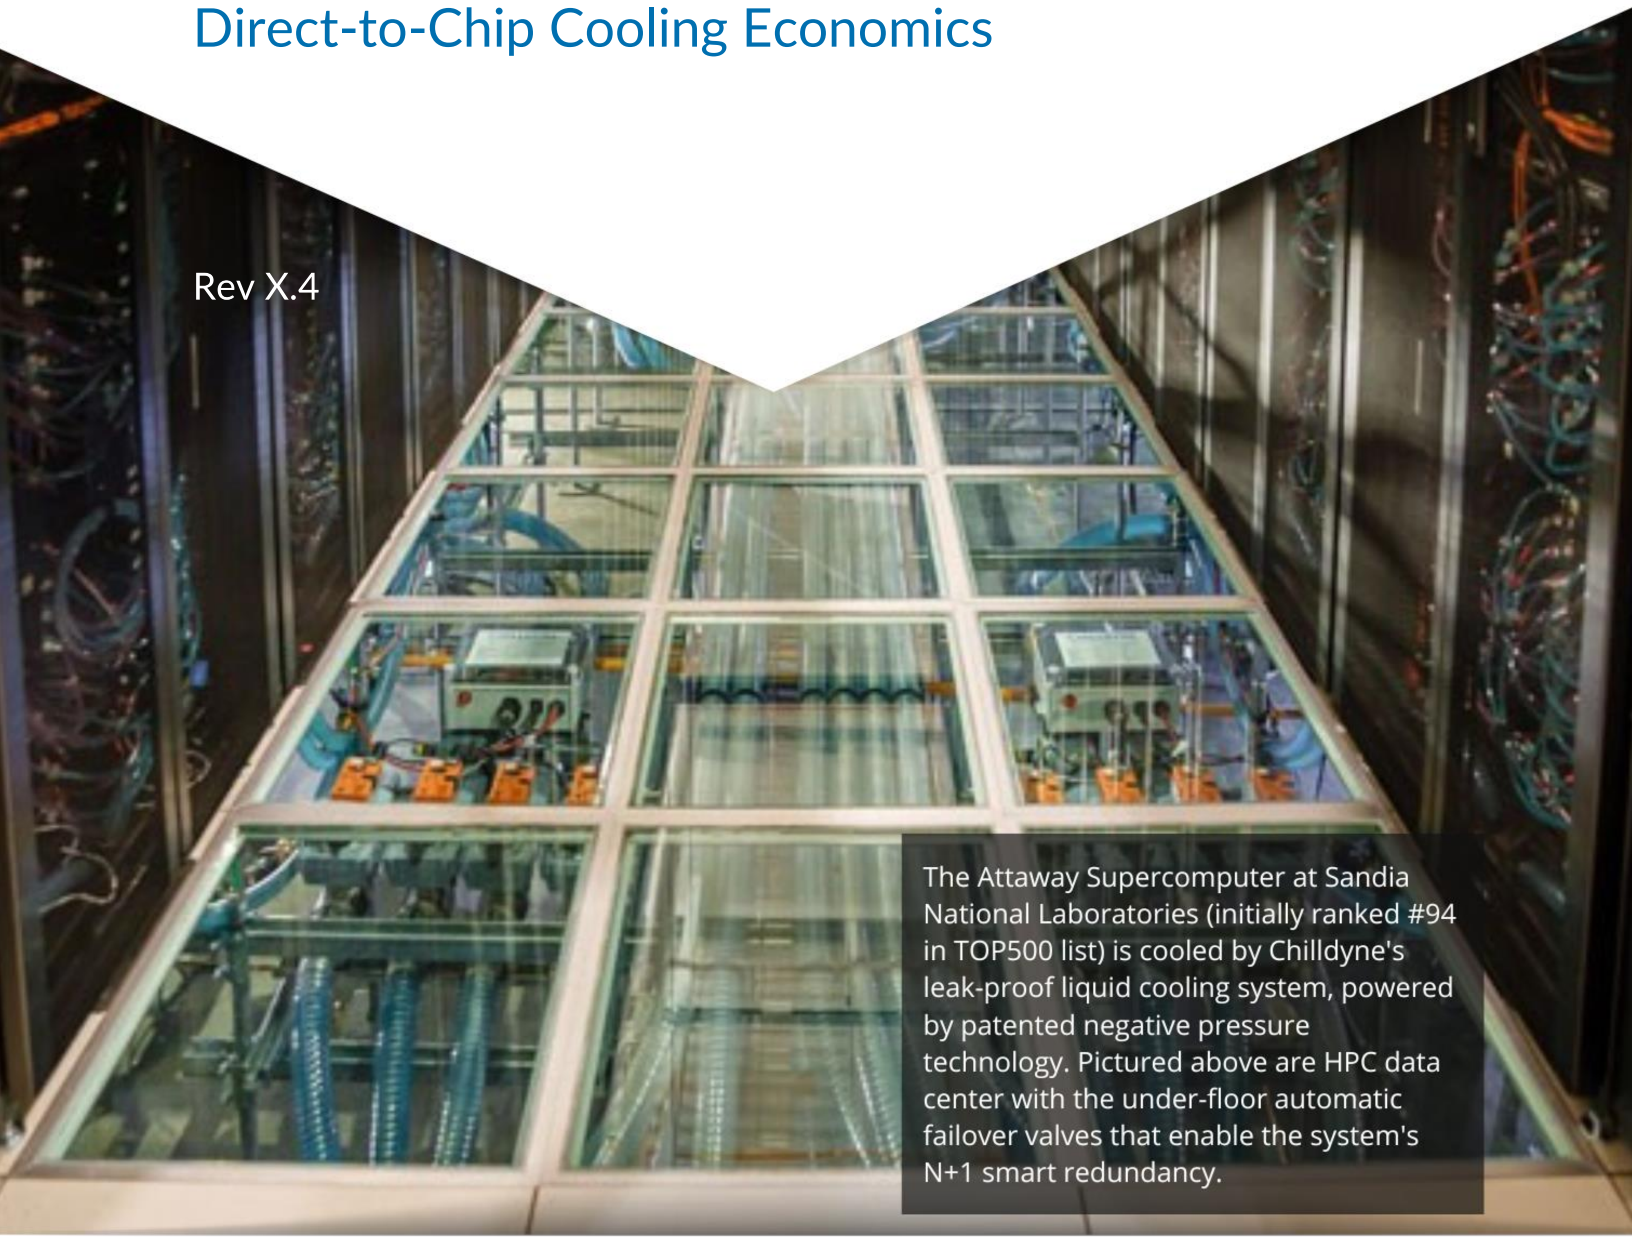

The Attaway Supercomputer at Sandia National Laboratories (initially ranked #94 in TOP500 list) is cooled by Chilldyne's leak-proof liquid cooling system, powered by patented negative pressure technology. Pictured above are HPC data center with the under-floor automatic failover valves that enable the system's N+1 smart redundancy.

Table of Contents

Section 0: About the FAQ

Executive Summary

This comprehensive FAQ guide addresses 42 critical questions about liquid cooling, exploring various methods, benefits, technical specifics, and the unique advantages offered by Chilldyne's technology. It serves as an essential resource for understanding the intricacies and advantages of implementing liquid cooling in data centers, offering insights into efficiency improvements, cost savings, and environmental impacts.

The guide covers a range of topics, from the basics of liquid cooling to advanced Chilldyne-specific technologies like negative pressure systems and cooling distribution units (CDUs). It provides detailed insights into how liquid cooling contributes to energy savings, enhances hardware longevity, and supports the acceleration of CPU performance. Additionally, the document delves into practical considerations, installation procedures, maintenance aspects, and leak management strategies, providing a holistic view of liquid cooling systems in modern data centers.

