NEXALUS
COOLING THE CLOUD

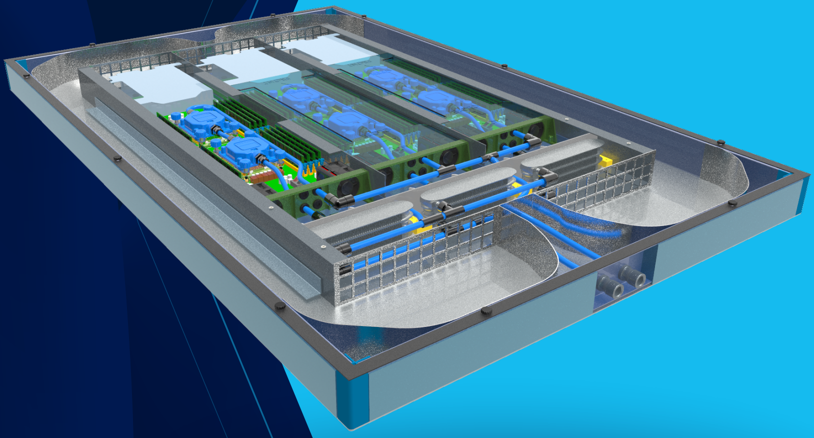

OCP Server Report

Performance Evaluation of OCP Leopard Server

Executive Summary

AN INTRODUCTION...

The following document steps through the re-engineering of OCP Leopard Servers with Nexalus sealed-server thermal management technology. The report outlines key results from tests performed at the Research Institutes of Sweden (RISE), which have independently verified that Combined Heat & Compute (CHC) is achievable with Nexalus' re-imagining of how data centres should manage heat energy from servers. This is the hard evidence that supports Nexalus' claim that Performance, Profit and Planet can all be prioritised without having to compromise. Specifically, compared with the way things are done now, the Nexalus sealed-server technology offers better cooling of server hardware whilst creating an asset that can displace grid-based heat energy consumption in neighbouring domestic, commercial and industry sectors.

Starting with the digital model of the OCP Leopard Server system, we developed this solution in accordance with the Nexalus 7-Step Design, Prototype & Consultation process. The foundation of this process is academic quality Multiphysics Simulation-Driven Design, which progresses to equally rigorous prototype build and testing stages before being handed over for independent validation.

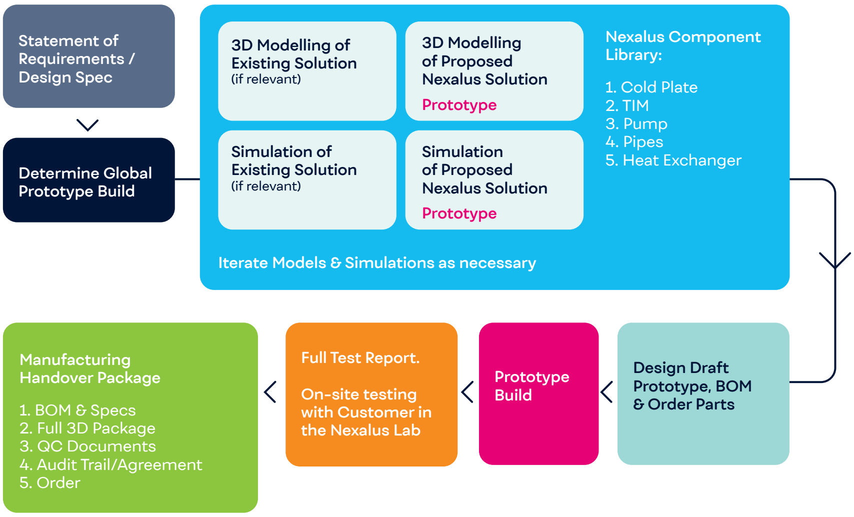

The main body of the report begins with a description of the way things are now, outlining key RISE measurements on the stock OCP server system to provide a sense of its thermal performance and behaviour. This is followed by describing the Nexalus approach, with our conclusions being supported by the independent measurements from RISE. We then wrap-up with a synthesis that attempts to answer the question ***'why should data centres do things differently?'***.

ACKNOWLEDGEMENTS

Nexalus would like to acknowledge the interest and engagement of Dr. Jon Summers and Mr. Jeffrey Sarkinen of RISE who carried out the independent testing. Thank you for rising to the challenge and understanding the standard of scientific integrity demanded by Nexalus.

1. / Performance Evaluation of OCP Leopard Server: Pre and Post Nexalus-Retrofit

OVERVIEW

A Multiphysics Simulation-Driven Design approach has been undertaken to augment the OCP Leopard Server thermal management hardware with the Nexalus liquid-cooled sealed server technology. Independent testing at the RISE Research Institutes of Sweden (www.ri.se) has proven that an average of 93% heat recovery is achievable at water temperatures up to and exceeding 65°C. The Nexalus solution can maintain this high water temperature over a range of CPU workloads, ranging from 8% to 75%, which captures the typical range in operation.

Importantly, the Nexalus solution achieves this within an architecture that is 1OU in height, half that of the stock 2OU server, potentially doubling the compute density of a rack.

1.1 / Thermal Characterisation of Stock OCP Leopard Server

1.1: THERMAL CHARACTERISATION OF STOCK OCP LEOPARD SERVER

Figure 1. shows photographs of the as-delivered OCP Leopard Server. As shown, each of the three servers slots into a channel in the Cubby. The stock OCP Leopard Server is an air-cooled system. Air is drawn into each individual channel-type server enclosed in the ORV2 Cubby through the rear panel by two 80 mm fans, as depicted in Figure 1. (a&b). The air drawn into the front of the server is forced over the internal electronic hardware, with critical cooling of the two CPUs facilitated by fixing heat pipe augmented finned heat sinks that extend close to the top of the 2OU case (Figure 1.(c)).

