BOSTON
Servers | Storage | Solutions

Asperitas
IMMERSED COMPUTING®

WHITEPAPER

IMMERSED COMPUTING®

THE EVOLUTION OF COOLING – WHY IMMERSION IS THE FUTURE

EXECUTIVE SUMMARY

> *For as long as computing has taken place, there has been the debate about how to efficiently and safely cool systems and data centres. This paper discusses the history of cooling to present day innovations and technologies, addressing common fears and issues associated with mainstream cooling practices, products and technologies.*

The faster the processor is, the hotter it tends to be when in use. Computer cooling is required to remove the excess heat produced by computer components, to keep them within safe operating temperatures. Central Processing Units (CPUs), graphic cards and hard disk drives are likely to temporarily or permanently fail if overheated. All modern-day processors are designed to cut out or reduce their voltage or clock speed if the internal temperature exceeds a specific limit. The use of heatsinks can also reduce any temperature rises by carrying the heat away through its thermal conductor. Too much heat can lead to a shorter life-span of your hardware, therefore it's imperative that real thought is put into the right cooling for your data centre.

For as long as computing has taken place, there has been the debate about how to efficiently and safely cool systems and data centres. This paper discusses the history of cooling to present day innovations and technologies, addressing common fears and issues associated with mainstream cooling practices, products and technologies.

THE EVOLUTION OF COOLING

The obvious way to cool anything is with air. Computational air cooling is achieved typically by numerous fans within a system chassis working to lower air temperature by evaporating heat – the faster the fan spins, the faster it evaporates. Within a data centre a common idea is to separate cold and hot air which is achieved by facing the cold sides away from the hot sides and this ultimately leads to the cabinets cooling themselves. Air cooling systems provide increased air flow and reduced temperatures with the use of cooling fans. However, most air-cooled systems use a combination of fans and heat sinks. Unfortunately, traditional air cooling uses fans for a lot of different components such as graphic cards, CPU's or processors and for drawing air inside and outside the case. Therefore, this makes it bulkier and a messy set up, not to mention inefficient and costly.

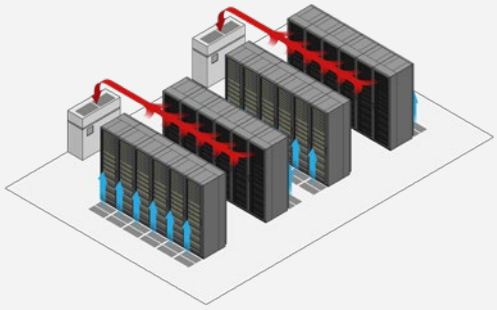

Figure 1: Data centre hot and cold aisle rack/cabinet configuration/layout. Arrows show flow of hot and cold air. Cold air enters from raised floor. Hot air is drawn into air conditioners.

THE ADVENT OF LIQUID COOLING

Liquid cooling is more efficient in removing any excess heat. This is due to liquid coolant having higher heat capacity, density and thermal conductivity than air which allows the system to transfer more heat over larger distances. Users can choose to cool specific components that are prone to heating such as hard drives and power supplies.

The search for an alternative cooling technology can be traced back to the 1960's, it was IBM that introduced liquid cooling to large mainframe computers. The concept was simple; cold water would run through a contraption that cooled it down to below room temperature and then the water would run directly to the computer system. Whilst it cannot be doubted that this innovation made its mark, the industry was still wary to adopt liquid cooling because of ongoing maintenance and support concerns.

During the 1990's 3M created an immersion cooling fluid that was used to aid the cooling of industrial equipment. This product has evolved to be known as Novec™, a non-conductive, non-flammable and non-toxic liquid. When utilised, according to 3M1 "Heat generated by the electronic components is directly and efficiently transferred to the fluid, thereby reducing need for other active cooling components, such as interface materials, heat sinks and fans. These improvements increase energy efficiency and enable higher packingdensities".Despiteinnovation,immersion cooling in the data centre was not widely adopted.

TYPES OF LIQUID COOLING

There are a vast variety of liquid cooling technologies available, and it is a field that is constantly evolving. Below we discuss the prominent technologies available today, inside of which sit several sub-categories. That said, of these varieties, not all liquid cooling systems are created equal!

CHILL PLATE

A chill plate cools components via surface conduction. As the name suggests, a cool heat sink or chill plate is fixed to the component, heat is then transferred conductively through the plate to the cooling medium via cooling tubes or internal cavities in the plate.

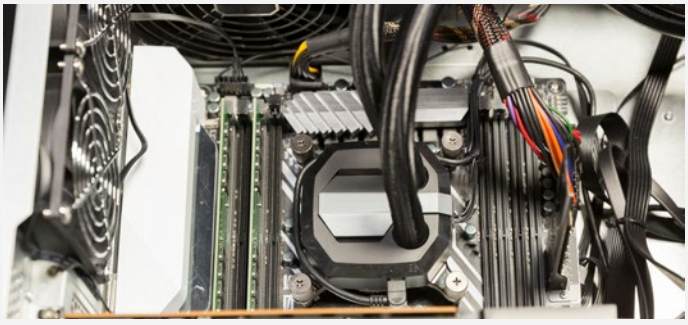

Image 2: A chill plate cooling a CPU in a gaming workstation with liquid pipes in and out of the plate.

