



WHITE PAPER: OPEN EDGE SERVERS MEET LIQUID COOLING - A FEASIBILITY STUDY

Author (s) :

Lucas D. R. Ferreira, Nokia

Samuli Toivola, Nokia

Date: February 2024

Executive Summary

This study updates the exploration of applying advanced cooling concepts to the Open Edge servers (OE). Exploring the use of direct-to-chip liquid cooling to the server CPU, which is the component with the highest cooling demand. The required design modifications to the server sled and to the OE chassis are discussed. In addition, the results of the study are presented, showing the potential of this application in edge scenarios, and the open challenges for the productization of such systems. With this, Nokia also aims to foster collaborations in developing future specifications of advanced cooling solution for the Open Edge platform, through the Open Compute Project community, which include the perspectives of hardware vendors, hyperscale data center operators, telecom and colocation providers.

Table of Contents

1. Introduction	4
2. Compliance with Open Compute Project Tenets	7
2.1. Openness	7
2.2. Efficiency	7
2.3. Impact	7
2.4. Scale	7
2.5. Sustainability	7
3. Open Edge cooling	8
3.1. The air-cooled Open Edge	8
3.2. The liquid cooled Open Edge	10
4. Performance of the liquid cooling solution	14
5. Open challenges	20
6. Conclusion	22
7. Glossary	22
8. References	23
9. License	25
10. About Open Compute Foundation	25

1. Introduction

In this white paper, Nokia continues to explore the possibilities of applying advanced cooling concepts to the Open Edge (OE) servers, this next step builds on an earlier white paper contribution to the Open Compute Project (OCP) [1]. In the previous contribution, the need for exploring advanced cooling solutions to the OE servers was discussed in detail, presenting some base level requirements to allow an evolutionary approach to advanced cooling in OE. In this feasibility study the use of direct to chip liquid cooling to the server CPU is explored.

The CPU is the server component with the highest cooling demand, and that continues to evolve towards the increase of processor's core count, I/O speeds, and higher power consumption. The trend towards increased power consumption promotes higher thermal dissipation needs, defined by the thermal design power (TDP). This can be confirmed by the data shown in Figure 1, in which, the evolution of the TDP for the different generations of Intel® CPUs in the server application vertical is shown. The trendline presented in this TDP evolution points towards an exponential growth of the average TDP values, which has accelerated for the highest performance CPUs in the last five years.




Figure 1: Evolution of TDP for Intel® CPUs in the server application vertical. Center marked are average values and bars indicate total range for each generation, with data from [1].

While the energy efficiency of the processors has improved over time, the trend towards ever higher CPU performance has offset those gains, resulting in the observed exponential increase of the average TDP. Additionally, a second trend can be observed in the reduction of the maximum allowable case temperature (T_{CASE}) for highest TDPs, as shown in Figure 2. The combination of these two trends promotes the need for improved CPU cooling, also impacting the need for cooling of other server components, such as the CPU VRMs.




Figure 2: Maximum allowable case temperature (T_{CASE}) as a function of Intel ® CPU thermal design power, for different generations in the server application vertical (497 data points). Trendlines are defined as linear interpolation of maximum and minimum values of T_{CASE} for each CPU TDP range. With data from [1].

From the analysis of Figure 2, the conclusion that the need for improved cooling can be evident. This trend has also been indicated by the 5th edition of the ASHRAE's "Thermal Guidelines for Data Processing Environments", which includes a new class of IT equipment, the H1 class. This new class is defined by ASHRAE for environments with a tight integration of high-powered components, being therefore described as high-density systems. The H1 class recommends a narrow air temperature band, from 18°C to 22°C, with an allowable envelope limited to 25°C, to meet the needed cooling requirements. [3]

While the reduction of the data center inlet temperature is being promoted for higher density systems, the opposite can be proposed for the increase of site energy efficiency. Additionally, there are physical limitations for the heat transfer with gases, as in air cooled servers. This can also be corroborated by the expectations that air cooling might not be able to support higher IT loads of 1 MW per rack or greater, which are expected to become typical in the next decade [4]. Notably there is also the need for high performance heat sinks, which can be both larger and more expensive.

As an example of these cooling challenges, consider that the CPU package sizes have significantly increased with the new generations of Intel® CPUs for server application, as presented in Figure 3.a. This package surface area increase could be beneficial for spreading the heat produced. Though, the increase in the TDP has progressed at a faster rate, causing the heat flux density (ratio of CPU TDP to package surface area – W/cm^2) to increase, particularly in the last five years, as shown in Figure 3.b.