Figure 1.
As-delivered stock OCP Leopard Servers in Open Rack V2 Cubby
(a) front view, (b) back view, and (c) Internals of stock OCP Leopard Servers.

1.1 / Thermal Characterisation of Stock OCP Leopard Server

Experiments were performed at the RISE test facility in Lulea Sweden by the ICE Data Centre research group led by Jon Summers (https://www.ri.se/en/person/jon-summers). The wind tunnel facility provides a testbed whereby a matched airflow to the stock server fans was supplied at a controlled air temperature of 30°C. The test method involved fixing the server workload and decreasing the airflow to increase the CPU average operating temperature. Server workload was varied between 8%, 50% and 75% to mimic the typical operating swings experienced in application.

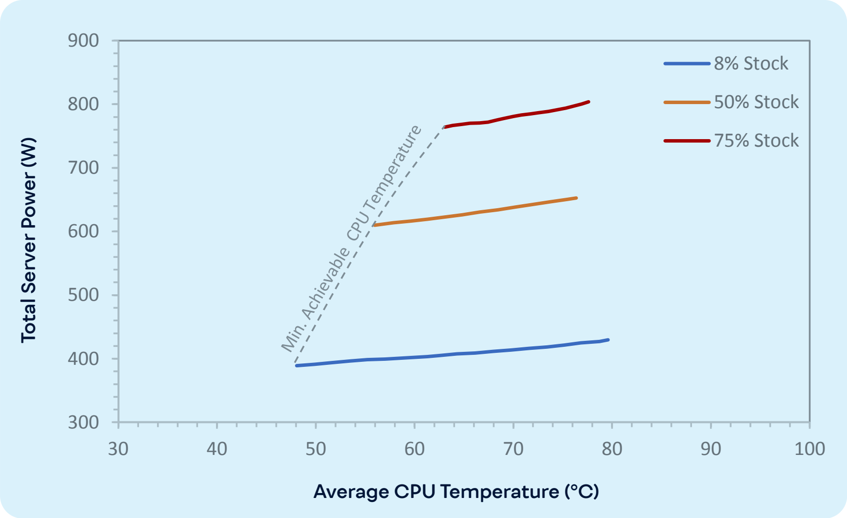

Figure 2.
Total measured power to three Leopard servers versus average CPU power for stock air-cooled system for workloads of 8%, 50% and 75% and air inlet temperature of 30°C.

Figure 2. plots the key thermal performance results for the three air-cooled stock servers in the Cubby. As shown, as the workload increases from 8% to 75%, the total power to the Cubby increases from an average of 410 W to 780 W, mainly due to higher CPU power draw. For each workload, the total power increases by about 40 W with increasing CPU temperature. Since silicon-based electronics tend to generate more leakage current at higher temperatures, some small efficiency losses are typical.

Also marked in the Figure 2. plot is the minimum achievable CPU temperature for the stock air-cooled system with 30°C inlet air. This is achieved with the highest fan speed and thus air flow rate through the servers, which provides the maximum level of heat transfer. As shown, the minimum achievable average CPU temperature increases from 48°C to 63°C between the lowest and highest workloads tested.

1.1 / Thermal Characterisation of Stock OCP Leopard Server

In the context of heat recovery, one important parameter is the exit temperature of the coolant. The example results presented in Figure 3. are for a fixed average CPU temperature of 78°C and varying workload. As the figure shows, only at the lowest workload of 8% i.e. the lowest power level, does the air temperature reach a level approaching 60°C, which can be considered reasonable from a heat recovery perspective. Regardless, transport of this heat energy over any reasonable distance would require an air-to-water heat recovery heat exchanger to transfer the heat into a water stream, since the energy density of water is thousands of times higher than air.

Air-to-water heat exchangers have typical efficiencies of 40%-85%, meaning that, at best, about 15% of the total heat energy is not recoverable. It also means that the heat in the water stream will necessarily be, at best, 10-15°C cooler than the supply air temperature. This is a hard penalty since it reduces the quality and usefulness of the recovered energy. The knock-on effect is that additional mechanical infrastructure, in the form of heat pumps, need to be installed for heat recovery from air-cooled data centres. Although heat pumps can boost the water stream temperature to usable levels, they add to infrastructure complexity & cost and consume significant electrical power.

In fact, adding heat pumps is a step backward on the global movement towards eliminating compressors from data centre infrastructure, a direction motivated by the need to reduce energy consumption.

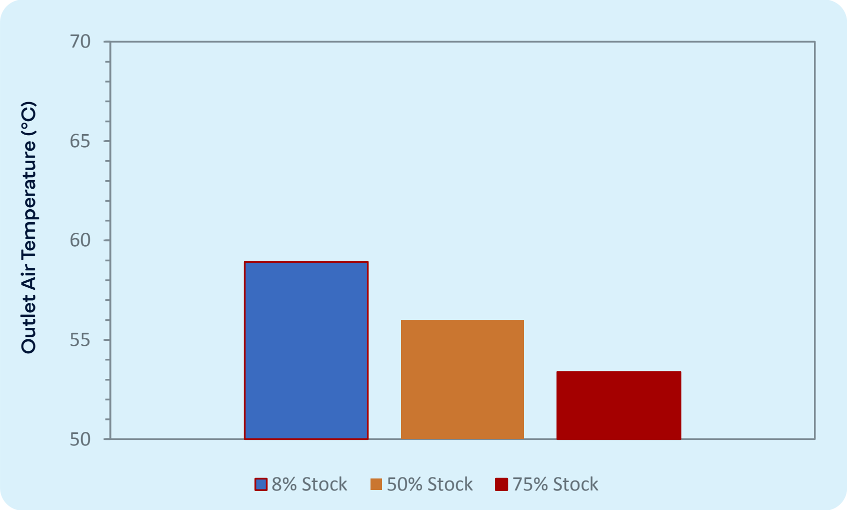

Figure 3.
Experimental outlet air temperature for stock air-cooled OCP Leopard Servers for workloads of 8%, 50% and 75%, average CPU temperature of 78°C, and air inlet temperature of 30°C.