Authored by Dr. Steve Harrington, CEO of Chilldyne, this guide draws upon his extensive experience in fluid dynamics and electronics cooling. Dr. Harrington's expertise shapes the content, ensuring that readers receive accurate, up-to-date, and relevant information. The FAQ aims to empower data center operators, IT professionals, and decision-makers with the knowledge to make informed choices about adopting liquid cooling technologies, ultimately contributing to more sustainable and efficient data center operations globally.

About the Author: Dr. Steve Harrington, CEO of Chilldyne

Dr. Harrington started Flometrics in 1995 while working on a Ph.D. in aerospace engineering, focusing on fluid dynamics. At Flometrics, he helped design laser cooling systems, medical ventilators, unmanned aerial vehicle (UAV) cooling systems for Northrop Grumman, improvised explosive device (IED) excavation robots for the Department of Defense, and rocket fuel pumps for DARPA.

Throughout his professional and personal life, Dr. Harrington worked with airflow and liquid flow technology. He taught aerospace engineering at UCSD, where his students designed, built, and flew liquid-fueled rockets. Dr. Harrington holds more than 30 patents in total.

In 2011, Dr. Harrington founded Chilldyne to bring his engineering, fluid dynamics, and electronics cooling expertise to the data center. His personal goal is to reduce energy consumption and the carbon impact of data centers globally by deploying liquid cooling to as many servers as possible. Widespread adoption of liquid cooling could save 2.5 gigatons of carbon over the next 30 years. Chilldyne's goal is to make this carbon reduction a reality by offering an easy-to-use, low-cost, and reliable liquid cooling system.

Contextual Note

In this document, the term 'liquid cooling' primarily refers to the direct-to-chip liquid cooling method (single phase), unless specified otherwise. For detailed information on this specific liquid cooling method, please refer to pages 8-10.

Section 1: Introduction to Liquid Cooling

What are the most important considerations for data center liquid cooling?

1. **Uptime**
    - The most important aspect of data center liquid cooling is ensuring continuous operation without interruption.
    - **Business continuity:** Data centers are critical for the continuous operation of numerous businesses and services. Any downtime can lead to significant disruptions, affecting everything from financial transactions to essential communications, as seen in this recent example: https://www.reuters.com/business/finance/singapores-central-bank-tells-dbs-citibank-investigate-system-outage-2023-10-19/
    - **Financial impact:** Data centers are expensive (millions of dollars) and downtime can lead to substantial financial losses.
2. **No Leaks**
    - Leaks can cause damage to expensive server equipment, leading to costly repairs or replacements.
    - Leak prevention is crucial to prevent operational disruptions and equipment damage.
3. **Server-side Cost:** Balancing performance with the cost for servers that are replaced approximately every 4 years.
4. **Infrastructure Cost:** Considering the long-term investment in cooling infrastructure over multiple server refresh cycles.
5. **Heat Capture Ratio:** % of heat captured by the liquid cooling system, which depends on server and data center design.
6. **Cold Plate Thermal Performance:** Effective heat transfer, balancing efficiency, cost, and reliability.

Why liquid cooling?

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.

How do I figure out the 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

Section 2: Liquid Cooling Methods and Additional Info

What types of data center liquid cooling exist?

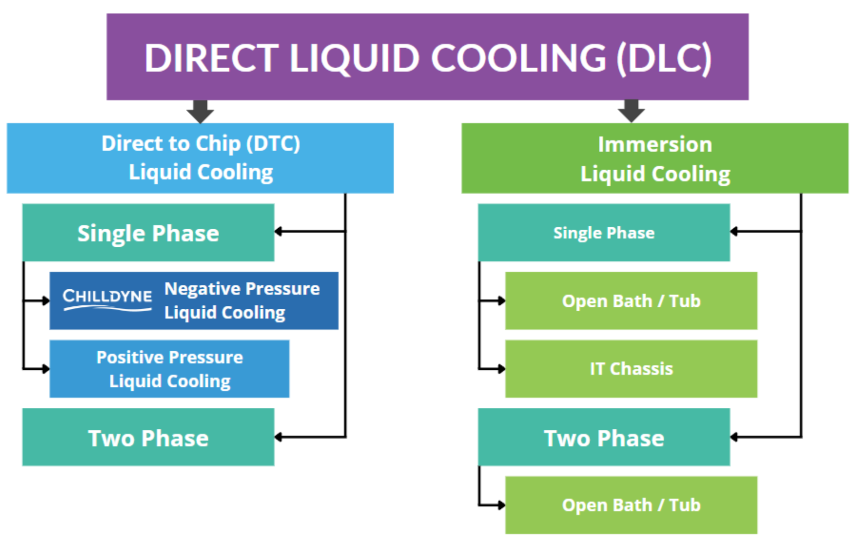

Figure 1: Liquid cooling classifications

Direct-to-Chip Liquid Cooling Method

General information

With direct-to-chip cooling, liquid coolant flows directly to hot components like CPUs and GPUs via cold plates mounted on top of the processors inside the server, without putting IT components in direct contact with the liquid. (Please refer to Figure 4)

Not all the heat is captured. With Intel servers using 205-watt CPUs, about 80% of the heat is captured by the liquid cooling system. As the CPU power increases in the future, the percentage of heat captured will go up to 90% or more. If the facility cooling is based on cooling towers or dry cooler, and the data center runs warm, very little power is needed to remove the heat (about 4% of the server power).

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.

In this document, the term 'liquid cooling' primarily refers to the direct-to-chip liquid cooling method (single phase), unless specified otherwise.

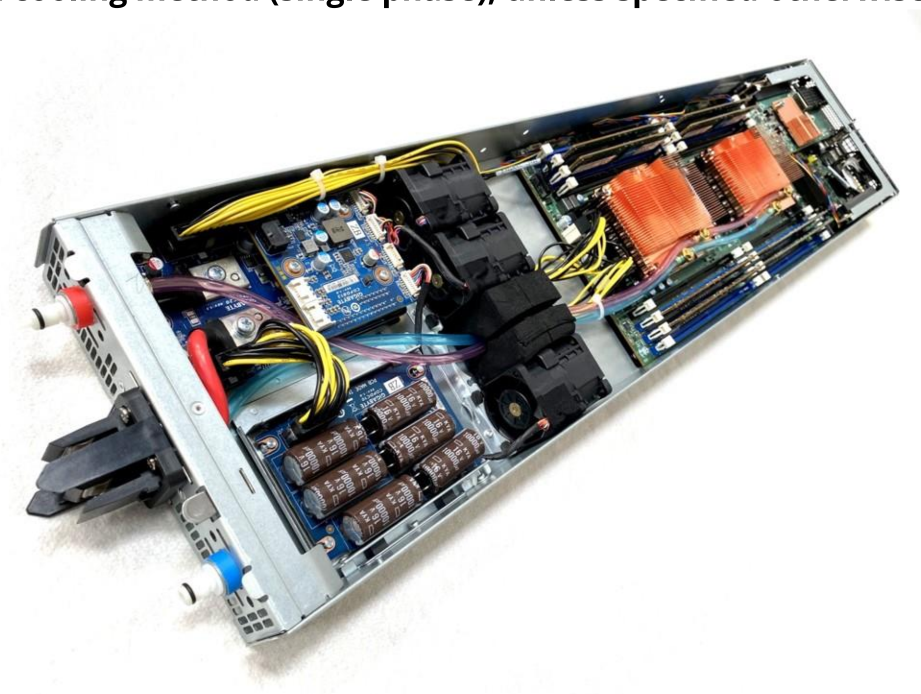

Figure 2: An OCP node featuring copper cold plates specifically for direct-to-chip liquid cooling.