Figure 3: Intel® CPUs in the server application vertical, in (a) the evolution of the package surface area, and in (b) evolution the ration of CPU TDP to package surface area – heat flux density [W/cm^2]. Center marked are average values and bars indicate total range for each generation, with data from [1].

It is then clear that there are reasons to consider the feasibility of advanced cooling solutions, such as the use of direct liquid cooling (DLC). In the next sections, the OE cooling is presented, together with the design modifications required to add liquid cooling capabilities, followed by the results of the adoption of DLC, and finally discussing additional challenges that could be expected for the productization of this solution.

2. Compliance with Open Compute Project Tenets

2.1. Openness

This contribution shares openly design selections and technical aspects of the application of liquid cooling to the Open edge server sled. While this white paper is focused on the feasibility study level, it enables other users to benefit from the design effort and build new advanced cooling solutions for the platform.

In system level Open edge solution components specifications are shared openly, the open edge chassis (OCP accepted) and the open edge server (OCP inspired) are open designs.

2.2. Efficiency

The Open edge server is designed to operate in telco site environments with limited power feed capacity and higher thermal requirements, aiming at high efficiency through technology selections. This feasibility study aims to take a further step into increasing the system efficiency and thermal performance through the application of liquid cooling solutions.

2.3. Impact

The Open edge server allows OCP designs to address market segments such as telco and edge computing. The proposed liquid cooled Open edge server contributes to meet the needs of new adopters and demanding environmental conditions, which may lead to more contributions and integration between Edge and Advanced Cooling Environments.

2.4. Scale

An Open edge system is highly scalable ranging from low end, single server to multi-chassis, multi rack systems. Server design is fully toolless and can be deployed easily to different customer environments with provided supporting documentation. The liquid cooled solution proposed will require modifications in server and chassis designs, also requiring adaptations to rack and facility levels.

2.5. Sustainability

The open edge physical footprint and material selections target to minimized use of natural resources, and the system design considers the energy efficiency and thermal performance parameters. The consideration of adding liquid cooled capacity to system then aims to further improve its cooling capacity, allowing the use of higher power units in the same footprint, and providing additional opportunities for facility level free cooling and heat reuse applications.

3. Open Edge cooling

The Open Edge server sleds are mounted to the Open Edge chassis [5]. While the chassis provides a steel enclosure, support brackets for the OE servers, as well as power distribution and grounding, it does not have fans nor integrated air filters. Instead, the cooling is integrated on the Open Edge sleds and PSU's, with all thermal events being dealt with at the sled level. This way, each sled has its own and independent cooling control. The sleds must operate at full specified capacity, within the required environmental conditions.

Open edge server sleds are required to be compliant with key environmental and regulatory standards. Some of these key operating conditions are:

- Normal operating temperature range from +5 to +40°C [ETSI EN300 019-1-3 Class 3.2 and GR-63-CORE]
- Short term operating temperature from -5 to +5°C and from +40 to 50°C [GR-63-CORE]
- Operating relative humidity from 5% to 90% [ETSI EN 300 019-1-3, Class 3.1]
- Seismic tolerance (Zone 4 in GR-63-CORE)
- Safety (IEC 62368-1:2014 and GR-1089-CORE)
- Fire resistance (shelf level criteria in GR-63-CORE)
- Acoustic noise (equipment room criteria in GR-63-CORE)

These specifications have an important impact in defining the working conditions from the cooling perspective. While also impacting the development of any innovative solution. Following we present a perspective of the current adopted air-cooled design of the Open Edge server sleds, and a view of the modifications required for the adoption of liquid cooling in the OE platform.

3.1. The air-cooled Open Edge

The current Open Edge server sled cooling design is based on air heat transfer, aiming to maximize heat sink capacity in a limited physical footprint. This is most critical for the 1U open edge server, as illustrated in Figure 4.




Figure 4: Open edge 1U server sled [6]

The air-cooled Open Edge server sleds are specified to a TDP of 500 W, for the 1U variant, and 800 W for the 2U version. The maximum CPU TDP is of 230 W for the 1U, and of 300 W for the 2U. Additionally, the maximum power levels for PCIe cards is of 100 W in 1U, and up to 300 W in 2U. To achieve this cooling capacity, the Open Edge servers make use of an Enhanced Volume Air Cooling (EVAC) heat sink for the CPU and VRMs, which is presented in Figure 5. This EVAC heat sink, presents heat pipes and a vapor chamber, in a maximum height of 24.2 mm. There are also heat sinks for the platform controller hub (PCH) and ethernet chipsets, in which extruded aluminum heat sinks are used.

