Implementation of the waste heat recovery measures at Aalborg – Action C2

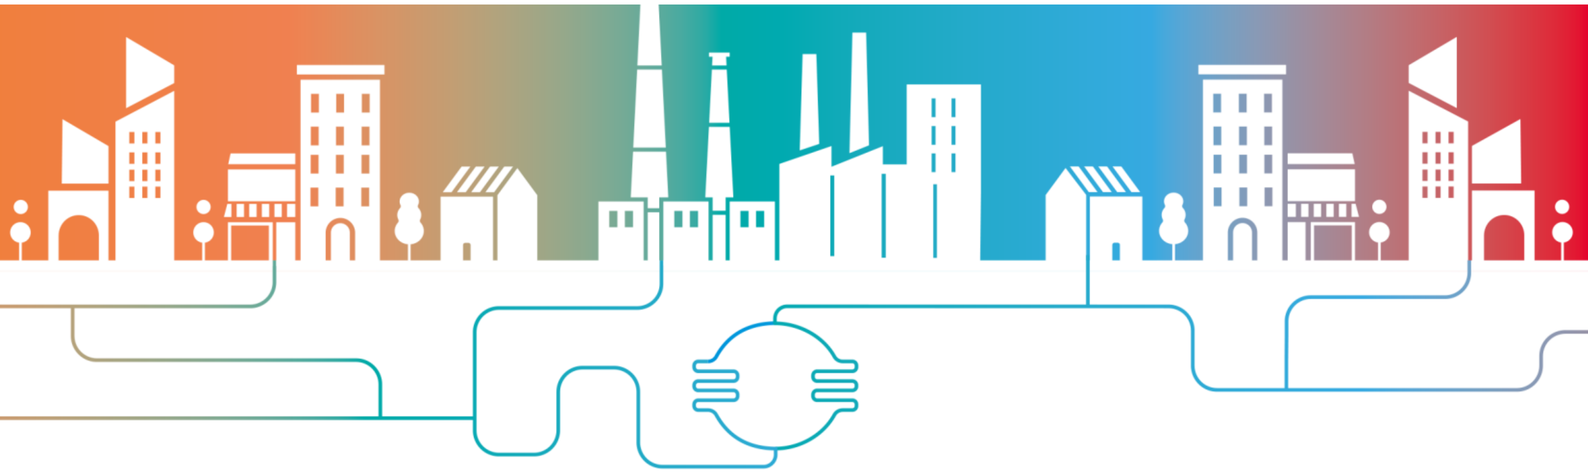

Low temperature, urban waste heat into district heating and cooling networks

as a clean source of thermal energy

LIFE4HeatRecovery

Project Title: Low temperature, urban waste heat into district heating and cooling networks as a clean source of thermal energy

Project Acronym: LIFE4HeatRecovery

Deliverable Title: Implementation of the waste heat recovery measures at Aalborg

Lead beneficiary: Heatflow ApS

Dennis Jensen, Heatflow ApS

Date: September 14, 2023

This document has been produced in the context of the Life4HeatRecovery Project.

The project has received funding from the LIFE programme of the European Union under contract number LIFE17 CCM/IT/000085

Table of Contents

1 Introduction

This deliverable describes the implementation and commissioning experience of the LIFE4HeatRecovery demo case at Aalborg University (Action C.2). In order to make the document more self-contained, this introduction recalls the general description of the system as reported in the corresponding design deliverables.

The mode case includes a complete waste heat recovery system (the prefabricated skid) designed and developed to meet the Tier III Standard in the current data center at Aalborg University (AAU). Skid design and preparatory construction works (room adaptation, connection with district heating) have been respectively described in the two deliverables “Design of prefabricated skid for waste heat recovery at Aalborg” and “Construction works for connection to the network at Aalborg”, both related to Action A.1 of the project.

The skid will demonstrate a highly efficient heat transfer from the data center server racks to the local district heating supply. It has been designed as a generic, scalable, and reproducible solution that can be installed in various data centers configurations.

The skid is based on an innovative cooling technology combined with an intelligent control system provided by ENISYST. This system allows for highly efficient cooling of server racks while simultaneously enabling the recovery of the generated heat into the local district heating network. The cooling system, developed by HEATFLOW, consists in a two-phase passive cooling system installed within server racks in an operating data center infrastructure in Aalborg, Figure 1 and Figure 2.

