WHITE PAPER

# Keeping Emmy Cool

## How Innovative Hybrid Data Center Cooling supports energy efficient operations for Göttingen University's leading Supercomputer in a new prefabricated Modular Data Center

Release 1, April 2024

STULZ MODULAR

# Table of Contents





sales@stulz-modular.com | www.stulz-modular.com

# Management Overview

High performance computing has become incredibly energy intensive due to more powerful CPUs and GPUs being specified as well as higher rack densities being incorporated into data centers and server rooms. As the heat load rises, so the cost and effectiveness of traditional air cooling in data centers is also increasingly under stress.

It's not just inside the data center that things are warming up. Extreme ambient temperatures experienced throughout Europe during the summer of 2022 also added to the cooling burden. As the outside temperature increased, so data center cooling infrastructure came under additional pressure, resulting in some very high-profile outages.

The reality for many legacy data centers is that climate change will create operational challenges for which they have not been designed to cope.

Energy efficiency is vitally important for cost efficient data center operations. However, the primary goal of most data center designs is to ensure reliable operation of the IT equipment and maximise uptime. Increasing heat density means that traditional room-focussed approaches to cooling using air as the medium for heat removal are no longer fully effective and present a growing risk of thermal shutdowns.

This white paper details a hybrid data center cooling solution, comprising the use of direct-to-chip liquid cooling for hotter server components including processors and memory devices, operating in concert with row-based air cooling to remove heat caused by other components and server room equipment, as an effective and efficient response to the outlined challenges.

The solution is in operation in a modular facility designed and installed by STULZ Modular at the University of Göttingen, where it houses and provisions Emmy, a leading supercomputer. The data center comprises four prefabricated modules; two larger modules joined along the spine to accommodate the liquid cooled supercomputer, and two smaller modules also joined along the spine to accommodate air-cooled IT equipment. A single door allows access to the complete data center (see Appendix 4).

The hybrid cooling system features CoolIT passive cold plate Direct Liquid Cooling (DLC). The overall solution proposed by STULZ was required to meet the University's demand for a low, <1.14 PUE and reduced carbon footprint for its data center operations. In fact, the efficiency of the hybrid liquid and air-cooling system has delivered an overall annual facility PUE of 1.13, with 1.07 PUE for the liquid cooled supercomputer room alone.

Improving energy efficiency is a top priority task, as increasingly complex simulations and bigger AI models cause computational demands to rise, while legal, economic, and environmental restrictions limit the possible energy usage. In concrete terms, the University of Göttingen and the UMG want to use the measures available to them to achieve climate neutrality by 2030. A corresponding Climate Protection Statement was adopted in June 2021.

For more details visit https://www.uni-goettingen.de/en/654994.html

# Introduction

The University of Göttingen is home to Emmy, one of the top 100 most powerful supercomputers in the world. The name was chosen to honour the German mathematician, Emmy Noether, described by Albert Einstein (amongst many other peers) as one of the most important women in the field of mathematics.

As part of Germany's National High-Performance Computing (NHR) alliance, the University of Göttingen assigned the operation of Emmy to the Gesellschaft für Wissenschaftliche Datenverarbeitung mbH Göttingen (GWDG) which additionally operates the local Scientific Compute Cluster (SCC), enabling researchers at both the University and the Max Planck Society – renowned for its cutting-edge research - to run parallel data and compute intensive applications.

The focus of the NHR site at the University of Göttingen and GWDG is to expand its competencies in the areas of life sciences, earth system sciences, fluid mechanics, artificial intelligence (AI), big data and digital humanities.

## Building a new home for Emmy

**Figure 1**

*Delivered and installed in less than two months, a series of STULZ MicroDCs within a STULZ Modular Data Center are the prefabricated home for the Emmy Supercomputer cluster at the University of Göttingen. The facilities utilise a hybrid liquid and air-cooling system to ensure maximum cooling efficiency and effectiveness for both high and low density equipment racks*

To facilitate the ongoing development of the supercomputer, the University of Göttingen and the GWDG needed a new data center to house Emmy, as existing onsite facilities could not provide the required space and cooling infrastructure without significant reconstruction and retrofitting work.