The airflow in the Open Edge server sleds is achieved using dual rotor fans, in the 1U version four 4056 fans, and in the 2U version two 8056 fans. The cooling design allows the airflow to be both front-to-rear and rear-to-front. With the fans at maximum speed an airflow of 100 CFM is expected for the 1U server sled, and 160 CFM for the 2U server. The cooling algorithm is designed to adapt the fan speeds to the environmental conditions, providing the adequate airflow with the minimum possible fan power consumption. The sleds are specified to tolerate the failure of a single fan rotor.



Figure 5: Open edge thermal solution for the CPU and VRMs, using an EVAC heat sink design.

3.2. The liquid cooled Open Edge

A proof-of-concept (PoC) design is proposed to analyze the feasibility of adopting liquid cooling for the Open Edge server and chassis system. This PoC study is mainly focused on testing the potential of the cooling solution, while analyzing the possible challenges and drawbacks of implementing such solution in an actual edge environment application. The base platform for the liquid cooled server was the AirFrame Open Edge Whitley Server Barebone 1U . As it presents a mature system, to which different hardware configurations could be tested, focusing the study mainly on the contribution of liquid cooling to the server performance.

The modifications required for the implementation of liquid cooling involved both the OE server sled, and the OE chassis. Which included, providing the system with hydraulic connections, manifolds for liquid distribution, and the direct-to-chip cold plate for cooling the CPU.

From the server perspective the modification for liquid cooling has an impact on the fans assembly, as it is necessary to have the available space for the hydraulic connections. One of the dual rotor 4056 fans is removed, and replaced by two hydraulic support blocks, these blocks provide the connection from the server to the chassis manifolds, and for the internal hoses on which the liquid is pumped to the CPU cold plate. One additional modification is the inclusion of a small profile heat sink to help cool the VRMs, as the original EVAC heat sink present in the air-cooled OE servers also connects to the VRMs (Figure 5). A view of the modified liquid cooled OE 1U server sled is presented in Figure 6.

The selection of the cold plate for this PoC was based on the availability of components for the CPU socket used in Open Edge server, and the available space in the 1U server sled. Leading to the selection of the cold plate EK-Pro CPU Water Block LGA-4189 Rack (Ni+Acetal), from EK Water Blocks. [7]. Compatible G 1/8" hose fittings are connected to the cold plate, a flexible EPDM rubber hose and adequate hose clamps are used inside the server.




Figure 6: OE 1U Server with the cold plate, hoses, and modified fans module with integrated hydraulic connectors.

The hoses are then connected to the hydraulic support blocks inside the server. These support blocks are designed to replace one of the 4056 dual rotor fans in the 1U OE server, providing a fixed point for the hose fittings and to the quick-release couplings. The support blocks are fixed to the Open Edge sled tray, being aligned to the back plate, as shown in Figure 6. There is a 2 mm gap between the two hydraulic support blocks, avoiding the direct heat transfer between the inlet and outlet blocks.

The connection between the Open Edge sled and chassis is done through a back plate mounted PCB, in which the power and connectivity plugs are aligned by a guiding pin and receptacle construction [5]. The same system needed to be kept for the PoC, maintain the restriction that all normal operations are done only at the front of the chassis, without the need of access to the back of the rack. This design was made possible using blind-mate quick-release couplings, which also needed to provide non-spill and non-drip connectivity, and have a small footprint, to be compatible with the available space. At the server sled the connectors used were the Stäubli CGD-03 socket couplings [8], which provided connectivity with a misalignment of up to 1 mm, allowing the connection even with relaxed tolerances between the OE sled and chassis.

At the chassis level, the modifications to allow liquid cooling require changes to the backplate, with the addition of manifolds for the distribution of fluid, and the mechanical support structure. These modifications also include the addition of port holes for couplings between the manifolds and the sleds. The design adopted is shown in Figure 7. The supporting bracket is bolted to the chassis back plate, making it a single and independent unit, which should still be compatible with common 19" racks, given that the required rack depth is available. The manifolds present CGD 03 module plugs from Stäubli [8], coupled to the OE Server sleds. The manifolds are connected to the main inlet and outlet hoses that are connected to the coolant distribution unit (CDU). This connection is done with the use of hand-mate non-spill quick release couplings [9].