Figure 1 – Data center building in Aalborg.

Figure 2 – Data center building entrance.

The data center has an operating profile that can benefit from the usage of both a conventional air cooling system and second liquid-based cooling system, in combination with heat pumps to cool the servers and to recover waste heat.

A combination of these two cooling systems can increase the overall efficiency of waste heat recovery.

In the demo system, the current installation focuses on the innovative passive liquid cooling technology given by thermosiphon systems. The system recovers and transfers the waste heat to the district heating (DH) network (shown in Figure 3) and to the local heating system of the building hosting the data center, depending on the needs.

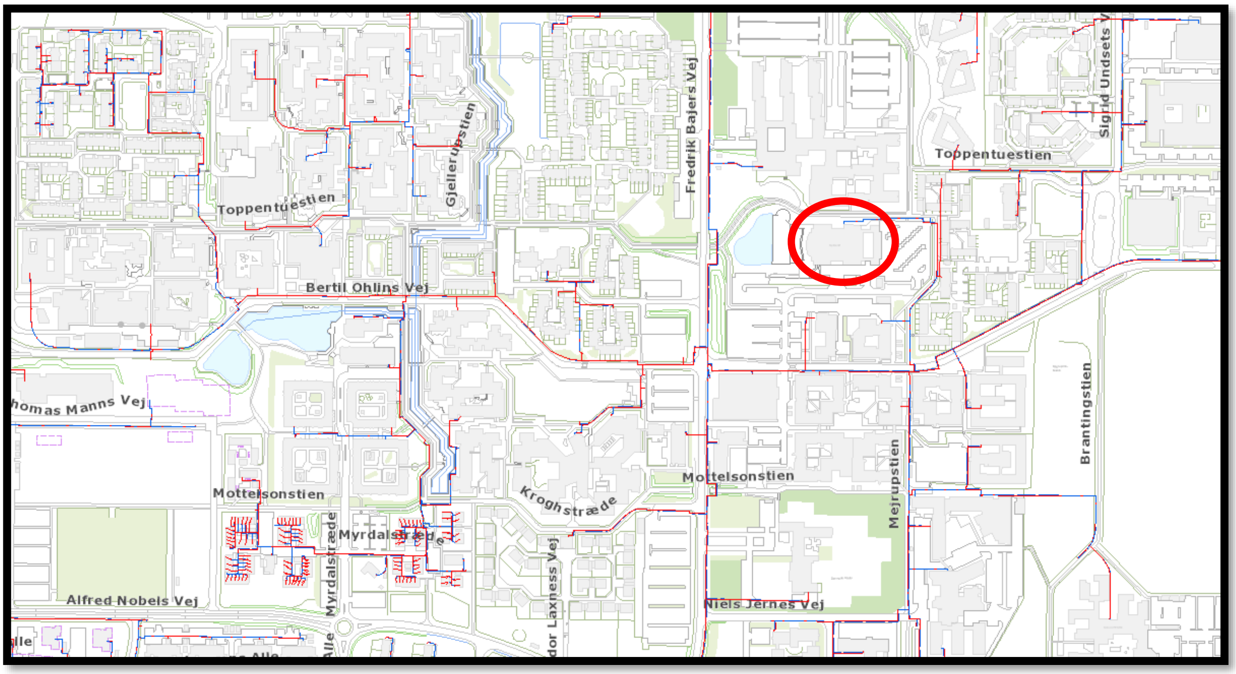

Figure 3 – District heating network (red lines). The data center building is circled in red.

The return temperature of the district heating network, of the order of 40°C or even lower, is heated up to the supply temperature, which is in the range 60-80 °C (higher in winter, lower in summer), in a specific waste heat recovery heat exchanger. On the primary side, the heat exchanger is powered by the waste heat recovered from the data center servers. The waste heat recovery is performed by a specific skid equipped with a heat pump to adapt the thermal levels to the required needs.

The two-phase cooling system has been installed on multiple server racks, providing a nominal thermal power of 13.6 kW (see below). According to market availability and in view of future expansions, a heat pump of about 30 kW of nominal cooling power has been installed. To decouple the heat pump from the HEATFLOW cooling system operation and cope with this power mismatch a thermal energy storage has also been installed.