To offer ongoing flexibility, the new data center needed to be a flexible modular construction that could accommodate further expansion as needed, with the deployment of a cooling system that could remove a heat density up to 100kW per rack and 1.5MW total capacity. At the same time, it had to provide maximum levels of resilience, continuity and availability, and ensure minimal downtime.

Since data center power consumption was also a concern, the implemented solution needed to be as energy efficient and sustainable as possible. The whole project had to be designed and delivered in a very short timeframe – less than two months were given to install the complete four-room modular data center.

## Cooling design considerations; air cooling lacking scalability for Emmy

Data center IT cooling loads change dynamically throughout the day as system resources are consumed. This means a cooling design must meet the challenges of ensuring sufficient cooling capacity to allow the heat load to flex on demand and offer scalable capacity to allow for future expansion while minimising first costs. The ability to satisfy short-term and long-term cooling requirements will have a major impact on future IT deployments, overall energy efficiency and potentially, the lifecycle of a data center.

Today, air is commonly used as the cooling medium for data center IT loads. It is a well understood approach which is economic in day-to-day application and likely to retain its popularity for the near future. However, there are weaknesses in air cooled designs when the heat flux of the central processing units (CPUs), graphics processing units (GPUs), dual in-line memory modules (DIMMs) and voltage regulators exceeds specified values.

Above certain watt densities, the physical limitations of air preclude its use as a cooling source as, if watt densities become too great, air cannot remove heat fast enough causing servers to throttle their performance or shutdown due to thermal overload.

Traditional air-cooled environments also present barriers for data centers that need to increase rack density while simultaneously lowering Power Usage Effectiveness (PUE) ratings; engineers must consider thermodynamics, the most efficient airflow and other operational factors. However, they're likely to only achieve a certain rack density before airflow dynamics are exceeded without the need for floorspace to be sacrificed or changes to the critical design. Put simply, air cooled solutions may struggle to cool highly dense server pods and clusters running intensive technologies.

Today, data center cooling designs are beset by different vectors converging to create a perfect storm for traditional air-cooled approaches. These include:

- Increasingly power-hungry CPUs and GPUs being designed by manufacturers to meet the data processing requirements of AI-driven software and emerging applications
- Increasing server rack densities as greater space utilisation is sought
- Extreme ambient temperatures and other effects of climate change
- Pressure on the sector to increase power efficiency and reduce the use of resources including water

## Efficiency and effectiveness demand liquid cooled IT at Gottingen

Using liquid as a cooling medium instead of air enables higher thermal transfer capacity and improved efficiency. The technology promises to reduce the reliance on some of the traditional energy hungry components of the cooling ecosystem including chillers, computer room air handlers (CRAH) and computer room air conditioning (CRACs) units. As liquid cooled technology continues its journey within high performance computing, it enables the mainframe level systems that reside in research laboratories and dedicated data centers to benefit from flexible platforms that have highly efficient, high-volume processing in standard rack or custom format configuration.

ASHRAE has issued a new H1 class of high-density IT equipment in the latest update to its data center thermal guidelines. Class H1 equipment requires inlet temperatures ≤ 22°C (71.6°F), well under the recommended upper limit of 27°C (80.6°F) for general IT equipment. Many modern data centers operate at temperatures above 22°C (71.6°F) to lower cooling energy and water consumption by minimising the use of compressors and evaporation. As the demands on cooling systems grow with every generation of technology, data center operators will find the advantages of liquid cooling harder to ignore.

No surprise that according to the Uptime Institute Cooling Systems Survey 2023, with rack densities rising and the advent of power-hungry chips, "direct liquid cooling of IT hardware is poised to displace air cooling as the dominant approach" as data center operators "grapple with the challenge of finding the right balance in cooling performance, energy efficiency and sustainability."

## Clear benefits of liquid cooling

Liquid has two properties that elevate its heat transfer capabilities above those of air: The thermodynamic property of specific heat capacity at constant pressure; and its increased density per unit of volume. Presented with density challenges, air cooling alone won't be able to handle the demands put on it by powerful processors, while pressure to reduce energy consumption will continue to increase.