Figure 7: Open edge chassis adapted for liquid cooling, with the chassis level liquid distribution using manifolds, and the supporting mechanical structure.

In the design adopted, the manifold plugs protrude to the inside of the chassis, while the sled sockets protrude to the outside. This is done in a way that does not allow the modified for liquid cooling OE server to be inserted into a non-modified OE chassis. On the other hand, a common air-cooled OE server sled can be installed to the modified chassis, without the need of any modifications to the air-cooled server sled. Allowing the interoperability of air-cooled servers in a liquid cooled rack, with a penalty of partially obstructing the air flow for one of the server fans.

One important point that was considered in the design of the chassis manifolds was the proper distribution of liquid between the servers. Considering that the chassis must be able to operate with one to five liquid cooled servers, without the need of an operator interference to equalize the fluid flow. This was achieved by designing the server sleds to have a determined pressure drop, and then designing the manifolds with a parallel hydraulic circuit, with uniform fluid

distribution, using computational fluid dynamics (CFD) tools. This was verified in practice using a combined CDU and heat exchanger test stand equipped with sensors, presented in Figure 8.




Figure 8: Coolant distribution unit base design, with integrated sensors and flow control.

The coolant distribution unit (CDU) and combined heat exchanger (HEX), Figure 8, was equipped with 3/2-way flow control diverting valve, allowing the change of the flow rate to test different modes of operation. This unit was also equipped with a flow meter (FM), a temperature sensor (TS), a differential temperature sensor between the inlet and outlet flows (dT_S), and pressure sensors on the inlet and outlet (PS). The unit also allowed the control of the heat exchanger to cool to a specified temperature setpoint, this way the temperature of the liquid could also be adjusted to test different operating conditions. The combination of the CDU sensors and controls with the OE server sensors will provide the necessary data collection for the analysis of the performance of the liquid cooling system.

One last and fundamental consideration about the liquid cooling system is the choice of the coolant solution, which must be done in account of the compatibility of the wetted materials, and the desired system performance given the tested operational conditions. Important work has been published by the OCP Cooling Environments project in this regard, notably [10], [11], and [12], which were referred to in the selection of the cooling fluid for this proof-of-concept.

Given the system restrictions and expected test conditions, the fluid selected was a mixture of deionized water, ethylene glycol with 30% concentration, and additive Sodium-2-Ethylhexanoate with a 2% concentration. This mixture provides a good combination of heat transfer and antifreeze properties, with corrosion protection. The development of liquid cooling mixtures was not the focus of this study, and additional considerations of safety, environmental questions, and long-term operability must be considered.

4. Performance of the liquid cooling solution

To analyze the viability of adopting direct-to-chip liquid cooling in the Open Edge system, a series of different tests were performed in a controlled thermal chamber. Exploring different hardware configurations, while executing stress loads that could emulate the expected conditions in demanding edge scenarios. The stress loads included the Intel Power Thermal Utility ® Core AVX-512 test, memory fast read/write cycles, continuous data transfer between SSDs, and ethernet connectivity test with Ping to all units NIC cards. This high load testing was conducted in a systematic manner, while recording the data from the sensors, of both the CDU and the OE servers.

One of the tests performed was a direct comparison between a liquid cooled OE server, and a traditional air-cooled unit. It's expected that the liquid cooled CPU will present lower temperatures, but the aim of this test was also to verify the impact to components in the mother board of the modified server sled, notably the VRMs that are present in the downstream of the airflow and close to the CPU. The hardware configuration adopted in this test was selected to represent a common use case of 1U OE servers and is presented in Table 1.

Table 1: Server configurations of the OE 1U server used to compare the air-cooled design and the modified liquid cooled version.

CPU	6338N Intel processor 32c 2.2GHz 185W
Memory (x8)	RDIMM DDR4 32GB 3200
OCP slot	100GbE dual port OCP3 NIC G4 E810
PCIe slot	25GbE dual port LP-MD2 NIC card CX5
Drive slot	SSD 4TB NVMe Gen4 1dwpd E1S
M.2 slot (x2)	SSD 480GB SATA 1dwpd M.2 2280

Notably, the first tests also served to verify the functionality of the modification design adopted for the PoC. This included initial leak test of the hydraulic connectors, the verification that there was a uniform distribution of fluid in all chassis positions, and the assessment of the locking force required for the blind-mate hydraulic connectors. These tests showed that the design adopted in the PoC could in fact provide a good solution, and the modifications done to the OE server sled and OE chassis did not impact the functionality of the system for the expected thermal tests.