2. / Thermal Characterisation of Nexalus-Retrofit OCP Leopard Servers

Figure 4. illustrates the modified stock OCP Leopard Servers within the Nexalus-retrofit sealed-case technology. Critically, this design is 1OU, doubling the number of servers and associated compute density in a rack. The case enclosure is air-tight, meaning that there is no through-flow of air into and out of the servers. In this way, all the required cooling as well as the recovery of thermal energy happens inside the sealed case, decoupling the server heat and energy management from the surroundings.

The electronics cooling as well as the heat recovery are facilitated by water which enters the case through quick-disconnect leak-proof fittings. The water is routed to six Nexalus Enflux direct-attach cold plates, which are mounted to the CPUs with a layer of Nexalus Hydronex thermal interface material and mechanically fastened in place with toolless Nexalus Torque Screws to ensure minimum thermal contact resistance. Direct liquid cooling of the CPUs simultaneously creates improved cooling whilst recovering all the heat generated by the server electronic hardware into the water stream. In these Leopard servers, the CPUs account for about 65% - 75% of the total energy of the servers, with the remaining being generated by supporting hardware.

Cooling and heat recovery of the remaining less critical components is achieved by the circulation of the enclosed air within the sealed server case. The trapped air is forced through the severs by six 40 mm fans, two per server, located just downstream of the DIMMs. The air temperature is managed by passing the heated air over three internal Nexalus NXQ air-liquid heat exchangers which act to absorb the remaining heat energy into the water, while at the same time cooling the air for repass over the air-cooled electronics.

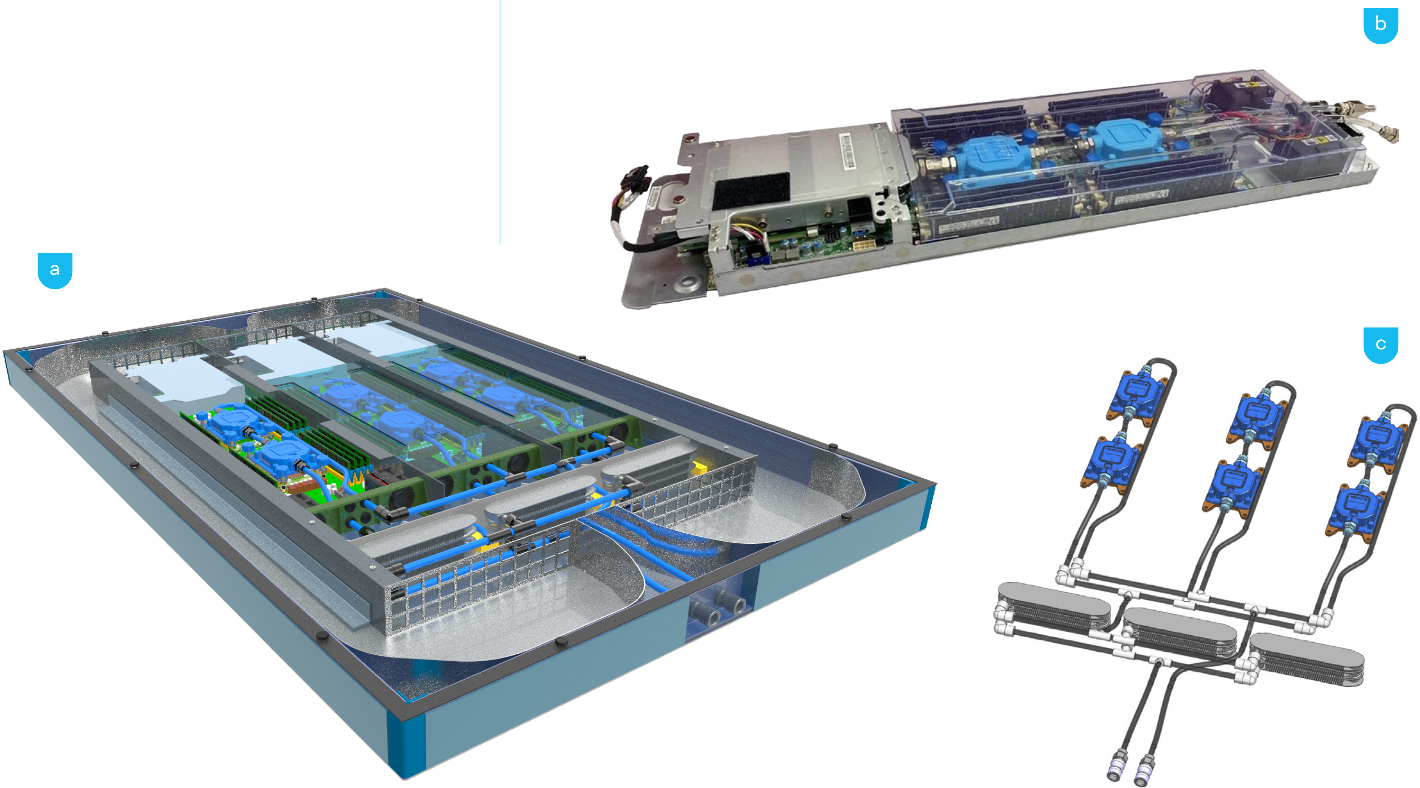

Figure 4.
(a) Nexalus-retrofit 3x Leopard Servers, (b) Internals of Nexalus-OCP Leopard Servers, and (c) Nexalus cooling hardware.