In terms of specific heat capacity, in pure water this is over four times greater than that of air at normal temperatures. On a per volume basis, the mass of water is over 830 times that of air at standard conditions and the rate of heat transfer is equal to the product of the mass flow rate of the cooling medium. This is why water has historically been the go-to cooling medium for removing large amounts of heat energy.

In the case of liquid cooled servers there is another key advantage over the use of air. The passive cold plates on the CPUs and memory DIMMs have high surface temperatures and this, coupled with the heat carrying capabilities of the liquid coolant enables much higher temperatures. It also enables economisation (compressor-less cooling) for many more hours, leading to greater energy saving without any decrease in compute operations.

## Direct, liquid-to-chip cooling, the way forward at the University of Göttingen

Direct liquid-to-chip cooling uses the exceptional thermal conductivity of a dielectric liquid to provide dense, concentrated and inexpensive cooling. As direct liquid-to-chip cooling systems can maintain chip operating temperatures far below that of air-cooled equivalents, they can also dramatically reduce static power. Fujitsu estimates, for example, that liquid cooling lowers processor power requirements by over 10 per cent when cooled from 85°C to 30°C.

By reducing dependency on fans and air handling equipment, extremely high rack densities are made possible, while the power consumed by the cooling system drops significantly. This results in more power availability for computing, as each server in each rack can be liquid cooled, as well as a significantly reduced carbon footprint and improved PUE. In fact, direct liquid cooled system implementations can help enable pPUE as low as 1.02, easily outperforming the most efficient air-cooled systems.

**Figure 2**

*A liquid-cooled passive "cold plate" is attached directly to the microprocessor – the computer component that does most of the important work, uses the most power and generates the most heat. A dielectric coolant is utilised so that in the unlikely event of a leak, no damage is done to sensitive electronic or electrical devices in the white space.*

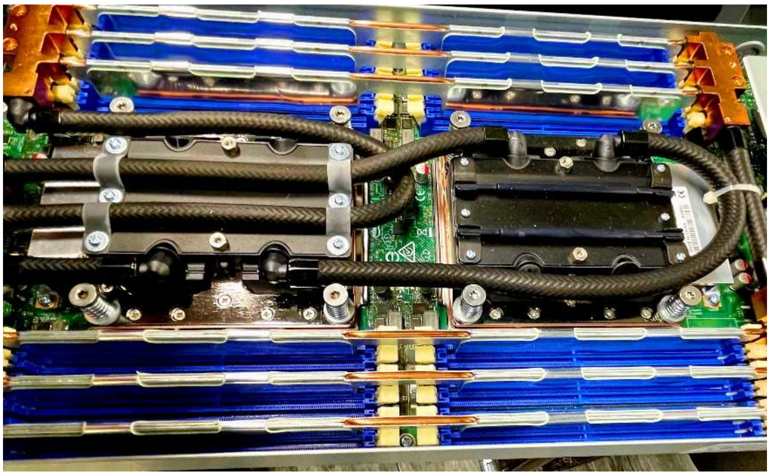

# University of Göttingen's liquid cooled data center

The University of Göttingen approached STULZ to configure a modular data center solution for a 70m$^2$ air cooled room together with an 85m$^2$ direct liquid-to-chip cooled room, with total capacity of approximately 1.4MW. In terms of energy efficiency, a desired PUE 1.14 was stipulated, with 75 per cent of peak load at 32°C ambient and 100 per cent at up to 19°C ambient.

The data center comprises high performance computers, 1,120kW direct liquid cooled systems with approximately 20 per cent residual heat, closed racks, a precision air conditioning unit with free cooling, row coolers, 96kW per full rack and 11 racks, with capacity for up to 14 racks.

The racks are configured with 96kW of IT equipment heat load. The direct-to-chip liquid cooled system removes 74.9kW (78 percent) of the server heat load. The remaining 21.1kW (22 percent) of heat produced by other electronic components of the servers, such as capacitors and resistors are air cooled using STULZ CyberRow cooling units. Final heat rejection to atmosphere for both the air cooling and the direct-to-chip liquid cooling systems is delivered by a shared fluid cooler and pump loops.