The demo room has been featured with separate double door entrance (for equipment and visitors to the demo site) and lattice walls and door for separation to other areas. See Figure 4 and Deliverable “Construction works for connection to the network at Aalborg” for more details.

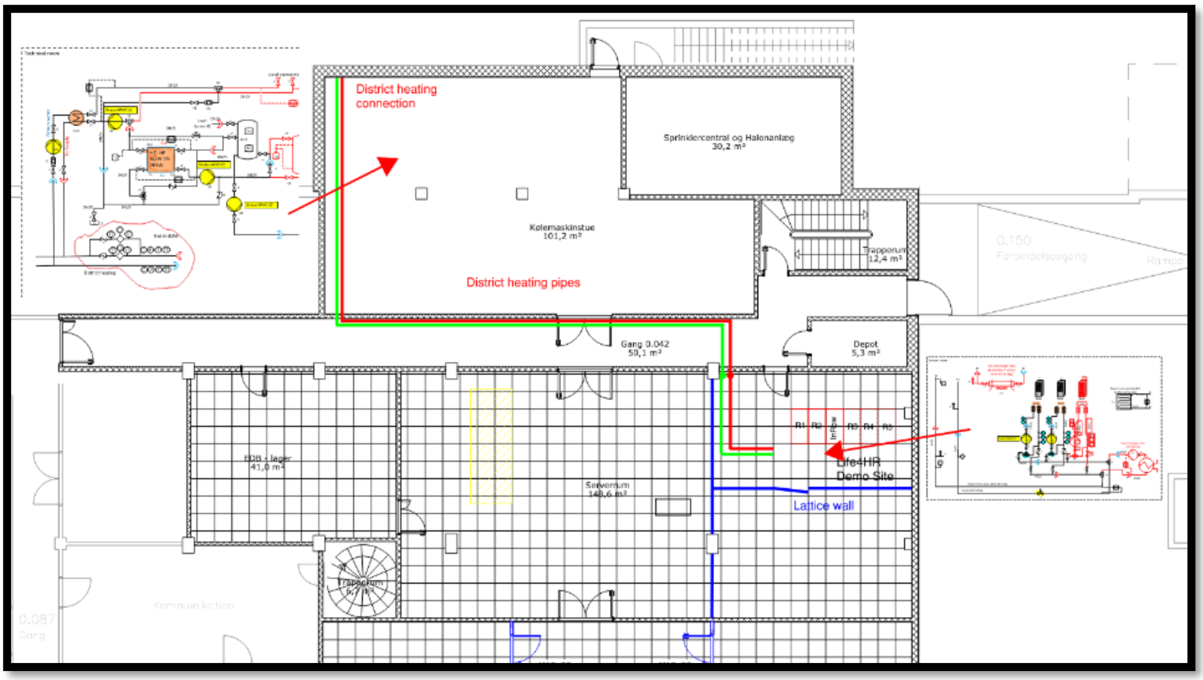

Figure 4. Floor map of the demo area.

2 Server description

Currently, the server room contains three server-racks installed: two of them are equipped with the HEATFLOW two-phase cooling system, one is available for future expansions.

For the connected racks, one rack is equipped with standard HPE DL380p 2U servers, the other one is equipped with “dummy” servers designed to simulate various CPU performance characteristics, with the objective to investigate system design optimizations. The higher controllability of dummy servers will allow more demanding stress tests, with more abrupt power profiles than with standard servers (equipped with more stringent safety controls).

In rack no. 1 there are 12pcs of HPe DL380p and in rack no. 2 there are 10pcs of dummy servers + 2pcs of HPe servers. The HPe servers each has 2x150kW of max CPU power dissipation where our dummy servers each have 2x500W of max “CPU” power.

There has been full support from the IT department at Aalborg University, which has provided support with both written and unwritten requirements and expectations for the installation of equipment with the aim of meeting current requirements for access to equipment and serviceability.

To mirror a state-of-the-art server room, an “In Row” AC unit to handle the air-conditioning of the room is installed. The installation is prepared for future back-up cooling from AAU facility cooling loop.

Figure 5 – Configuration of server racks in demo room.

3 Prefabricated skid description

3.1 Pump station

Figure 6 – Pump station.