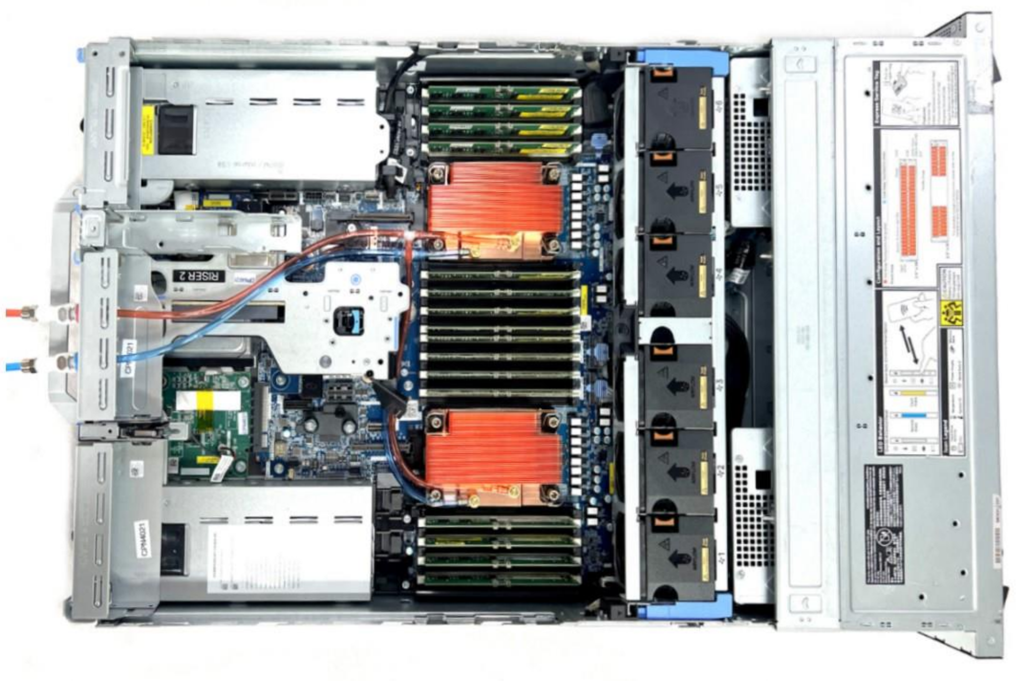

Figure 3: A Dell R760xp server modified for direct-to-chip liquid cooling, with copper cold plates prominently visible.

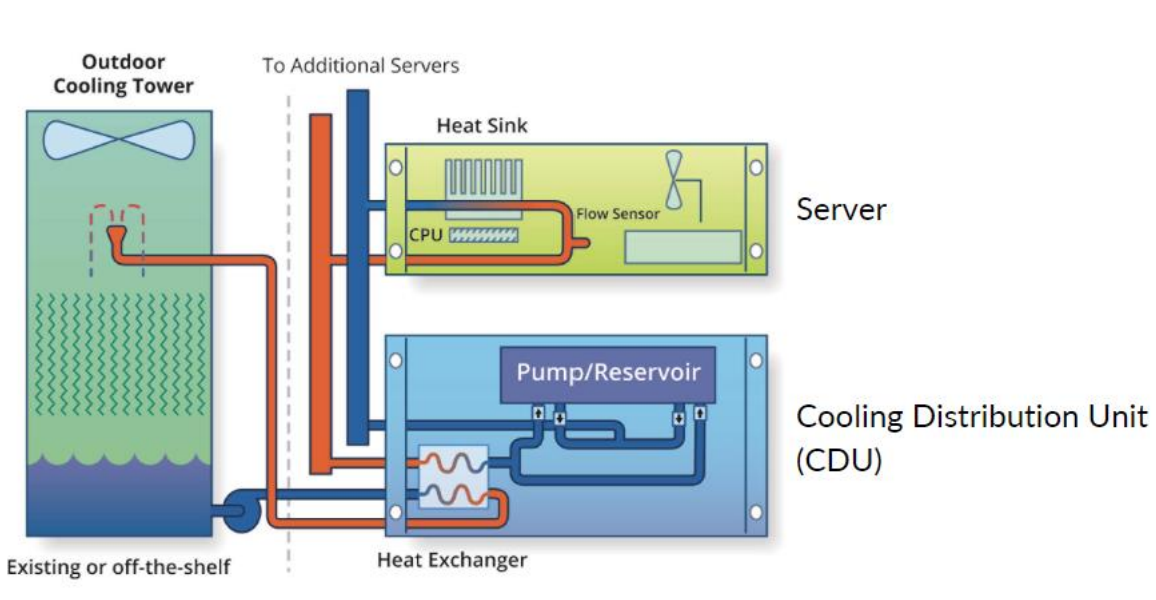

Figure 4: Basic Block Diagram of direct to chip liquid cooling system. Use a cooling tower or a dry cooler for best efficiency.

Direct-to-chip liquid cooling with negative pressure CDUs (Chilldyne) – Single phase

Chilldyne's negative pressure CDUs enhance the reliability of direct-to-chip liquid cooling.

- A leak does not damage the server; a leak allows airs into the liquid cooling system.
- This design allows the use of cost-effective, generic connectors, parts, and plumbing, reducing overall costs, complexity, and installation time. (Any data center liquid cooling system will have thousands of connections, and there are many ways that they can leak.)
- Chilldyne's approach is particularly advantageous, being failure-tolerant and suitable for rapid deployment, even with unskilled labor, without risking damage to expensive servers.

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

Here is a video of how negative pressure technology prevents leaks: https://youtu.be/St5cUGCjm0U

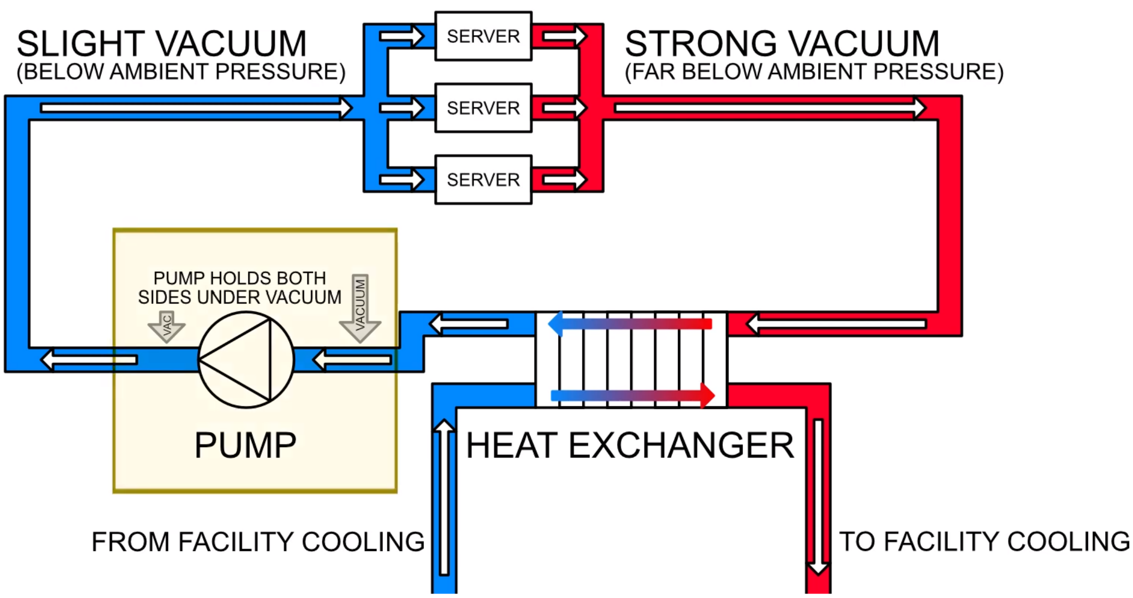

Figure 5: Diagram of liquid cooling system using negative pressure CDU technology, which pulls the coolants from both sides under vacuum.