The system also comprises a cooling distribution unit (CDU), pumps, and a dry cooler together with the overall hydraulics needed to operate the system. The dry cooler surface area is very large and therefore low fan speed is needed to achieve the correct capacity, and speed-controlled pumps are used to ensure efficiency under all conditions. It has been designed in such a way to ensure water flow and corresponding water temperatures have the most efficient operation.

**Figure 3**

*HPC racks at Gottingen might lack the drama of flashing lights but illustrate how liquid cooling eliminates the need for perforations or grilles to facilitate airflow across the server load. Not so easily photographed is the lack of noise from fans and other air handling equipment!*

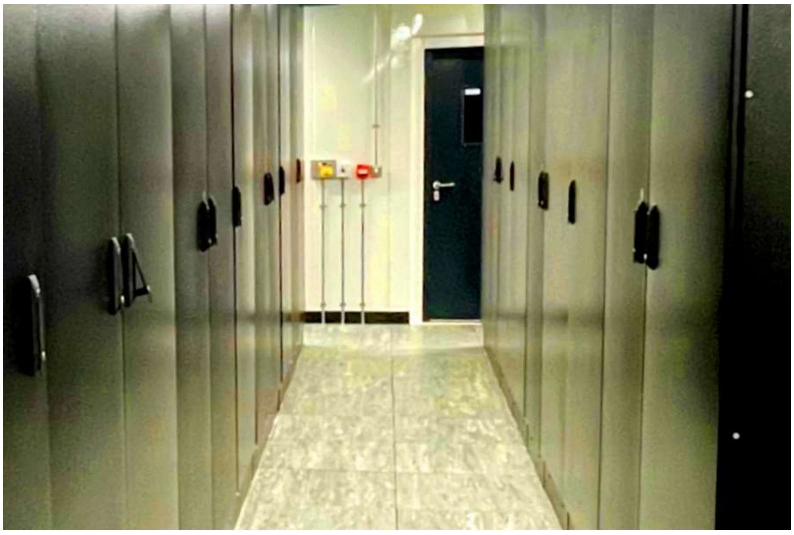

## Cooling – the numbers

Due to the specified layout, the precision air conditioning unit return air temperature is specified at approximately 48°C, supply air temperature at 27 - 35°C and water temperature at 32 - 36°C. As the return air temperature is quite high and exceeds the limits of some electrical components of the row coolers, special ductwork was developed to avoid the possibility of electrical component overheating.

The air conditioning unit is controlled according to the supply air temperature. However, when we consider the water loop, the returning cold water from the dry cooler is first pumped through the CyberRow cooling unit to ensure the lowest possible water temperature. This process step minimises compressor usage to maximising the free-cooling capacity and achieve the highest possible cooling system efficiency.

The exhaust (return) water temperature from the chip can rise to approximately 55°C. This secondary cooling circuit heat is then rejected at the CDU through a plate heat exchanger to the primary cooling circuit. The cold-water supply from the dry-cooler (primary cooling circuit) first flows through the STULZ CyberRow, row coolers to gain maximum cooling capacity, and second flows to the CDU heat exchanger to capture the heat from the liquid-to-chip cooling circuit (secondary circuit). This process helps maximise both the cooling capacity of the CyberRow units, and the efficiency of the entire system.

## Leveraging free cooling to maximise efficiency

A key objective of the design was to reduce compressor runtime, which can be the source of between 80 - 90 percent of overall cooling energy consumption. The cooling design therefore makes use of an indirect free cooling system to eliminate the most energy-consuming parts and maximise the efficiency of the system.

An indirect free cooling system consists of a closed refrigerant circuit, with the plate condenser connected to the dry cooler. To increase the free cooling capacity, a special coil was calculated and developed. The installed system is able to operate with no mechanical cooling for long periods of the year. When compared to a standard air-cooled system, energy savings of up to 70 per cent are possible at the same time ensuring dynamic operation of the primary fluid circuit for maximum efficiency.

The system utilises outdoor air for indirect free cooling in cooler months, switching to energy saving mode as soon as the outside temperature permits. The indoor unit has two cooling components – a direct expansion (DX) cooling coil and a free cooling coil. In warmer months, when the external ambient temperature is above the set point, the system operates as a water-cooled DX system and the refrigeration compressor rejects heat into the water.