Figure 6 shows the pump station for the heat pump and buffer tank which is mounted as a prefabricated skid. The skid is equipped with a heat exchanger for interface with the district heating network to avoid mixing of fluids. All major valves are actuator-controlled by the system control from ENISYST.

The pump station offers a good flexibility: it can operate both with and without the buffer tank (see Deliverable "Design of prefabricated skid for waste heat recovery at Aalborg" for the system P&ID); and it can switch the heat pump on and off so that the system becomes adaptive in relation to the cooling demand and the thermal levels.

Thanks to this flexibility, it is possible to test different scenarios and to find the optimal operating point.

3.2 Shut unit for server racks

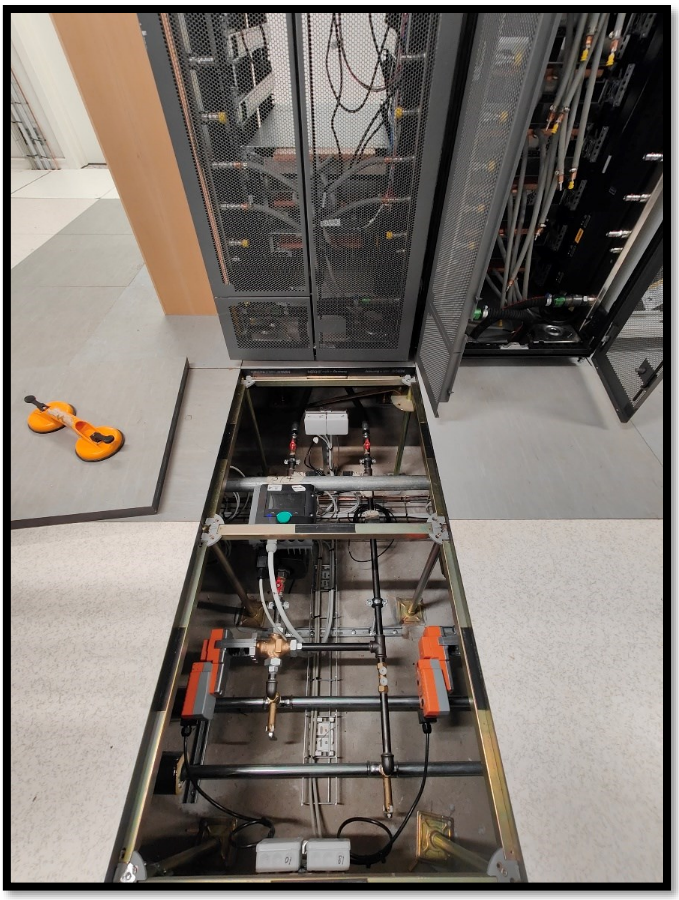

Figure 7 – Shut unit.

The installation is designed according to the Tier III Standard for data centers, and emphasis was placed on securing the installations against leaks, so that liquid cannot penetrate the racks and servers. The shunt unit, see Figure 7, can be prefabricated and installed under each server rack. It is equipped with a separate circulation pump so the main system pump can be used for backup. Hence, the skid is complying with the redundancy requirements for server farms.

3.3 HEATFLOW heat exchangers for servers cooling

In collaboration with the AAU IT department and with respect current requirements for the installation of liquid-containing equipment, a great deal of effort has been put into adapting the external part of the server cooling. Drip-free connections and approved pressure hoses are used to match the designed manifold. To ensure that the solution meets the service requirements, the length of the connecting hoses has been adjusted through several iterations, Figure 8.