TYPES OF LIQUID COOLING

PHASE COOLING

A single-phase coolant never changes state as it has a higher boiling point and always remains in its liquid phase. Single-phase is efficient and inexpensive when it comes to cooling electrical components as it eliminates all pressure and fumes that are created by a coolant's state transition from liquid to gas.

In two-phase cooling, the working fluid boils and therefore exists in both a liquid and gas phase. Throughout that change of state, it takes the heat to the top in a gasified form and then through condensation, the gas cools down and returns to liquid form which starts the whole process all over again.

THE ELEPHANT IN THE ROOM

Poorly installed systems are likely to result in leaks. It is crucial to ensure the system is resistant to any possible leaks by testing it outside the IT before installing it to avoid any damage. Even a single drop of liquid can ruin a component. Due to this, vendors are more hesitant to provide warranties... but what about if you were to immerse your entire data centre in liquid?

TOTAL IMMERSION

Total immersion cooling is the practice of completely immersing electric components in a dielectric liquid. Through this process all the heat generated via the hardware is captured in the liquid. Comparatively a suitable liquid is able to capture 1500 times more heat energy than air with the same volumes and temperatures.

Total immersion cooling systems tend to operate in one of two principles, either with a pump to circulate and exchange heat or via natural convection.

The primary whilst effective, uses more energy than the latter. Within the data centre there are also a number of additional infrastructure requirements to consider including access to and drainage of water, the active maintenance of submerged systems and cleaning of spillages. Whilst all of these are easily addressed, the adoption of totally immersed data centres has only recently begun to take place.

Air cooled rack	Immersed Computing®
An air cooled rack can commonly support about 5kW of IT power, and takes the space of 2 floor tiles. In order to get sufficient air through the rack, 1 or 2 floor tiles are needed in front and in the rear of the rack, also for service space. The total footprint of the rack now becomes 600x2400/3600 mm. 5kW of air cooled IT divided over 1U servers is comparable to 3 kW of immersed IT. This results in a power density of **1,5-2kW per m²**	Immersed Computing® can support 22 kW of pure IT power with a footprint of 600x1200 mm. No airflow is required but we'll assume a service area of 600x1200 mm with each module. This results in a power density of **more than 15 kW/m²**

Image 3: A table illustrating the major differences between an air cooled rack and Asperitas Immersed Computing®

BOSTON LIQUID COOLED SOLUTIONS

In 2011 at EAGE in Vienna it was announced that Boston Limited and Hardcore Computer had signed a deal to see Hardcore's machines built for the first time outside of the USA.

The Hardcore Detonator was billed in the dual press release as; "... a high performance liquid cooled workstation that has been optimised to push the performance boundaries of today's processor graphics controllers and storage technologies further than ever before enabling designers, engineers and scientists to tackle the most demanding of to modeling and visualisation, to accelerate development."

Primary heat generating components in the workstation form factor were submerged in CoreCoolant™, a safe, odor-free, dielectric liquid that absorbed heat more efficiently than conventional liquid cooling systems of the same era. Continuous exposure to heat has been proven to reduce the reliability and longevity of electronic components. Submersion cooling protects temperatures allowing them to run far cooler and at peak performance in comparison to solutions using traditional air cooling techniques.

AWARD WINNING RESULTS -2014

Boston continued to develop industry ready solutions for a variety of workloads and achieved success at the ISC Student Cluster Competition in 2014 in partnership with CoolIT. The student team from EPCC at The University of Edinburgh achieved a 1st place ranking for the Highest LINPACK in the ISC14 Student Cluster Challenge.

The student team, made their first place mark by recording a score of 3.38 Tflops/kW with the system ranking at an estimated #4 on the Green500 list. It was the first time a team has broken the 10 Tflops barrier in under 3kW. the CoolIT Rack "DCLC AHx cooling system", mounted directly onto the Intel Xeon E5-2680 v2 CPUs and NVIDIA K40 GPUs.

AWARD WINNING RESULTS - 2014

The student team celebrating their win at ISC14

The system allowed both the processor and GPU accelerator heat output to be directly absorbed into circulating liquid, which then efficiently transported the heat to a liquid-to-air heat exchanger mounted on the top of the rack. By pairing CoolIT's Rack DCLC technology with Boston's HPC servers the team were able to divert more power for compute and less on cooling. The student group was then able to reduce the power required for cooling by removing a number of fans from the server and allocate more of their resources to compute nodes.

LOOP HEAT PIPE (LHP) TECHNOLOGY

It was at another supercomputing tradeshow, this time SC17, that Boston announced a further liquid cooled solution. The Boston Quattro 12256-T "LHP edition" was based on a 2U, 4-node server featuring Intel® Xeon® Phi processors (codenamed Knights Landing).

In collaboration with LHP experts Calyos, Boston have developed this server SKU by utilising a Calyos cooling solution based on LHP (Loop Heat Pipe) technology. The innovative 2-Phase cooling solution was designed for the data centre and provided several advantages such as improved thermals, better efficiency, lower power consumption and extended server operating temperature when compared to a standard air-cooled equivalent.