The comparison analysis then proceeded by checking the CPU temperature, with the results presented in Figure 9. The CPU used in the tests was stressed with the Intel Power Thermal Utility® to a measured power consumption of 184 W. In this analysis the cooling liquid was kept at 25°C, while the air inlet temperature was controlled using a thermal chamber. As shown Figure 9, in the room temperature and idle conditions the CPU temperature was dominated by the liquid temperature, while the air-cooled CPU temperature was 14°C higher. With the system under stress loads the difference is more significant, with the liquid-cooled CPU keeping a constant temperature for all ambient temperatures, with notable margin from the maximum CPU temperatures.




Figure 9: Comparison of CPU temperatures for liquid and air-cooled OE server sleds, for different ambient temperatures, in idle and under stress testing.

Apart from the CPU temperature, all other available temperature sensors on the motherboard were monitored during the tests. The CPU VRMs were of particular interest since these components could be impacted by the limited air flow available in the modified liquid-cooled server. The comparison for the reported VRMs temperatures is shown in Figure 10.




Figure 10: Comparison of VRM temperatures for liquid and air-cooled OE server sleds, for different ambient temperatures, in idle and under stress testing.

The results presented in Figure 10 show that, contrary to what was considered in the design phase, the VRM temperatures were not affected in the room temperature idle case by the liquid cooling modifications to the sled. Additionally, the VRMs presented lower temperatures in the liquid cooled server than in the air-cooled 1U OE under stress loading. This can be attributed to two factors, the first one was the use of a VRM heatsink in the liquid cooled server, which helped to promote the cooling of these components. The second factor is the reduction of the overall motherboard temperature under liquid cooling, although only the CPU was directly liquid cooled, the reduction of its temperature helped to reduce the spread of heat for all close components.

Other factors to consider in the tests were the selection of the appropriate liquid temperature and flow rate, as both could be adjusted in the CDU and Heat Exchanger test apparatus. The temperature of the cooling liquid was determined to be mostly relevant in the idle case during the tests, in which case the CPU temperature was kept at the same temperature as the cooling fluid. When the system was under stress, the use of warm liquid helped to reduce the overcooling of the CPU, while still maintaining a significant margin over the air-cooled CPU, even when the cooling liquid inlet temperature was of +45°C.

To determine the adequate liquid flow rate, pressure drop and cold plate thermal resistance test were performed, as presented in Figure 11. These curves show that, at higher flow rates there is no significative reduction of the thermal resistance. On the other hand, the pressure drop presents an exponential increase with the flow rate, implying higher energy costs for fluid pumping in higher flow rates.

A fluid flow between 1 and 2 l/min per server was selected as presenting a good balance between promoting high heat transfer (with a lower thermal resistance in the cold plate), while having a low pressure drop, and therefore allowing the use of small pumps, with lower energy consumption.




Figure 11: Pressure drop of a single liquid cooled OE server sled including manifolds and connections to the chassis, and the thermal resistance of the cold plate to the CPU, as a function of the volumetric fluid flow rate.

After comparing the modified liquid cooled OE server to the traditional air-cooled design, the test of a higher performance system was proposed, giving a better idea of the limits of this PoC design. This system was still based on the 1U OE, however CPUs with a higher TDP and processing capacity were used. In this test, five servers were built using the liquid cooling design, and different system configurations tested, including variations in the CPU, network, and storage capacity. The configurations tested are presented in Table 2.

Table 2: Server configurations of the OE 1U servers modified for liquid cooling, for high thermal stress.