2. / Thermal Characterisation of Nexalus-Retrofit OCP Leopard Servers

As with all Nexalus technologies, the Combined Heat & Compute OCP Leopard system design was simulation-driven, using state-of-the-art Multiphysics CFD software. This included both thermal and fluid dynamics simulation for the air and water circulation inside the system, along with heat transfer in relevant solids including the heat exchangers and key electronic hardware.

The CFD simulation predictions pictured in Figure 5. illustrates the key engineering design features of this novel flow configuration. Instead of air flowing in-the-front and out-the-back, as is the case for the stock system, the Nexalus design recirculates the air within the sealed enclosure by framing the server sleds in a ducting system with front-to-back flow bypasses on each side. As illustrated in the figure, the air heats upon passing over the exposed electronic components, like the DIMMs, is passed over the embedded Nexalus NXQ air-water heat exchangers where it is cooled, and then routed front-to-back via the outer bypass vanes to re-enter the server sleds.

The water-side simulation shows that for a simulated maximum CPU power of 125 W per CPU, commensurate with their maximum power, the average CPU core temperatures were predicted to be just over 30°C above the coolant temperature for flow rate of just 0.33 Litres per minute per server, illustrating the effectiveness of the highly engineered microjet-based Nexalus Enflux cold plates.

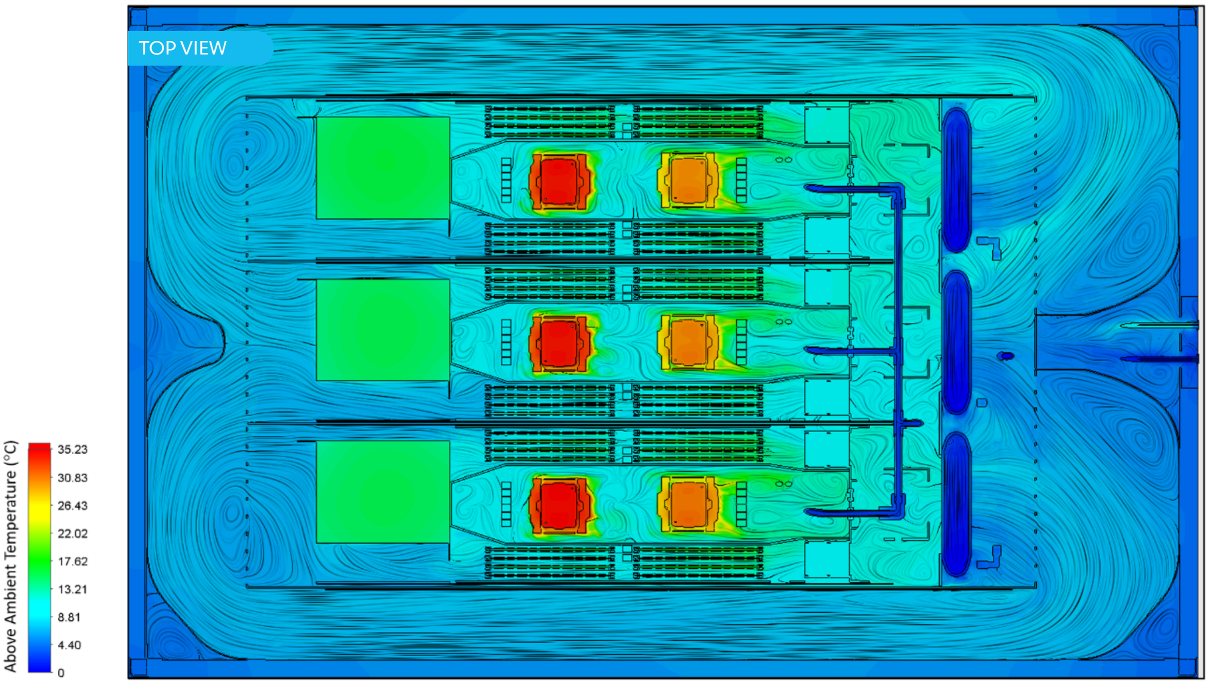

Figure 5.
Multiphysics CFD simulation of Nexalus-retrofit OCP Leopard Servers for 125 W per CPU power showing illustrations of flow distribution via streamlines, and above-ambient temperature difference distribution via false-colour mapping.

2. / Thermal Characterisation of Nexalus-Retrofit OCP Leopard Servers

Experiments were performed at the RISE test facility in Sweden by the ICE Data Centre research group. The RISE water-cooling facility provides a testbed whereby a monitored and controlled water flow was supplied at a variable temperature, here ranging from 30°C to 50°C. With the fans fixed at 50% speed, the test method followed a similar procedure to the air-cooled testing, whereby the server workload was fixed and the water flow rate decreased to increase the CPU average operating temperature. Server workload was varied between 8%, 50 % and 75% to mimic the typical operating swings experienced in application.

Figure 6. plots the key thermal performance results for the 1OU Nexalus sealed-server technology for the 30°C inlet water temperature case. Also plotted are the stock 2OU 30°C inlet air-cooled results for direct comparison[1]. As discussed earlier, as the workload increases from 8% to 75%, the total power to the three server's doubles owing primarily to higher CPU power draw, and this is true for both the stock and Nexalus-retrofit architectures. Unsurprisingly, both scenarios experience the same level of total power increase with increasing CPU temperature since leakage current is intrinsic to the operating temperature of the silicon devices, not the cooling solution.