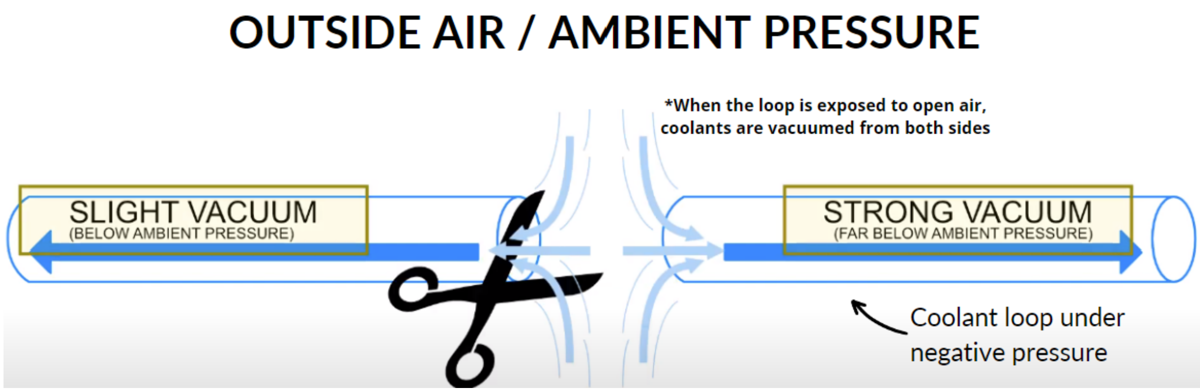

Figure 6: Detailed illustration showing how negative pressure technology prevents coolant leaks when a line is cut before the leak occurs. The cooling loop is under vacuum on both supply and return side.

Direct-to-chip liquid cooling with positive pressure CDUs - Single phase

Any liquid cooling system CDUs that are not explicitly "negative pressure" are positive pressure.

They require sweated copper plumbing, welded stainless or plastic for connections to the racks, which must be installed and pressure tested by professional plumbers.

A water leak can severely damage an entire rack of servers. Leak detection must be used. Failure isolation also proves more challenging compared to the inherently leak-proof negative pressure system.

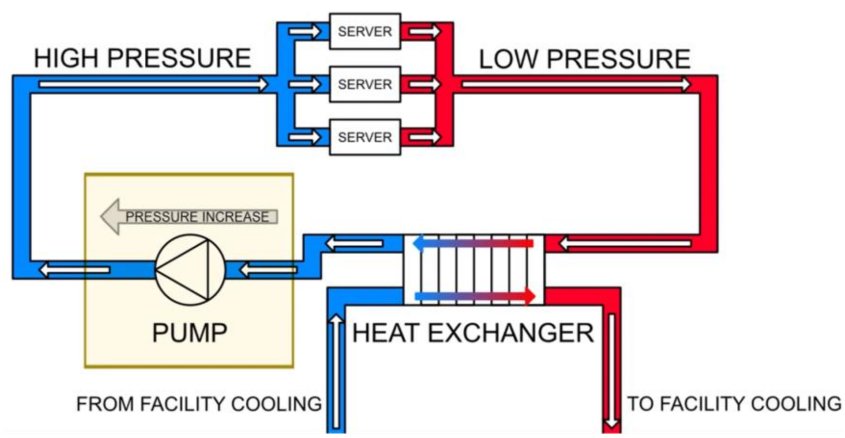

Figure 7: Positive pressure liquid cooling system pushes the coolant from one end to another, making it susceptible to liquid leakage.

Direct-to-chip liquid cooling - Two-phase

Two-phase direct-to-chip liquid cooling differs from single-phase by using a fluid that transitions between states - from liquid to gas - to remove heat. This is the same technology that is used in air conditioning systems. This method employs expensive engineered dielectric fluids containing fluorine-carbon bonds, which may be regulated or banned. The vaporized dielectric is either transferred to an external condenser or used to release heat into a building's water loop. Thermal performance is optimal for a narrow temperature range so the system is less flexible compared to water-based systems.

Immersion Liquid Cooling Method

General information

Immersion liquid cooling involves direct contact between the liquid coolant and IT electronic components. Servers are fully or partially immersed in a dielectric liquid.

The good thing about immersion cooling is that all the heat is captured from the server. However, it requires specific adjustments:

- 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.

Single-phase oil immersion cooling

With the tub method (also referred to as open bath), the IT equipment is completely submerged in a dielectric fluid, usually oil-based. With traditional IT racks, the servers are horizontally stacked from the bottom to the top of a rack. However, because this method uses a tub, it's like laying a traditional rack of servers on its back. Instead of pulling servers out on a horizontal plane, tub immersive servers are pulled out on a vertical plane. Often, centralized power supplies provide power to all the servers within the tub. The heat within the dielectric fluid is transferred to a water loop via heat exchanger using a pump or natural convection.

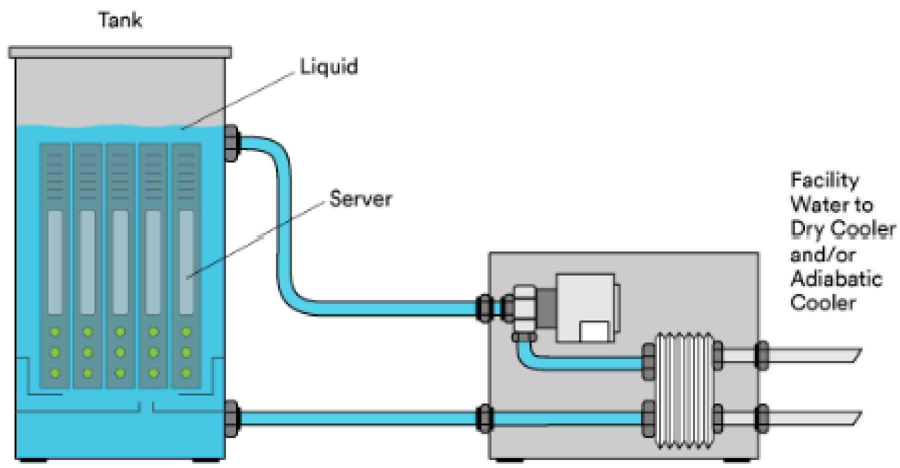

Figure 8: In single-phase immersion cooling, electronic components are directly immersed in a non-conductive, dielectric liquid within an enclosure. The fluid remains in its liquid phase. (Courtesy of 3M)

Chassis-based immersion cooling encapsulates servers within a sealed chassis, cooled passively by the dielectric fluid or actively with forced convection.

Immersion cooling is widely used to cool transformers and high voltage electrical devices. It can cause embrittlement of plastics when used to cool servers. Immersion works well for high voltage utility systems (like a power transformer) or avionics cooling, where the arc resistance at altitude and ability to work down to -40°C is needed.

Keep in mind that with immersion cooled systems, fixing them is like fixing an automatic transmission. It will be a messy and expensive process.

Two-phase immersion cooling (open bath / tub)

Just like the single-phase tub method, the IT is completely submerged in the fluid. However, it uses a specialized two-phase dielectric coolant which transitions from liquid to gas for heat absorption.