	Server 1	Server 2	Server 3	Server 4	Server 5
CPU	8380 Intel processor 40c 2.30GHz 270W	8368 Intel processor 38c 2.40GHz 270W	8380 Intel processor 40c 2.30GHz 270W	8368 Intel processor 38c 2.40GHz 270W	8380 Intel processor 40c 2.30GHz 270W
Memory (x8)	RDIMM DDR4 32GB 3200	RDIMM DDR4 16GB 3200	RDIMM DDR4 32GB 3200	RDIMM DDR4 32GB 3200	RDIMM DDR4 32GB 3200
OCP slot	100GbE dual port OCP3 NIC G4 E810	-	-	100GbE dual port OCP3 NIC G4 E810	25GbE dual port OCP3 NIC G4 E810
PCIe slot	25GbE dual port LP-MD2 NIC card CX5	100GbE dual port A-A FHHL NIC G4 E810	100GbE dual port LP-MD2 NIC G4 E810	100GbE dual port LP-MD2 NIC G4 E810	100GbE dual port A-A FHHL NIC G4 E810
Drive slot	SSD 4TB NVMe Gen4 1dwpd E1S (x1)	SSD 2TB NVMe Gen4 1dwpd E1S (x2)	SSD 4 TB NVMe Gen3 1dwpd E1S (x2)	-	-
M.2 slot (x2)	SSD 480GB SATA 1dwpd M.2 2280	SSD 480GB SATA 1dwpd M.2 2280	SSD 480GB SATA 1dwpd M.2 2280	SSD 480GB SATA 1dwpd M.2 2280	SSD 480GB SATA 1dwpd M.2 2280

The results from this higher performance test of the liquid cooled OE 1U server are shown in Figure 12. Once again, the CPU temperatures are the same as the cooling liquid in the Idle case. For the system under stress loads, the CPU power was up to 270 W, which is above the maximum 200 W allowable for the air-cooled 1U OE. Still for the system under stress the CPU temperatures have about 30°C of margin to air-cooled OE server, indicating that liquid temperatures of up to +55°C could be adopted. One additional advantage of the liquid cooled design is stability of the CPU temperature, even for ambient temperatures of +50°C, no significant impact is observed.




Figure 12: CPU temperatures of liquid cooled high power 1U OE server sleds for different ambient temperatures, in idle and under stress testing.

In this study, the modified 1U OE uses the mixed cooling approach, with the CPUs being liquid cooled and the other components in the motherboard being air-cooled. The analysis of the motherboard temperature sensors showed that the adopted design had an overall positive impact on cooling, with the components either being cooler (Figure 10) or having no relevant temperature differences than observed in the traditional OE server. As the system dependence in air-cooling is reduced, two alternatives could be explored.

The first possibility enabled by the liquid cooling modifications is the adoption of a modified cooling algorithm, which would allow the fans to operate with reduced duties in the lower ambient temperatures. This option would promote a reduction in the power consumption by the fans, improving the overall system efficiency. This modified cooling algorithm is presented in Figure 13. From the cooling perspective, even lower fan speeds could be used, but in that case, there is a higher risk of fan stoppage and instability.




Figure 13: Minimum fan duty required by the cooling algorithm for different ambient temperatures, with the original air-cooled OE server sleds, and modified for liquid cooled.

The second possibility enabled by the mixed cooling approach is the reduction of the number of fans, as in the design adopted in this PoC there are still three dual rotor 4056 fans in operation. The use of only two fans could provide enough air flow to promote the necessary cooling. In this case, the fan in the middle (Figure 6) could be replaced by air blocking baffle. The main drawback would be reduction of the system redundancy.

Apart from the server energy consumption the use of different fan configurations in this mixed cooling design also promoted a reduction of the system sound levels. The system sound levels were measured for the full chassis configuration, in an in-house facility, adopting the microphone positioning indicated by the ISO3744:2010 [13] for the small machines class. The results are presented in Figure 14, showing the significant sound reduction, which would put the OE system in the ETSI sound power requirement for commercial business area [14]. These are non-official results, and more tests should be conducted for proper certification.




Figure 14: Comparison of sound levels for different fan configurations in the liquid cooled 1U OE server, for room temperature environment (+25°C). *Sound pressure is measured from 0.6 m of the rear of the chassis.

5. Open challenges

This PoC showed the potential of applying liquid cooling to the OE servers, allowing the use of higher power CPUs, while keeping the motherboard temperatures low, and with minimum modifications to the server sled. There are also benefits in reducing the acoustic noise levels, and the reduction of the cooling energy consumption, resulting in an increase of efficiency. There is also the potential for reducing failure rate of the electronic components, as the system can be run at constant and lower temperatures. Despite these benefits, there are fundamental open challenges for the productization of a liquid-cooled Open Edge system.