BOSTON QUATTRO 12256-T "LHP EDITION"

THE BEST WAY TO PREDICT THE FUTURE IS TO CREATE IT - IMMERSED COMPUTING®

At CloudFest, Rust, Germany in 2018 Boston announced a strategic partnership with Immersed Computing® specialists Asperitas. The AIC24 is Asperitas' flagship module that requires only power, access to a water loop and data connectivity to operate. The simplicity of the design and installation means that the solution is essentially "plug-and-play" making it an ideal solution for the modern data centre.

LOOKING TOWARDS A GREENER DATA CENTRE

The Asperitas AIC24 is a closed system and is the first water cooled system which is driven by natural convection to circulate the dielectric liquid meaning no pumps are needed and mechanical components don't need to move.
The immersed environment allows for warm water cooling and has very minor temperature fluctuations. Total immersion ensures that no oxygen encounters any components inside the system which prevents oxidation. In order to ensure the system can run all year round, a simple dry-cooler system was deployed.

Asperitas say: The circulation of the liquid is where Immersed Computing® is unique and ground-breaking. Asperitas eliminates any infrastructure for the primary medium, oil. Instead, the oil is circulated by natural means and does not leave the immersion system. The oil circulates by the heat generated by the IT and water cooled "Convection Drives™". This means that the primary circulation is completely driven by the thermal expansion of the oil and gravity. The only requirement for heat rejection is common water infrastructure.

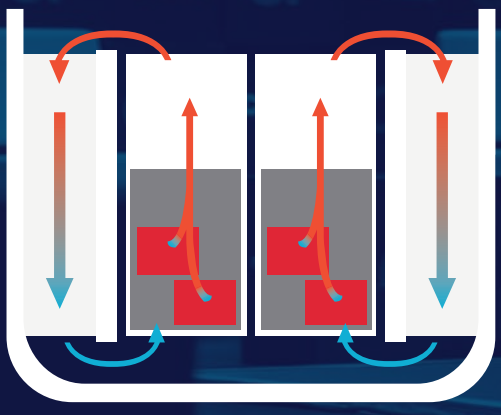

Image 8: An illustration of the circulation of liquid in the AIC24

To demonstrate the flexibility of the AIC24, Boston and Asperitas have conducted live demonstrations at a number of tradeshows using a portable AIC24 solution, populated with compute sleds built by Boston. The most recent demonstration took place at CloudFest 2019 where the companies officially launched four plug and play solutions for cloud and compute, CPU and GPU-based workloads.

Image 9: The AIC Demo unit at CloudFest 2019

The solutions are optimised for different purposes, ranging from generic platforms with full stack support and extensive integrated storage, to high density compute and GPU compute applications. All solutions are designed and optimised for synergy between server systems and the Computing® technology concept for optimal performance.

At the same event the companies laid down their vision for expanding this portfolio of application centered solutions with a specific focus on emerging digital technologies like AI.

CONCLUSION

As computer technology evolves and becomes more complex, the way we keep our systems cool is important now more than ever. Though air cooling methods have historically increased air flow and reduced system temperature, they've often been part of bulkier systems with a clumsy set up. Aiding the path of innovation, developments such as the AIC24 liquid cooling system from Asperitas, has emerged as a preferred solution eliminating the use of oil (without leaving the immersion system) and successfully removing heat. Not only that, the modern-day technology also keeps the system cool with zero noise level for any complex computer system. Therefore, the more component heavy the system, the more powerful the technology becomes to ensure the most optimal cooling for the system.

REFERENCES

https://www.3m.co.uk/3M/en_GB/novec-uk/applications/immersion-cooling/

https://multimedia.3m.com/mws/media/1127920O/2-phase-immersion-coolinga-revolution-in-data-center-efficiency.pdf

https://asperitas.com/wp-content/uploads/2017/06/Asperitas-Immersed-Computing.pd

https://www.3m.co.uk/3M/en_GB/novec-uk/applications/immersion-cooling/

https://www.datacenterknowledge.com/archives/2012/10/08/3m-demos-new-immersion-cooling-technique

https://pressreleases.responsesource.com/news/64945/boston-limited-launch-the-hardcore-detonator-the-ultimate-workstation-for/

https://www.hpcwire.com/off-the-wire/boston-limited-coolit-join-forces-sponsor-epcc-student-team/

https://www.epcc.ed.ac.uk/blog/2014/06/27/teamepcc-isc14-student-cluster-competition

https://www.calyos-tm.com/cooling-technologies

https://www.ctmmagnetics.com/different-types-liquid-cooling/

ABOUT BOSTON

Boston Limited has been providing cutting edge technology since 1992 using Supermicro® building blocks. Our high performance, mission-critical server and storage solutions can be tailored for each specific client, helping you to create your ideal solution. From the initial specification, solution design and even full custom branding – we can help you solve your toughest business challenges simply and effectively.

CONTACT

+44 (0) 1727 876 100
sales@boston.co.uk
www.boston.co.uk

BOSTON
Servers | Storage | Solutions