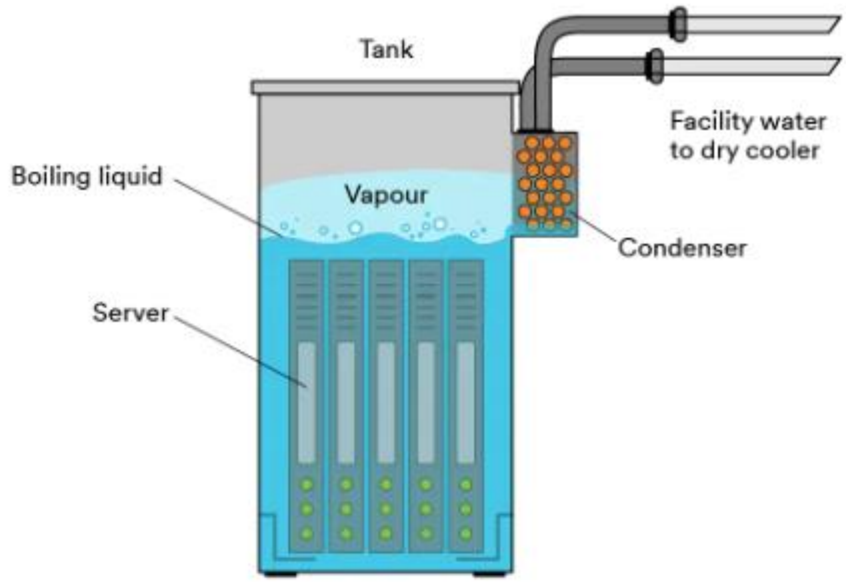

Figure 9: In two-phase immersion cooling, fluid is boiled and condensed. The vapor condenses on a heat exchanger (condenser) within the tank, transferring heat to facility water that flows outside of the data centers. (Courtesy of 3M)

The Wright brothers used boiling gasoline to cool the engine in their 1903 flyer. The Cray-2 used it. Modern applications including Bitcoin mining are cooled by fluorocarbon fluids. Novec and Fluorinert leach out plasticizer, leak easily, evaporate quickly, and are very expensive. (3M's announcement means they will eventually cease to be produced.) 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.

How is the heat rejected from a liquid cooling system?

Cooling towers are best, unless water is scarce, in which case dry coolers can be used. Evaporating water to cool servers is better for the environment than using power to run chillers, because power requires carbon, even in countries like France and Sweden (which use nuclear and hydro power).

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

Data center direct to chip cooling remains an emerging, rapidly evolving arena. Chip power is going up quickly, warm water cooling is still in the early phase. While liquid cooling principles are mature, translating innovations into reliable large-scale infrastructure is not always straightforward. Data center liquid cooling systems that are assembled from commodity hardware may not be reliable and vendors may not provide long term support. (No data center operator wants to be the person who ordered the system that caused downtime.)

Simply put, bargain hunting risks substantial, costly downtime from leaks, corrosion, hardware mismatches and other failures. It is too early to standardize, unless you want to use the Cray or IBM style of using bespoke systems with aerospace fittings and connections, which are too expensive.

Chilldyne systems are designed for hassle-free operation with features like leak prevention, automatic filling and draining, and automatic coolant monitoring for 100% uptime liquid cooling. For instance, Sandia's DOE Attaway cluster operated for several years without a single leak and no cooling-related downtime since its inception.

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

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

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

1. *Leak in the server*

Problem: Leaks can occur due to improperly connected hoses or broken Quick Disconnects.

Chilldyne Solution: Coolant leaks are prevented through negative pressure technology. Chilldyne equips each server with flow-limiting valves to ensure a single server failure doesn't impact others. In certain designs, servers can switch to air cooling if needed by using fins on the cold plate. (Refer to Figure 14 for the flow-limiting valves.)

2. *CDU failure*

Any CDU failure will result in the activation of a switchover valve to switch to a backup CDU (this is our current recommended practice for system redundancy with N+1 CDUs). At the CDU level, the software monitors levels, temperatures, and pressure, alerting for irregularities. 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. (Refer to Figure 16 for the automatic switchover valve.)

CDUs are designed to avoid applying positive pressure to servers during outages. They have specific valve positions to vent the system, preventing damage. 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 proactively 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 use an hour meter on the CDU to schedule maintenance and replacements, and determine appropriate replacement intervals for components subject to wear.

What will my PUE be with liquid cooling?

Liquid cooling can achieve a PUE of less than 1.1. However, PUE isn't the most effective for liquid cooling efficiency. Liquid-cooled computers require minimal fan power, and CPUs often consume less power, paradoxically leading to a higher PUE. For general comparison, an air-conditioned cluster with good containment typically has a PUE of around 1.3, one cooled by fans and misters will be at 1.15, and a liquid cooled one, which will still use 5-10% less power at the server, will have a PUE of 1.1 or less.

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.

Section 3: Benefits of Direct-to-Chip Liquid Cooling

How efficient is the liquid cooling system?

If the data center is air-cooled, air cooling will use 10-30% of server power for fans (server and data center fans) and air conditioners. Liquid cooling reduces this to about 2%, assuming no chillers are used for cooling water. The efficiency difference between efficient and less efficient liquid cooling systems is minor.

The heat capture ratio is more significant:

- Suppose an air conditioning system has a COP of 3.3.
- Let's assume the liquid cooling system is 25% efficient in terms of coolant pumped, running at 1 kW per 1 lpm (14°C rise). The deltaP is 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, making the liquid cooling system's power 9.2 watts or less and using about 0.9% of server power.
- If heat capture is 80%, HVAC power needed is 6% of server power; at 85% heat capture, it's 4.5%.
- The power for pumping in the liquid cooling system is negligible compared to removing the heat not captured.

- The processor temperature effect: a processor might use 5% less power when liquid-cooled, meaning more flow (and power) in the liquid cooling system can result in lower overall data center power.

So, balancing the benefit of running processors hotter (and saving on pumping power) is crucial against the extra power for processors running hot and additional cooling needed to maintain the data center air temperature at a comfortable level for personnel. Contact Chilldyne if you want us to run these calculations to optimize your data center liquid cooling system.

Will my hardware last longer with liquid cooling?

Yes, liquid cooling can extend hardware longevity. Heat and vibration reduce hardware life. Liquid cooling significantly reduces these; especially in a system where fans are often idle, meaning the hardware will last much longer.

Will my CPUs speed up?

Potentially. They can if they are unlocked. While most server CPUs have locked speeds, liquid cooling can lead to faster performance, especially in clusters. Research indicates that the slowest CPU in a cluster tends to run faster when liquid-cooled. This can be particularly beneficial in programs requiring all nodes to complete one step before moving to the next, as it may lead to an overall increase in processing speed. Performance increases of 7% have been measured at SDSC and LANL.

How does liquid cooling contribute to energy savings compared to traditional cooling methods?

Liquid cooling is significantly more energy-efficient than traditional air cooling due to the superior thermal conductivity of water. Direct to chip water is about 700% more efficient in terms of thermal resistance compared to air. This means it can absorb and transfer heat more effectively, making liquid cooling systems much more capable of dissipating heat from high-performance computing components. The higher efficiency in heat transfer leads to substantial energy savings, as it reduces the need for additional cooling mechanisms that consume more energy.

What are the environmental impacts of using liquid cooling in data centers?