One of the challenges observed is related to the control of the CDU. In the design used the cooling liquid is evenly distributed to all the servers in the chassis, in a hydraulic parallel configuration. This way, if for example, there are 5 servers in the chassis and the design requires a flow rate of 1 l/min per server, the CDU must be regulated to provide 5 l/min for the chassis. In this scenario, if one of the servers is removed from the chassis, the leakproof hydraulic couplings would close the flow for the removed server sled, and each server would now have a flow rate of 1.25 l/min, with an increase of the system pressure.

In the tests conducted for this study, this removal or addition of servers to the chassis during operation (hot swapping) did not result in any damage to the system, or any other disturbance apart from the pressure and flow rate variations. To keep a uniform flow rate during the tests, the manually operated flow control valve had to be adjusted. In a field application this could be challenging, particularly considering a full rack configuration. Therefore, the automation of the CDU, and the integration to the Rack Mounted Controller (RMC), would be necessary. This would imply the need of a common communication standard between CDU vendors and hardware manufacturers.

The design of the cold plates could be another source of challenges for the productization of liquid cooled edge servers. In this PoC only the CPUs were liquid cooled, but there could be increased benefits of including other components, such as the VRMs, the disks, and PCIe accelerators. This would demand a closer integration to the design of the cold plates. There is also a need for standardization of cold plate connectors, as there are several hydraulic connections options, which are not compatible between different cold plate vendors. One

additional point to be considered would be the adoption of standard rate classes for the cold plates, that would facilitate the selection, based on pressure drop and thermal resistance from the liquid to the base.

While the tests conducted in this study showed different benefits, these were all realized in a limited time frame. Although the system accumulated more than 200 hours of operation during these tests, the long-term analysis was not in the scope of this project. The enduring stability and Mean Time Before Failure (MTBF) of the hydraulic couplings and hoses must be analyzed before any actual productization phases. This is particularly relevant for edge uses in the telecom market, in which components can be subjected to varying environmental conditions, and the equipment are expected to operate for longer time periods.

One last open challenge is related to the site level implementation of the liquid cooling systems. While this study used an integrated CDU-Heat Exchanger system, in practice these components would need to be independently procured and integrated at the facility level, considering different installation site requirements and environmental conditions. These challenges are common to all advanced cooling solutions and have started to be addressed by the OCP Cooling Environments project, particularly with the recent creation of the Open Technology Cooling System (TCS) workstream [15].

6. Conclusion

This white paper presented the adoption of a liquid cooling solution for the Open Edge servers, detailing the modifications necessary to the server sleds and OE chassis, and included results from the analysis of the system performance, showing the potential of this application in edge scenarios. While also discussing some of the open challenges observed in the feasibility study.

The benefits observed included the reduction of the CPU temperatures, and stability even at high ambient temperature conditions. Also, the capacity to cool higher thermal design loads, which could lead to the improvement of the system density when compared to the air-cooled servers. Other advantage is the reduced dependency to the facilities air cooling in this mixed cooling approach, promoting lower sound noise levels due to reduced fan speeds or even the reduction of the number of fans, which also leads to the reduction of the system energy consumption and increased efficiencies.

Despite the observed benefits, there are still important challenges and research to be conducted before a commercial product could be available. These are not definitive limitations, and it is expected that a form of advanced cooling solution will be part of future edge hardware products. In this context, there could be a beneficial collaboration between the OCP Edge project, and the work conducted by the OCP Cooling Environments community.

7. Glossary

CDU	Coolant Distribution Unit
CPU	Central Processing Unit
DLC	Direct Liquid Cooling
OE	Open Edge System
PoC	Proof-of-concept
T _{CASE}	Maximum allowable case temperature
TDP	Thermal Design Power
VRM	Voltage Regulator Module