In free cooling mode, dry cooler fans are allowed to run and cool the water to approximately 5°C above ambient temperature before being pumped through the free cooling coil. Dependent on water temperature and/ or heat load demands, the water can be used in mixed mode, where it is directed through both proportionally controlled valves, thus enabling proportional free cooling and water-cooled DX cooling to work together. DX mode will start at an ambient temperature below return air temperature.

In order to prevent the dry cooler from freezing when the outdoor ambient temperature is below zero, 35 per cent ethylene glycol is added to the primary water loop as antifreeze. This is the lowest concentration of refrigerant which can be added to ensure the effectiveness of the antifreeze agent, and also helps to reduce the carbon footprint of the cooling solution in comparison with alternative approaches.

In summary, the cooling design at University of Göttingen is optimised to maximise the number of hours of free cooling operations throughout the year. When compared with the power consumed by a typical DX (direct expansion) air-cooled system, the free cooling chiller with chilled water solution delivers approximately 40 percent savings, while the glycol economizer delivers up to 55% power savings.

## Detailing the liquid cooling circuit

**Figure 4**

*Cooling distribution units (CDU) are installed in the rear of equipment cabinets to ensure that waste heat is conducted away from the processor boards, and a cooled dielectric is returned for reliable and ongoing HPC operations*

STULZ worked with CoolIT to incorporate Direct Liquid Cooling (DLC) to the supercomputer microprocessors. The direct-to-chip liquid cooled system is composed of two liquid loops. The secondary loop provides flow of cooling fluid from the CDU to the distribution manifolds and into the servers, where heat is transferred through the cold plates into the coolant.

The secondary fluid then flows into the heat exchanger in the CDU, where it transfers heat into the primary loop and the absorbed heat energy is carried to a dry cooler and then rejected. The water flow direction in the dry cooler outlet is configured to ensure highest capacity and lowest power consumption with the water temperature set to ensure maximum water Delta T.

In terms of maintenance, cooling loops can be disengaged from a piece of equipment without affecting nearby equipment, so technicians can easily service as needed without having to power down the whole rack. Importantly, there's no risk of water spillage and potential damage, as the fluid is safe, non-flammable and dielectric.

CoolIT Quick Disconnects (QDs) are fitted for ease of servicing. QD flush face metal connectors enable fast and simple service of hot-swappable servers while providing the structural integrity data centers require. They are also higher in quality and reliability than plastic counterparts, an important consideration when installing liquid cooling in high-density clusters designed for HPC loads.

The reliability and no-touch features of liquid cooling solutions match the needs for extended mean time to maintenance, and the longer intervention intervals needed for viable operation and management of the in-situ equipment.

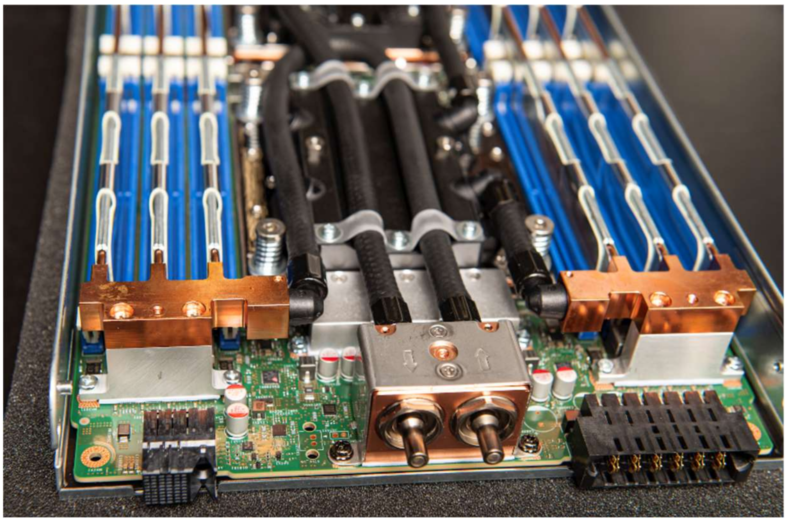

**Figure 5**

*Picture of the processor board showing supply and return connections for the direct-to-chip dielectric coolant, illustrating ease of servicing IT equipment moves, adds and changes.*

# Results

1. 1.5MW of supercomputing has been contained in a series of STULZ micro data centers, optimized for 100kW per rack. The 100kW Micro DCs are housed in 85$m^2$ of modular data center space, and utilise direct liquid-to-chip cooling for precise heat removal and reliable operations of the CPUs, together with air-cooling for the rest of the server components (lowering the risk of thermal shutdown, and providing the ability to overclock processors for maximum supercomputing performance).