- Individual Data Center: 450 tons of carbon emission reduction per year (based on based on OCP Cluster Manzano located at Sandia rated at 960 kW)
- Global: 2.5 gigatons of carbon emission reduction in 30 years

$$\sum_{k} GCO2 \cdot P_{data} \cdot Savings \cdot Fraction_{liquid_k} \cdot ((1 + 6.4\%))^{k} = 2.504 \text{ gigaton}$$

- Key assumptions
  - Global carbon emission = 11 gigatons
  - Global data center power consumption % to global power consumption = 1%
  - Data center growth = 6% per year
  - Carbon intensity of electric power decreases 1.25% per year
  - Power savings = 35%
  - 6.4% annual growth in data centers
  - Liquid cooling adoption for 75% of data centers in 8 years

Section 4: Chilldyne Liquid Cooling Technology (Negative Pressure and CDU)

Why is the Chilldyne liquid cooling system better?

- **Leak-proof through Negative Pressure Technology**: Eliminates leak risks, enhancing system reliability.

  - See it in action, here: https://www.youtube.com/watch?v=552tzND2Xx0
- **Ease of Maintenance**: Simplifies node accessibility and replacement.
- **Simple Leak Detection**: Air bubbles in the return line indicate leaks, facilitating quick response.
- **Energy Efficiency**: Lower power draw compared to positive pressure CDUs.
- **Design Versatility**: Supports both air and liquid cooling with automatic switchover valves and fans.
- **System Monitoring**: Integrated with Modbus, SNMP, FTP, and web interface for comprehensive control.

Figure 10: Screenshot from a video showing a cooling line being cut on a Dell R760xp. As soon as the coolant loop is exposed to air, the liquid vacates from both sides, eliminating the risk of coolant leaks through negative pressure technology. Watch the full video here: https://youtu.be/_eq8BPy5tzM

What are the top 3 reasons to use negative pressure technology for your liquid cooling system?

1. Using negative pressure virtually eliminates risk of (coolant) leaks.
2. Lower cost components; cold plates, plumbing and connectors are 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.
3. Easier, faster, less complicated installation and maintenance.

Reducing the threat of leaks, coupled with easier and less expensive installation, lower cost components, better reliability, and lower total cost of ownership make the negative pressure technology to consider in building next generation data centers, cloud computing facilities and supercomputers.

What are the four key components of a Chilldyne liquid cooling system and their functions?

1. Cooling Distribution Unit (CDU)

The CF-CDU300 cools up to 300kW of servers. CDU acts as both the pump for the system and the heat exchanger.

2. Rack Manifold

Delivers coolant to more than 100 servers while installed inside a standard server rack.

3. Cold Plates

Mounted directly to server processors, to remove heat at its source.

4. Cool-Flo Software

Remotely monitor and operate the Chilldyne CDU using a web-based interface, industry-standard protocol, or integration with DCIM and BMS software.

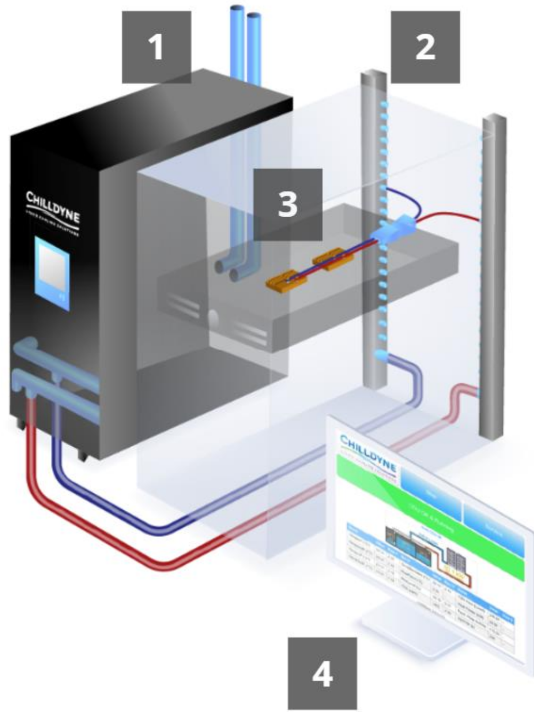

Figure 11: Anatomy of a Chilldyne Liquid Cooling Solution, illustrating the four main components – CDU, rack manifold, cold plates, and the software.

How does the Chilldyne's Cooling Distribution Unit (CDU) work?

Chilldyne's negative pressure CDU operates under a vacuum, enabling leak-free operation. The CDU's chamber system, known as the "ARM'' chamber (Auxiliary, Reservoir, Main), both pumps the coolant and stores it. It's divided into three smaller chambers: Auxiliary, Reservoir, and Main. The CDU's pumping action 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 to 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. 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. Disconnecting a server automatically evacuates the water, leaving the server dry for maintenance.

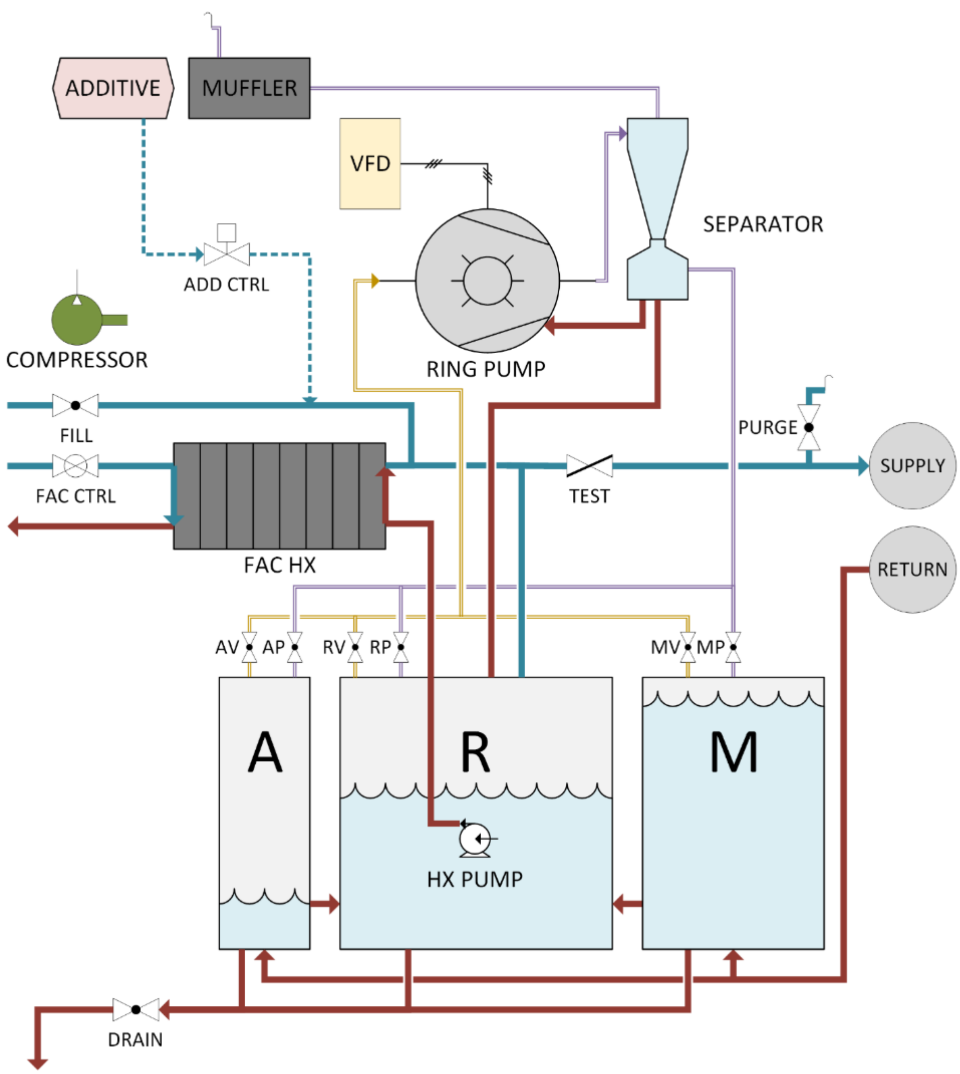

Figure 12: Chilldyne negative pressure cooling distribution unit (CDU) schematics.