8. References

- [1] L. Ferreira and S. Toivola, "WHITE PAPER: ADVANCED COOLING CONCEPTS FOR OPEN EDGE SERVERS - MOVING TOWARDS MORE SUSTAINABLE AND EFFICIENT SOLUTIONS," September 2022. [Online]. Available: <https://www.opencompute.org/documents/advanced-cooling-concepts-for-open-edge-servers-moving-towards-more-sustainable-and-efficient-solutions-white-paper-v-0-4-pdf-1>. [Accessed July 2023].
- [2] Intel , "Intel Core Processors - Product Specifications," June 2023. [Online]. Available: <https://ark.intel.com/content/www/us/en/ark/products/series/122139/intel-core-processors.html>.
- [3] D. Bizo, "New ASHRAE guidelines challenge efficiency drive," Uptime Institute, 21 October 2021. [Online]. Available: <https://journal.uptimeinstitute.com/new-ashrae-guidelines-challenge-efficiency-drive/>. [Accessed July 2023].
- [4] L. Simon, "Direct liquid cooling bubbles to the surface," Uptime Institute, 31 March 2022. [Online]. Available: <https://journal.uptimeinstitute.com/direct-liquid-cooling-bubbles-to-the-surface/>. [Accessed July 2023].
- [5] T. Männikkö, "Open Edge chassis - Revision 1.4," Open Compute Project, 27 10 2020. [Online]. Available: <https://www.opencompute.org/documents/open-edge-chassis-ocp-contribution-v1-4-1-pdf>. [Accessed July 2023].
- [6] S. Toivola, "White Paper: Next generation Open Edge Server designs," Open Compute Project, 2021.
- [7] EK Water Blocks, "EK-Pro CPU Water Block LGA-4189," [Online]. Available: <https://www.ekwb.com/shop/EK-IM/EK-IM-3831109844700.pdf>. [Accessed July 2023].
- [8] Stäubli, "CGD Quick-release couplings," [Online]. Available: <https://www.staubli.com/content/dam/fcs/brochures/products/cgd/cgd-thermal-management-quick-couplings-staubli-en.pdf>. [Accessed July 2023].
- [9] Stäubli, "Quick-release couplings SPT/CG," [Online]. Available: <https://www.staubli.com/content/dam/fcs/brochures/products/spt-cg/spt-non-spill-thermal-management-staubli-en.pdf>. [Accessed July 2023].
- [10] D. Sartor, F. P. Yu, C. Myers, S. Barkett and S. T. Sivapalan, "Guidelines for Using Water-Based Transfer Fluids in Single-Phase Cold Plate-Based Liquid-Cooled Racks," 2022. [Online]. Available:

<https://www.opencompute.org/documents/guidelines-for-using-water-based-transfer-fluids-in-single-phase-cold-plate-based-liquid-cooled-racks-final-pdf>. [Accessed July 2023].

- [11] K. Yaroch, J. Cuthbert, D. Miller, S. T. Sivapalan and S. Barlett, "Guidelines for Using Water-Based Transfer Fluids in Single-Phase Cold Plate-Based Liquid-Cooled Racks," 2022. [Online]. Available: <https://www.opencompute.org/documents/guidelines-for-using-propylene-glycol-based-heat-transfer-fluids-in-single-phase-cold-plate-based-liquid-cooled-racks-final-pdf>. [Accessed July 2023].
- [12] J. Gullbrand, N. Gore, J. Matteson and E. Langer, "ACS Liquid Cooling Cold Plate Requirements Document - Revision 1.0," 2019. [Online]. Available: <https://www.opencompute.org/documents/ocp-acs-liquid-cooling-cold-plate-requirements-pdf>. [Accessed July 2023].
- [13] International Organization of Standardization, "ISO 3744 - Acoustics: Determination of sound power levels and sound energy levels of noise sources using sound pressure - Engineering methods for an essentially free field over a reflecting plane," *International Standard*, no. Third edition, p. 86, 2010.
- [14] European Telecommunications Standards Institute , "ETS 300 753 - Equipment Engineering: Acoustic noise emitted by telecommunications equipment," October 1997.
- [15] Open Compute Project, "OpenTCS Mission Statement and Scope," July 2023. [Online]. Available: https://docs.google.com/document/d/14jO0p2IC57dM03RF8sb_onTbdogSZ9NKGoZdRFmywmg/edit. [Accessed August 2023].

9. License

© Copyright Nokia Solutions and Networks Oy 2022. All rights reserved.

10. About Open Compute Foundation

At the core of the Open Compute Project (OCP) is its community of hyperscale data center operators, joined by telecom and colocation providers and enterprise IT users, working with vendors to develop open innovations that, when embedded in product are deployed from the cloud to the edge. The OCP Foundation is responsible for fostering and serving the OCP Community to meet the market and shape the future, taking hyperscale led innovations to everyone. Meeting the market is accomplished through open designs and best practices, and with data center facility and IT equipment embedding OCP Community-developed innovations for efficiency, at-scale operations and sustainability. Shaping the future includes investing in strategic initiatives that prepare the IT ecosystem for major changes, such as AI & ML, optics, advanced cooling techniques, and composable silicon. Learn more at www.opencompute.org.