What is most notable in Figure 6. is the superior cooling effectiveness of the Nexalus solution, as demonstrated by significantly lower minimum achievable CPU temperatures for an identical coolant temperatures Notably, for 75% workload, the Nexalus-retrofit CPUs are 15°C cooler than the stock air-cooled system.

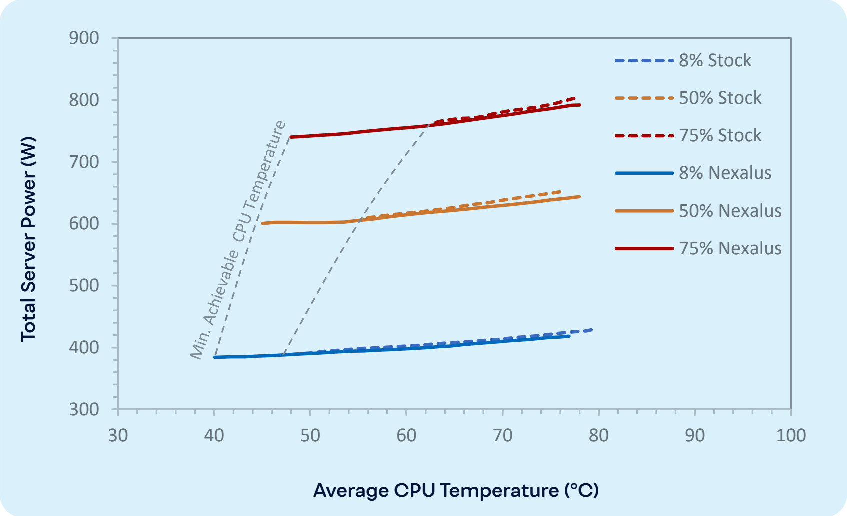

Figure 6.
Total measured power to three Leopard servers versus average CPU power for Nexalus-retrofit OCP Leopard Servers compared to the stock air-cooled system for workloads of 8%, 50% and 75%. Both air and water inlet temperature are 30°C.

---

[1] Albeit the same Leopard server type, system-to-system variability in power (up to circa -/+ 10 W) is possible owing to small inconsistencies in hardware, installation etc. All comparisons should be considered in this context, as well as understanding that all measurements have associated experimental uncertainty.

2. / Thermal Characterisation of Nexalus-Retrofit OCP Leopard Servers

Figure 7. compares the coolant outlet temperatures for the stock air-cooled and Nexalus-retrofit servers for a fixed average CPU temperature of 78°C, despite a significant swing in workload and associated thermal power.

The Nexalus technology achieves temperatures approaching and exceeding 65°C, not only adequate for space and process heating applications, but is stored in water which has an energy density about 4000 times greater than that of air. Remembering that the constant average CPU temperature (78°C in this case) for varying workload is achieved by simply varying the water flow rate, it is then important to note that the Nexalus solution provides a near constant hot water supply temperature regardless of wide swings in workload, with only 2.5°C separating the 8% and 75% workloads, which is important in the context of demand-side energy management.

A key differentiator of the Nexalus sealed server technology, whether for OCP-type servers or any other, is the ability to capture the heat energy not only from the high-powered processors, but also the supporting lower powered and less critical (from a cooling perspective) hardware. This includes though is not limited to DIMMs, VRM and sometimes power supplies. For the Leopard servers tested here, Figure 8. demonstrates this by comparing the percentage heat recovery of the Nexalus sealed server technology with what would be achieved in the conventional liquid-cooled server scenario, where cold plates are used to cool the CPUs and conventional flow-through air cools the remaining hardware. As shown, for the Leopard servers, between 25% to 35% of the recoverable heat is wasted if the Nexalus solution is not deployed, which is a significant lost opportunity. On average, 93% of the heat is recovered in the Nexalus system, with the leftover simply due to experimental heat leakage to the surroundings, which would not be relevant when multiple servers are stacked in a rack.

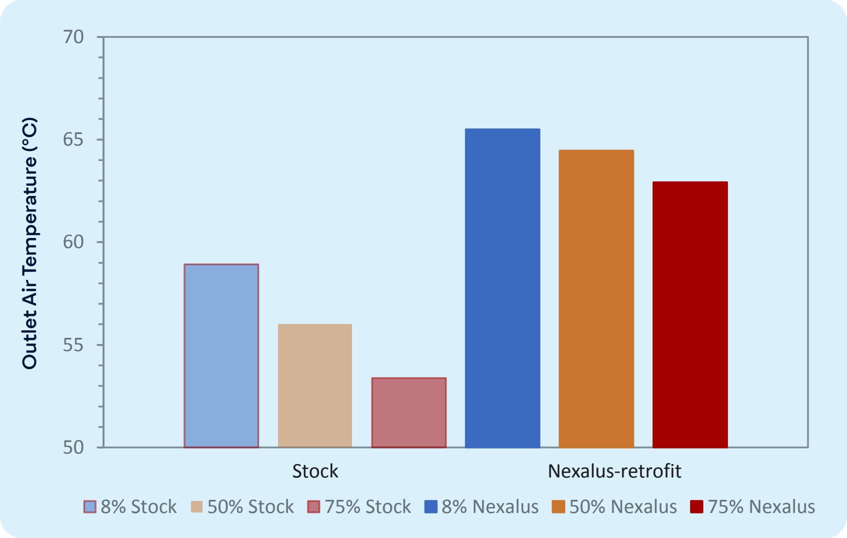

Figure 7.
Experimental outlet air temperature for power for the Nexalus-retrofit OCP Leopard Servers compared to the stock air-cooled system for workloads of 8%, 50% and 75%, average CPU temperature of 78°C, and air/water inlet temperature of 30°C.