2. A secondary, 70m2 air cooled STULZ modular data center was designed and built to house network and storage equipment, located adjacent to the high-density environment (see rendered images in Appendix 4).

3. In comparison with a standard air-cooled data center with a PUE of 1.56 (current industry average according to the Uptime institute), the STULZ modular data center provides 27% savings in electricity at an average 75% (1050kw) load, equating to 3.96 GWh per year.

4. The use of a combination of STULZ Micro DC and prefabricated modular data center technology ensures:
    a. High density 100kW/ rack heat has been captured 78% through the liquid and 22% through air in an effective way.
    b. The benefits of prefabrication throughout this application includes increased speed of deployment; the requirement for delivery and installation in a two-month window was met by STULZ.
    c. By completing much of the engineering work in a factory setting, onsite work was carried out far more quickly and accurately than most traditional data center build or upgrade projects.

5. The quantity of refrigerant used has been greatly reduced by up to 80 percent, further reducing the CO2 footprint of Emmy.

6. The requirement to meet annual operating PUE <1.14 has been met, helping the University meet its sustainability targets for energy savings.

# Conclusion

The University of Göttingen is dedicated to reducing the CO2 footprint its energy consumption, including that of its high-performance computing-based operations. Not only does the STULZ modular data center operate reliably with even better energy efficiency than requested, the efficiency of the hybrid liquid and air-cooling system has provided an overall annual facility PUE of 1.13, and 1.07 PUE for the liquid cooled supercomputer room alone.

Using expertise from STULZ means that the University of Göttingen now has a compact, flexible, incredibly efficient and reliable data center with high performance density in a small footprint. STULZ and the University are currently working together to further improve the efficiency of Emmy’s hybrid cooling system. This white paper will be updated when the results are ready.

# APPENDIX 1 – ABOUT AMELIE EMMY NOETHER

Emmy is named after the pioneering German mathematician Amalie Emmy Noether, who was born in Erlangen in 1882 and died in 1935. She blazed a trail for women in the field of mathematics and during her career achieved several landmarks.

In 1915 she joined the Mathematical Institute in Göttingen – despite the objections of some faculty members to a woman teaching at the university. Whilst there, Emmy started working with Felix Klein and David Hilbert on Albert Einstein's general relativity theory. In the process, she proved a revolutionary mathematical theorem that has changed the way physicists study the universe and kickstarted an entire discipline of mathematics called abstract algebra.

In a letter to Hilbert, Einstein himself praised Emmy's work as a piece of 'penetrating mathematical thinking'.

# APPENDIX 2 – EMMY SUPERCOMPUTER IN NUMBERS

## The need for speed

Emmy is ranked in the latest Top 500 list of the world's fastest computers worldwide. In Germany, at her inauguration, Emmy was ranked sixth as the most powerful computer in the north of the country, serving regions including Berlin, Brandenburg, Bremen, Hamburg, Mecklenburg-Western Pomerania, Lower Saxony and Schleswig-Holstein.

Emmy can perform several trillion arithmetic operations per second and has a theoretical peak performance of 8.78PFlop/s. It has reached 5.95PFlop/s as a maximum value so far – that is 5.95 billion million. Its compute nodes are:

## Phase 1

- 432 nodes with 192Gb memory (medium node)
- 16 nodes with 768Gb memory (large node)
- Per node 2 CPUs and 1 Intel Omni-Path host fabric adaptor
- Per CPU 1 Intel Skylake Gold 6148 20 cores
- 1 480Gb SSD hard disk