Tag	Name	Description
MUFFLER	Muffler	Prevents droplets from escaping and reduces audible volume of the 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 points for supply coolant.
RETURN	Return Manifold	Multiple connection points 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 13: Chilldyne CF-CDU300 with its side covers off, showing off the ARM chamber.

Is the Chilldyne CDU certified?

It meets UL and FCC requirements for emissions and safety.

What coolant does Chilldyne system use?

The Chilldyne system uses either 1) reverse osmosis water with anti-corrosion and anti-bacterial additives or 2) polyglycol 25% (PG25) as coolant.

For most applications, we recommend reverse osmosis water because it provides superior thermal performance compared to glycol mixtures in several ways:

- Over 4% higher heat capacity, absorbing more heat
- Half the viscosity, meaning 15% more flow for PG25 is required for the equivalent cooling capacity of 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, better coolant is more important.

How does the Chilldyne system work with a leak?

- *Server flow limiting*

The system's server connections consist of two key components: a check valve on the supply side and a sonic nozzle Venturi on the return side. These components play crucial roles in managing air flow during a leak.

In normal conditions, the check valve's flow resistance is about 0.1 inches Hg, and the Venturi's is around 1 inch Hg. If there's a significant air leak, the check valve's controlled leak feature restricts air flow into the manifold to 2 liters per minute (lpm). (Refer to Figure 14 for the check valve.) This results in coolant bubbles for servers downstream of the leak. However, the lower bulk density of the coolant increases the volume flow rate, slightly raising the downstream server temperatures by 1 to 3 degrees Celsius. Despite this, the system continues to effectively cool all servers, except the one with the leak, using liquid cooling.

On the return side, the Venturi caps air flow at about 10 lpm. The air flow here is limited by the sonic nozzle's design, which prevents flow from exceeding the speed of sound in the Venturi's narrow section. (Refer to Figure 15 for the sonic Venturi.)

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 leak prevention or failure tolerance.

For a visual explanation of how these flow-limiting valves function, refer to this video: 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

Figure 14: Check valve with precise leak flow limiter on supply side

Figure 15: Sonic Venturi flow limiter on return side

Section 5: Installation, Maintenance, and Operation

How should I prepare the site for Chilldyne liquid cooling?

The CDUs need access to facility cooling water, tap water (reverse osmosis is recommended) and drain pipes. Typically, these water supply systems are already in place for existing CRAH, 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 easily by your data center technicians, without needing specialized plumbers. 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.

Can you compare the installation processes for positive pressure and negative pressure liquid cooling systems in a real use case?

Positive Pressure Government Use Case: For a positive pressure cooling system, the installation process is quite extensive and labor-intensive. It took over 2 weeks to set up the plumbing for 25 racks. This involved a professional plumber performing tasks such as cutting and soldering pipes, threading, and using expensive, pressure-rated materials. The complexity and need for specialized skills in this process mean that scalability is often limited by the availability of skilled labor in the area.

Negative Pressure Government Use Case with Chilldyne: It took 1 week to install; the process was much cleaner and simpler. Notably, no plumbing professionals were required. The installation was carried out by unskilled labor using flexible plastic tubing and basic tools like a knife and wire cutter. Note: for the negative pressure system to function properly, the CDUs need to be within 100 feet of the racks and on the same level.

How do I service a Chilldyne liquid cooling system?

The Chilldyne liquid cooling system is designed for continuous operation with no shutdown required for changing filters, coolant additive or coolant.

Each server can be easily removed, replaced, and hot-swapped with only a few seconds of time required to disconnect the fluid couplings.

How do I 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 swimming pool chemicals and 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 does the Chilldyne system drain the servers and racks when they are disconnected during operation?

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.

- How the CDU keeps 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, 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 and added racks with no downtime. It also automatically adds more coolant additives as required.

Section 6: Monitoring and Leak Management

How does Chilldyne's negative pressure system handle leak detection and management?

The Chilldyne negative pressure system never leaks onto a server so there is no requirement for an external leak detection system. Each rack may 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.

Impact on Operation: 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 (air, not coolant), 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. For leak detection, we utilize an air flow sensor in the CDU which measures flow of air out of the system and alerts the operator if more than 10 lpm of air is coming out. Once a leak is detected it can be located by following the bubbles in the coolant return line to the source of the leak.

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.

Backup Cooling: Depending on the configuration, the server with the leak will have backup air cooling via fins on the cold plate so that it can continue to operate at reduced clock speed and power dissipation. In the event of a rack-level leak, the system may stop cooling.

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 has already been raised due to the CDU failure.

How does an automatic switchover valve prevent downtime?

We recommend our automatic switchover 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. This way, a CDU downtime does not cause cluster downtime. (Refer to Figure 16 for the automatic switchover valve.) This is similar to the standard procedure for air cooling with N+1 HVAC systems.

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

Chilldyne's system uses precise thermistor-based temperature sensors in rack manifolds to measure inlet and outlet temperatures with 0.2°C accuracy. Flow is gauged using a vortex flowmeter, accurate to 1.5% full scale (1.-26.4 GPM). The CDU also measures flow rate, input/output temperatures, and calculates 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.

Section 7: Chilldyne Technical Specifications and Other Features

Why is condensation prevention important in liquid cooling? How does Chilldyne address condensation prevention?

Condensation prevention is crucial in liquid cooling systems to protect servers from potential water damage, which can lead to hardware failure and data loss. Condensation occurs when coolant temperature falls below the dew point, leading to moisture accumulation on server components, so the liquid cooling system must never provide the servers with water that is below the dewpoint, even if the servers are off.

The Chilldyne system uses an onboard humidity sensor to measure the dewpoint and controls the coolant output temperature with redundant systems. Even if a temperature sensor fails, the system still provides coolant above the dewpoint.

How does Chilldyne prevent corrosion and chemical reactions with coolant additives?

The cold plates all have 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 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 from debris getting in the system.

How does the Chilldyne 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, meaning if inner diameter, ID, of a tube goes from ½ inch to 5/8 inch (13-16mm) the pressure drop is 3x less.) 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 in Figure 16, which prominently features the tubing designed for minimal pressure drop.

The pressure drop accounting is in Figure 17. The low pressure drop has the added advantage of reducing the power requirements for the liquid cooling system so that energy consumption on the negative pressure system is 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.

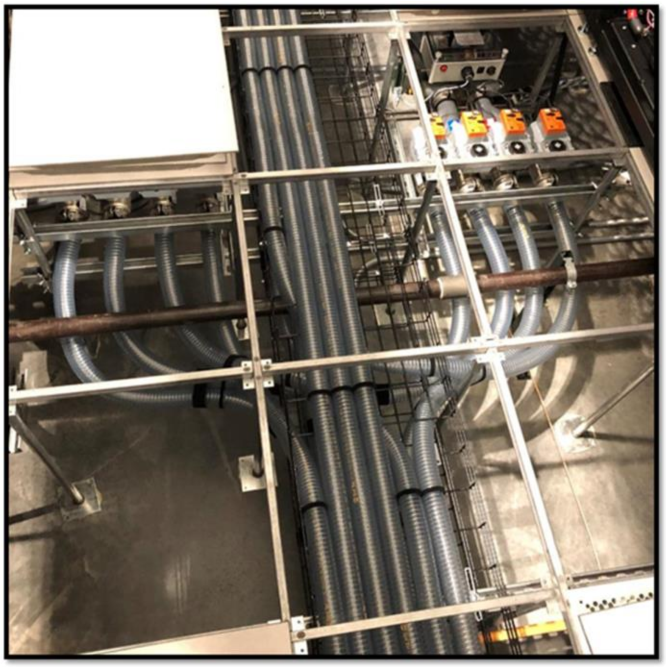

Figure 16: Chilldyne's liquid cooling deployment at Sandia National Laboratories, showing the flexible PVC tubing to server racks and automatic switchover valves.