2. / Thermal Characterisation of Nexalus-Retrofit OCP Leopard Servers

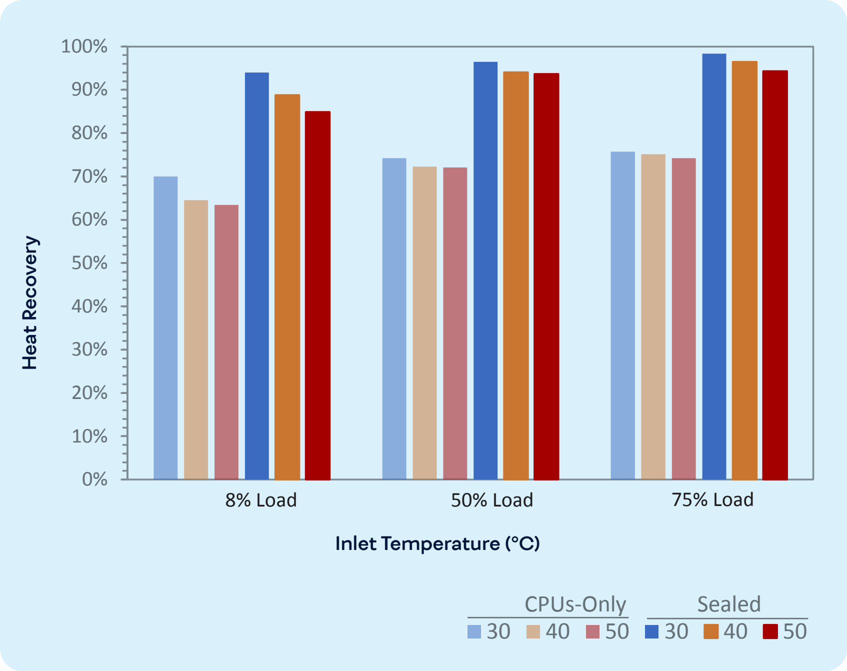

Figure 8.
Experimentally measured heat recovery percentage for the sealed Nexalus-retrofit OCP Leopard Servers compared to scenario of direct liquid cooling of the CPUs-only, for workloads of 8%, 50% and 75% and inlet temperatures of 30°C, 40°C and 50°C.

3. / Benefits & Synthesis

The key benefits of the Nexalus sealed server technology can be summarised as follows:

HELPING TO SOLVE THE ENERGY CRISIS

In today's society there is no doubting it - we need data centres. When we look at data centres, one of the biggest issues in this industry is their energy consumption.

At Nexalus, we are excited to introduce a new energy sequence which will help to solve this energy crisis. Our cooling solutions apply the concept of 'energy borrowing'. Not only does our cutting-edge technology reduce the energy use requirement by 30%, it also recovers over 80% of this energy and converts it into heat, giving energy a second life as heating & hot water.

COMPLETE HEAT RECOVERY IN ONE STEP

It is of course possible to recover heat from air-cooled data centres, but this can not be done efficiently. Practically, when transferring heat from air to water for transport and reuse, there are penalties both in terms of double-digit percentage losses in the amount which can be recovered and in the degradation of the already low quality (temperature) energy. Although heat pumps have been deployed in retrofit-type data centre heat recovery systems to increase the energy quality, they are expensive and power hungry and opposite to the requirement of reducing data centre electricity consumption.

3. / Benefits & Synthesis

Direct water cooling of server CPUs is not new, nor is it by itself innovative. However, when considering water cooling from a heat recovery perspective, one must engineer with a holistic view that encompasses both the heat and energy problem at the same time. Most obviously, when only cooling and capturing heat from the processors, at least one-quarter of the energy in the servers is not recovered. Beyond this, mechanical systems for air handling and conditioning are still required in the data centre infrastructure.

The Nexalus sealed server solution not only captures all the heat generated in the servers, but does so by completely decoupling the thermal management of the servers from the data hall environment. Critically, heat energy has both quantity and quality, and the higher the temperature the higher the quality and associated utility and value. This is the engineering crux, since cool processors and hot water are, on the face of it, conflicting. The Nexalus technology is a new level of liquid cooling performance, engineered to be so effective that safe CPU operating temperatures can in fact be maintained whilst producing hot (~65°C) water.

HIGH COMPUTE DENSITY

The Nexalus sealed server is 1OU compared with the 2OU stock Leopard Cubby. In a rack, this doubles the compute power in the same floor space, before gains in the hot and cold aisle real estate are considered. Furthermore, this solution does not require overhead or underfloor space for hot/ cold air management.

COOLER CPUS, HIGHER RELIABILITY AND LOWER ENERGY COST

For identical coolant temperature and server operating conditions, the Nexalus sealed server liquid-cooled solution can realise much lower CPU operating temperatures, if component reliability and downtime is the priority. Additionally, dust and moisture ingress are eliminated, further improving reliability.

Superior cooling also creates a wider thermal envelope that can be leveraged to either (i) facilitate higher powered processors, or (ii) allow server deployment in hotter climates.

The excess energy required to manage and distribute air to and from the servers in a data centre is, by comparison, significantly higher than water distribution. For water-cooled data centres, PUEs of 1.05 or less are easily reachable.

WHISPER QUIET OPERATION

Acoustic noise levels from server fans and other air handling equipment are notoriously high, affecting health and safety in the workplace.
The Nexalus servers are quiet, owing to acoustic containment in the sealed case. This not only improves data centre health and safety, but opens new doors in terms of deployment and location of computing systems (Office, Edge, 5G etc.).