## Phase 2

- 1004 nodes with 384Gb memory (standard node)
- 16 nodes with 768Gb memory (large node)
- 2 nodes with 1.5Tb memory (huge node)
- Per node 2 CPUs and 1 Intel Omni-Path 100G host fabric adaptor
- Per Intel Cascade Lake Platinum 9242 (CLX-AP) CPU  48 cores

## 3 GPU nodes

- 2 Intel Skylake Gold 6148 CPUs (40 cores per nodes)
- 768Gb memory
- 1 480Gb SSD hard disk
- 1 Intel Omni-Path host fabric adaptor
- 4 NVIDIA Tesla V100 32Gb

For more details :visit https://www.top500.org/system/179883/ - Emmy's position in the Top500 was 47 at inauguration, today she is #133

# APPENDIX 3 – LIQUID COOLING BENEFIT CHECK

As servers consume more electrical power, they also generate more heat – an effect that has increased the need for innovative approaches that go beyond the limitations of conventional air cooling. The use of direct liquid to chip cooling technology dramatically increases data center density and enables maximum sustained CPU throughput. Here are the key benefits:

1. Liquid cooling lowers energy use
2. Liquid is a highly efficient way of transferring heat away from electrical components, reducing the need for high energy consuming fans. Lower energy bills, combined with low maintenance and high reliability, lead to lower ongoing facility costs.
3. The technology provides a fast return on investment (ROI)
4. Direct liquid-to-chip cooling uses liquid coolant to provide adequate cooling for servers. Most systems have an ROI of around a year and offer substantial ongoing savings when compared to air cooling systems.
5. Space saving
6. Liquid cooling improves cooling performance
7. Precise cooling is needed when the latest chips are operating at their highest performance. Direct liquid-to-chip cooling eliminates hotspots and thermal shutdowns, keeping things running smoothly.
8. Simple servicing - server level components such as CPUs and DIMMs are easily accessible and quick to service
9. Enables processing capacity to be maximised
10. Direct liquid-to-chip cooling is precisely targeted and highly effective in isolating and dissipating heat. This allows greater CPU and GPU densities.
11. Reduces downtime
12. Direct liquid-to-chip cooling means less likelihood of equipment overheating, reducing the risk of thermal shutdown and helping to avoiding downtime.
13. Facilitates heat reuse - outgoing heated liquid from servers can be utilised by a secondary system for greater circularity and reduction of waste
14. Liquid cooling can be integrated with air-cooling - if a completely liquid cooled data center isn't needed, existing air-cooling infrastructures can be integrated to keep capital expenditure down, while maximising rack density.