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" Tubing	49.6	0.4
Failover Valve	198.4	0.1
40 ft 2" Tuving	198.4	2
Rack (subtotal)	1.6	
CDU-Rack (subtotal)	3.1	
Total ΔP	**8.5**	
Available ΔP @ 200 lpm	**14 avg, 19 max**	

Figure 17: Pressure drop accounting table. Available delta P is based on sea level elevation.

The Chilldyne system uses a separate heat exchanger pump to pump water from the reservoir through the heat exchangers and back to the reservoir (Please refer to Figure 12). This way 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 °C at 200 kW per CDU.)

What connectors/interconnects are suitable for Chilldyne liquid cooling?

The Chilldyne negative pressure system uses common, affordable generic liquid connectors because a leak is not a high priority issue. The negative pressure system does not require more expensive aerospace-quality connections.

Figure 18: On the left is a more expensive, dripless liquid cooling connector. On the right is a generic fluid connector. Generic parts are possible with Chilldyne's negative pressure liquid cooling system because leaks or broken connectors do not cause downtime on other servers.

Here's an article how negative pressure technology can reduce costs by 1,400% on connectors: https://chilldyne.com/2023/09/07/liquid-cooling-economics-reducing-costs-by-1400-on-connectors-through-negative-pressure-technology/

What flow rate does Chilldyne recommend for cooling specific server power, and what are the optimal water temperatures?

We recommend one liter per minute of flow to cool one kilowatt of server power (1 LPM / kW). This leads to a temperature rise of approximately 14°C. Users can adjust the temperature rise depending on the server and available facility water. We can use water from 2°C to 45°C as the input. Colder water results in more heat capture at the server level, lower server power consumption, and greater reliability for the semiconductors.

What are the water pressure levels and types of water used in Chilldyne's cooling system?

Chilldyne's liquid cooling system operates under vacuum conditions, and the water pressure is with the water pressure being approximately 4 inches of mercury vacuum on the supply side and 18 inches of mercury vacuum on the return side. While tap water can be used, we recommend reverse osmosis water in the server along with the suggested additives.

Section 8: Practical Considerations

Should I implement heat recovery with my liquid cooling system?

Implementing heat recovery in a liquid cooling system requires considering the balance between the increased power usage at higher CPU temperatures and the efficiency of heat pumps. (Please note that a one to one comparison of watts of heat to watts of computer power is not appropriate) For example, a hotter running processor might consume 5% more power. Given that a typical heat pump has a Coefficient of Performance (COP) of 4, this translates to 4 kW of heat added to the building for every 1 kW of electricity used.

If a 100-watt processor runs hotter, consuming 105 watts to produce usable heat, and the heat being replaced requires 25 watts of power for 100 watts of heat, the 20-watt power saving might not justify the system cost. In some areas the warm water coming out of the servers can be used as a source for a heat pump that raises the temperature to a higher value.

Moreover, integrating an uncontrollable heat source like a computer into a heating system necessitates a backup heat source. The feasibility depends on factors like local climate and consistent computer load. For example, it might not be practical in San Diego, but could be beneficial where heating costs are high and computers consistently run at full load.

What are the long-term cost-benefits of switching to liquid cooling? How does Chilldyne's liquid cooling system reduce TCO specifically?

1. Liquid cooling systems will reduce air conditioning loads, 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 temperatures and lower fan speeds.
2. The Chilldyne liquid cooling system uses less energy than air-cooled systems. The Chilldyne system also uses less energy than other liquid cooling systems. At Sandia, we measured 1.6% less power consumption of the liquid cooling system.
3. The system includes onboard coolant additives to energy reduce the need for periodic water testing to be done by facility employees. 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.

How does liquid cooling affect the overall data center infrastructure?

Better Energy Efficiency: Liquid cooling is more efficient than air cooling, primarily because water has a higher heat capacity than air. This means it can remove more heat per unit volume, reducing the energy required for cooling. The power to pump air for cooling is much more than pumping water for the same amount of cooling.

Increased Hardware Density: Liquid cooling allows for more efficient cooling of high-performance components, which means data centers can host more powerful hardware in a smaller footprint.

Reduced Cooling Infrastructure: Traditional air-cooled data centers require extensive cooling infrastructure, including air handlers, HVAC systems, and extensive ductwork. Traditional chillers, compressors, condensers etc. can be downsized by enabled by liquid's efficiency.

Optimized Real Estate / Layout: Back-to-back server rows are possible. Increased density can lead to better utilization of space and potentially reduce the physical size of data centers.

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

Yes, we support both raise-floor and overhead configurations.

In the Sandia National Labs installation, there are 62 nodes per OCP rack, dual 205-watt Xeons, with a raised floor.

Does the Chilldyne system work 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.

- Example 1: Los Alamos (elevation: 7,000 ft / 2,100 m)

In Los Alamos, New Mexico, at 7,000 ft altitude with 55°C return water, the vapor pressure is 4.7 in Hg (Source: Wikipedia article "Vapour pressure of water"), and the absolute pressure is 23.1 in Hg (Source: Wikipedia article "U.S. Standard Atmosphere"). The available pressure is calculated by subtracting vapor pressure from absolute pressure, so the available pressure is 18.4 inch Hg. The highest temperature there is about 32°C, so the system works with cooling water at 45°C.

- Example 2: San Diego (elevation: near sea level / 0 m)

In San Diego, California, with 60°C return water, the vapor pressure is 5.9 in Hg and the absolute pressure is 29.9 in Hg, so the available pressure is 24 in Hg. The highest temperature ever recorded in San Diego was 44°C, so with 10°C heat load, the system would work with cooling water temperature at 50°C.

Editor's note: The point of the question is just to address the viability of high altitudes, not to spec out an entire system. There are many other considerations that go into deciding final temperatures for a complete system.

Section 9: About Chilldyne

What is Chilldyne's approach and long-term vision when it comes to driving wider liquid cooling adoption?

The goal of the Chilldyne 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 (liquid cooling system components constitute less than 5% of the total data center cost) and it has a major impact on cost reduction at the server level (i.e., much lower electricity bill compared to air-cooled data center). In addition, it scales easier because expert technicians and professional plumbers are not required for setup, and fill, drain, and coolant additive control are all automated.

Why is your technology better than other companies?

Chilldyne liquid cooling systems are designed by rocket scientists and built by medical device engineers. 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 developed cooling systems for rocket engines, lasers, and medical devices.

- We recently developed avionics cooling system for a defense contractor, and they did not have any other decent bidders on the job.

Cooling business is complex:

- Parker, Ametek, and Lytron are involved in avionics cooling. It seems to us 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, but none are inclined 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.

Chilldyne's advantage lies in our over 25 years of thermal and fluid design expertise, plus over 10 years of data center liquid cooling experience.

Chilldyne, Inc.
5900 Sea Lion Place #150
Carlsbad, CA 92010
www.chilldyne.com
+1 (760) 476-3419 | info@chilldyne.com