Scalability

THE RUSSIAN-DOLL APPROACH TO SCALABILITY

We have discussed the technical engineering nuances of the Nexalus technology, its benefits and the impact it can have by disrupting the status-quo of data centre heat and energy management. The question that remains is; ***'is it scalable?'.***

The answer is Yes, and this once again comes down to holistic engineering. Scalability is embedded in the Nexalus technology because it was designed using the Russian-Doll approach of nested-modularity. Whether a few servers, hundreds of servers or thousands of servers, the only thing that fundamentally changes are the boxes that contain them. From a single rack in an office, through containerised computing at the Edge, to brick-and-mortar hyperscale facilities, it is always the same solution to the same problem, just progressively more and bigger boxes.

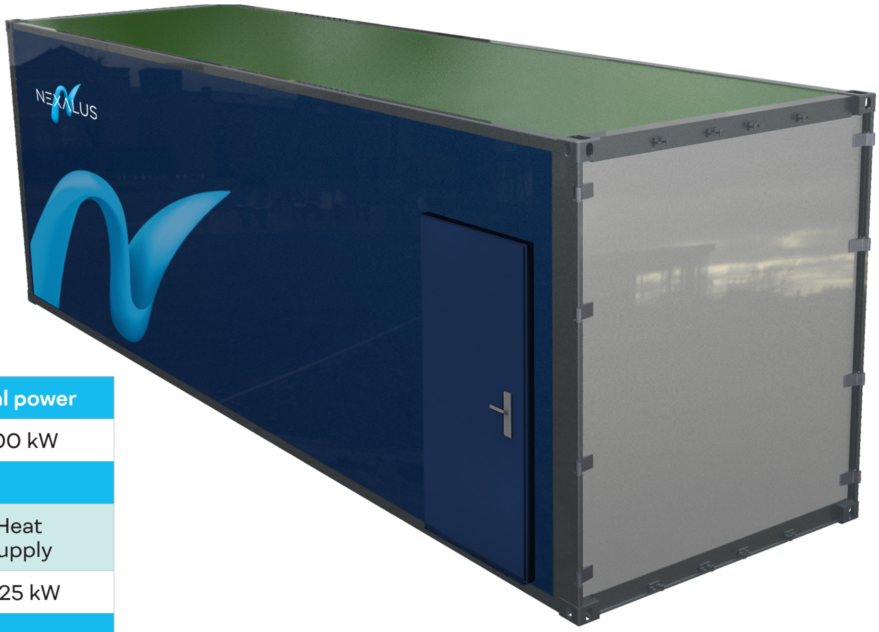

Servers per rack	Number of racks	Total power
108	13	500 kW
Cooling Phase		
Coolant temperature	Coolant flow rate	Heat Supply
≥65°C	25 m^3/hr	≥425 kW
Heat Recovery Phase		
Recovered water temperature	Recovered water flow rate	Power available for reuse (MW)
≥60°C	25 m^3/hr	≥425 kW

To illustrate this, Nexalus have completed the design of a Combined Heat & Compute modular 500 kW containerised data centre. The nested-modularity involves (i) enclosing each server in the Nexalus sealed case, (ii) encasing 108 servers in racks, and (iii) housing thirteen racks within a 40ft container. The transportable container can now be viewed as a Lego-brick module, in the sense that the containers can themselves be stacked on top and beside one another to required scale. The scale can be driven by compute demand, heating demand, or both.

From the brick-and-mortar hyperscale facility perspective, it is of course the same concept, potentially minus the shipping container. Thinking of it this way, a 20 MW facility is simply 40 x 500 kW Lego-bricks arranged in data hall aisles.

Nexalus

ABOUT NEXALUS

Nexalus is a leading thermal design consultancy and solution provider for sustainable computing, with patented technology that prioritises not only performance and profit, but also the planet. With our foundations firmly rooted in both science and engineering, we are not your typical consultants and system designers. At Nexalus, we set out to explore the thermal issues affecting the electronics market, namely ‘heat problems’ and ‘energy problems.

Our team of researchers and engineers have the thermal expertise and capabilities to conceive, design, and build products that focus on solving the thermal engineering and energy problems affecting the electronics industry. Nexalus is here to drive thermal engineering from a place of scientific understanding to commercial deployment of new technologies from Hyperscale to the Edge.