sales@stulz-modular.com | www.stulz-modular.com

# APPENDIX 4 – RENDERED IMAGES OF EMMY'S NEW HOME

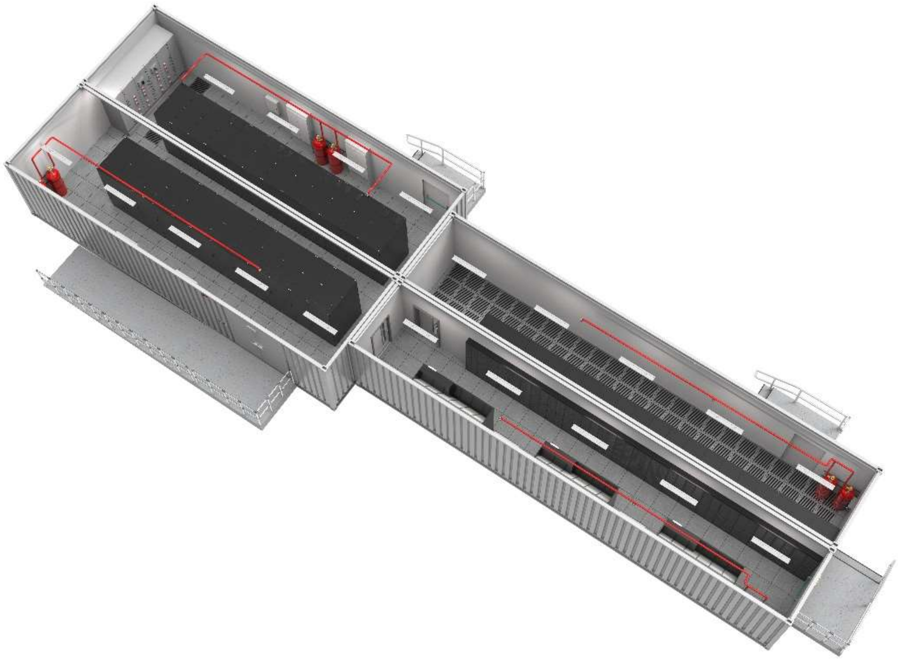

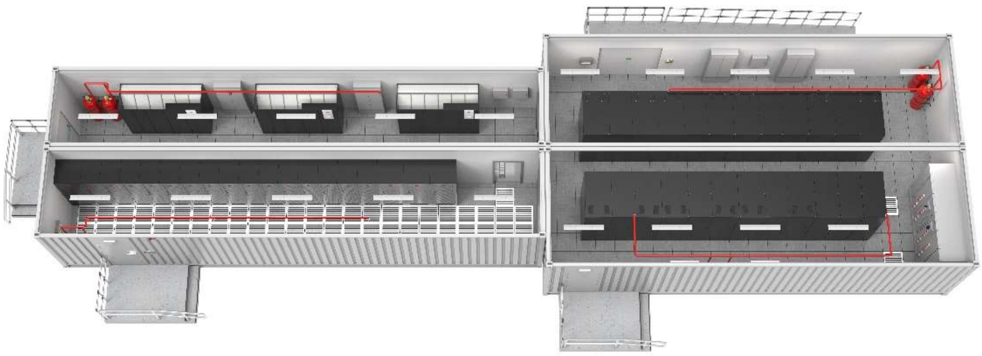

# ABOUT THE AUTHORS

## Dr. Christian Boehme

**Deputy Head HPC Group, GWDG**

Christian Boehme has been with GWDG since 2003 and introduced and operated the resource management for the institution's first Linux based HPC cluster. He also coordinated planning and operation of several HPC systems, including the NHR system "Emmy", as well as University of Göttingen's Modular Data Center (MDC). He has coordinated national research projects on HPC-as-a-Service and Performance Monitoring.

## Dr. Sebastian Krey

**HPC System Architect, GWDG**

Sebastian Krey studied statistics at the Technical University of Dortmund, graduating in 2008. He successfully completed his PhD one year ago. As a member of the GWDG scientific staff he works in a number of positions, including working as an HPC system architect planning and developing the HPC infrastructure at GWDG. His practical work is focused on operating the HPC storage systems and high-performance networks. He also invests his time in research and is currently working on energy efficient HPC system operations as well as performance optimization of large scale HPC storage systems.

## Dave Meadows

**Director of Technology, STULZ-ATS**

A respected engineer, Dave Meadows is the holder of a BSc in Mechanical Engineering as well as a graduate of the United States Navy's Nuclear Power School. He has over 24 years of data center HVAC design experience and is a US patent holder as co-inventor of a dual mass precision cooling device. He is an active member of ASHRAE TC 9.9, ASHRAE SSPC 90.4, ASHRAE SPC 37, ASHRAE SSPC127, the Green Grid, and the AHRI data com standards committee.

## Norbert Wenk

**Head of Product Management, STULZ GmbH**

With 25 years of experience in the data center cooling industry, Norbert Wenk is head of Product Management at STULZ. Following a degree in Mechanical Engineering, he joined the company in 1999, starting his career in the R&D department. As a technical thought leader, he is the author of many engineering articles as well as being an enthusiastic contributor to white papers.

# ACKNOWLEDGEMENTS:

Major contributors to this white paper include its authors:

Dr. Christian Boehme - Deputy Head HPC Group, GWDG
Dr. Sebastian Krey – HPC System Architect, GWDG
Dave Meadows - Director of Technology, STULZ-ATS
Norbert Wenk – Head of Product Management, STULZ GmbH

This white paper was reviewed by:
Dushyantha Goonawardhane – Managing Director, STULZ Modular Ltd



Direct-to-chip liquid cooling - more details regarding the Direct Liquid Cooling system are provided in a white paper published by CoolIT Systems and available for download from their website https://www.coolitsystems.com/