Components

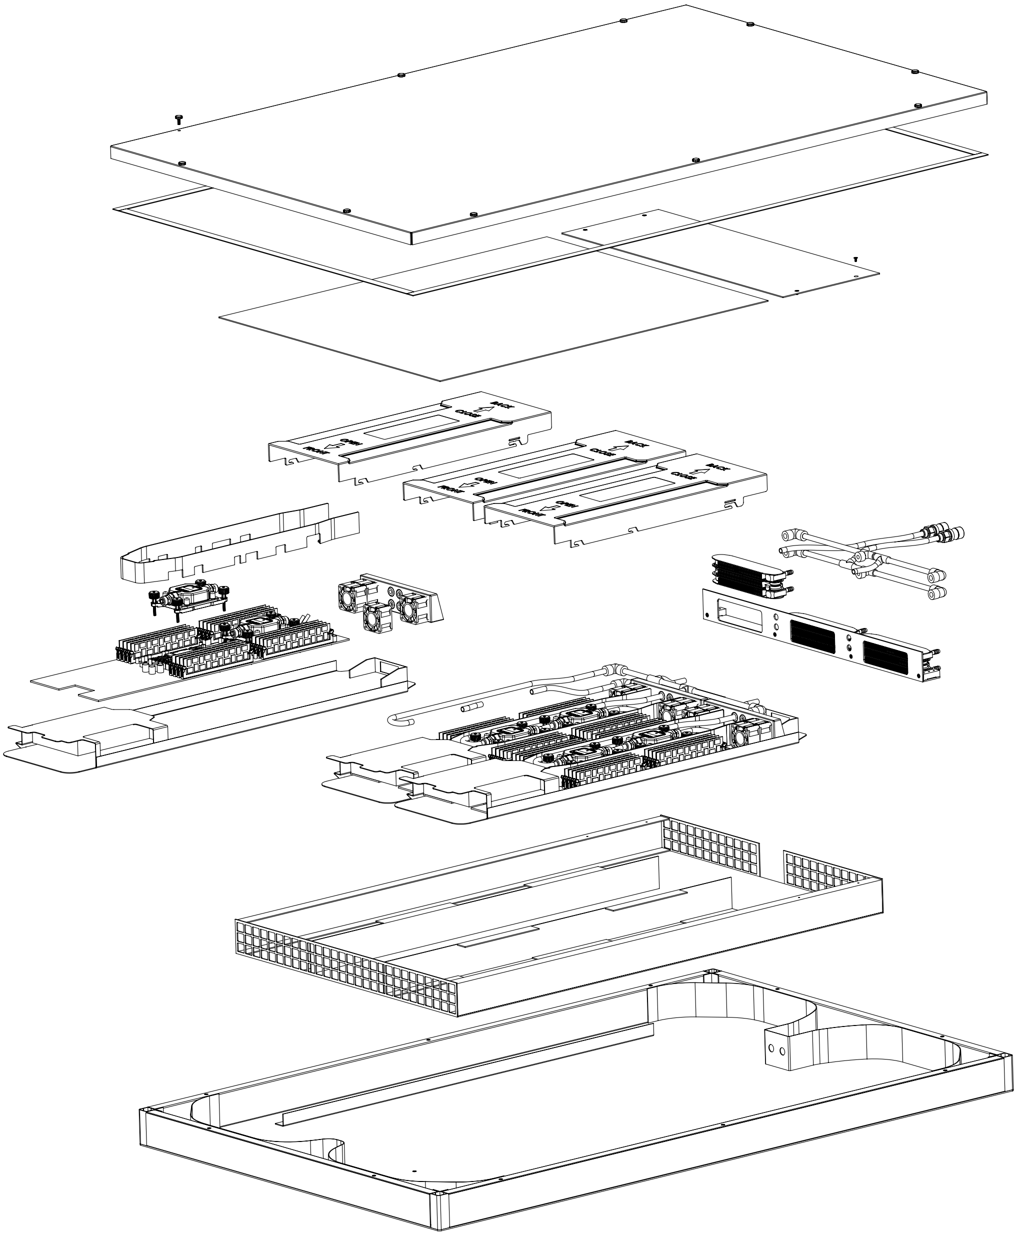

Key Nexalus Components

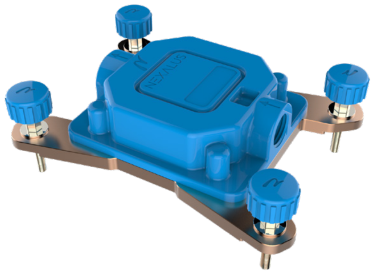

ENFLUX COLD PLATE

A high-end CPU cold plate, engineered for server-based applications with the Xeon chipset. More than two years' worth of research, calculation, experiment, CFD simulation and testing now results in the highest performing cold plate on the market for 1U server applications, where space is of a primary concern.

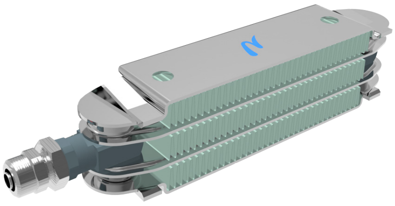

NXQ HEAT EXCHANGER

Our forced air-to-water heat exchanger is suitable for use in 1U server applications with limited real estate, while maintaining high-performance. This unit has two water inlets and one outlet, or vice versa with G1/4 fittings for pipe connection.

NEXALUS TORQUE SCREWS

Our TORQUE SCREWS allow a "tool-less" install of our cold plates with consistent and even force applied to the CPU, ensuring optimal TIM performance. A torque limiting mechanism is located within the head of each screw. This mitigates excessive chip temperatures resulting from incorrect installation of the cold plate.

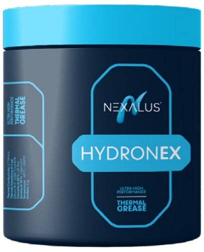

HYDRONEX TIM

The thermal interface material has a significant impact and can reduce thermal performance if the application is not carried out appropriately. The Hydronex thermal grease was chosen to aid thermal transfer between the IHS and the copper plate for Nexalus cooling solutions.

www.nexalus.com

CONTACT US:

Nexalus Ltd,
Unit 13, South Bank,
Crosse's Green,
Cork, Ireland,
T12 XT71

Nexalus Labs,
Water St,
Loughanalla,
Castlepollard,
Co. Westmeath, Ireland,
N91 EP90

Nexalus Research,
Trinity Research & Innovation,
O'Reilly Institute, Trinity College,
Dublin 2, Ireland,

info@nexalus.com
www.nexalus.com

Registered in Ireland No. 628880,
VAT No. IE3573607SH

IN PARTNERSHIP WITH:

Munters

mergon

arteco
ENGENIOUS COOLANTS

Trinity College Dublin
Coláiste na Tríonóide, Baile Átha Cliath
The University of Dublin

8PACK
APPROVED
WORLD RECORD BREAKING PERFORMANCE

WITH ENDORSEMENTS AND SUPPORT FROM:

ENTERPRISE IRELAND
where innovation means business

Science Foundation Ireland sfi
For what's next

QUALITY CERTIFICATIONS:

IATF 16949 / Automotive Standard
ISO 13485 / Medical Device Standard
ISO 50001 / Energy Management Standard
ISO 45001 / Occupational Heath & Safety Standard
ISO 14001 / Environmental Standard

NEXALUS
COOLING THE CLOUD