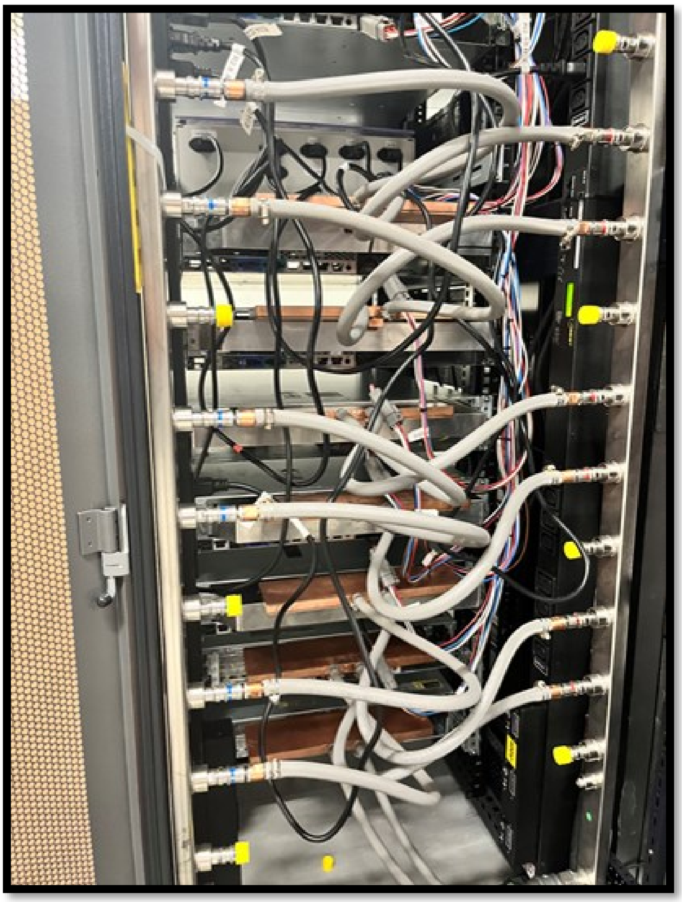

Figure 8 – Connection of HEATFLOW heat exchangers to the manifold.

As shown in the picture, there is a great flexibility in changing the servers and even adding more to the rack. This installation is based on U2 servers.

During the first installation of a new server, a process whereby all air is vented out locally is required. This prevents the formation of air bubbles into the system. The process is very easy thanks to the drip-free quick connectores.

The final CPU cooler design is shown in Figure 9.

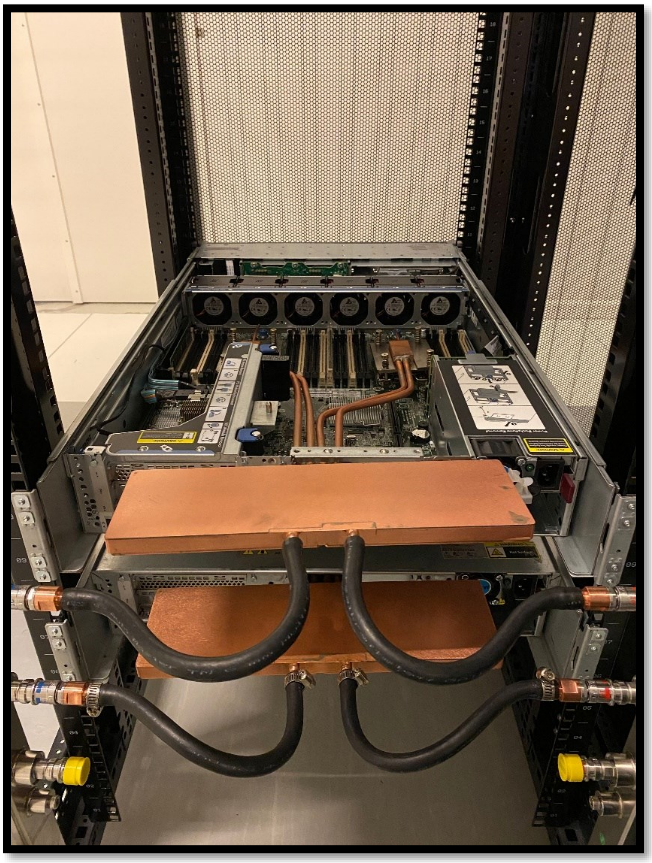

Figure 9 – Inside view of the CPU cooling configuration and connection to the heat exchanger.

The current cooling solution is using the refrigerant R134a to ensure an optimal performance. The design is prepared for R718 (water) going forward in the commercial replica: indeed, thanks to the full brazing, absence of leakage can be guaranteed and the usage of water within the thermosyphon circuits will reduce operating pressure thereby reduce costs and improve environmental parameters.

The cooling system design has been optimized throughout the construction of the demonstration site.

3.4 Heat pump and circuit

Based on the cooling demand, the thermal levels and the availability on the market, the chosen model of heat pump is KRONOTERM WPB-35-1 XHT, with a nominal thermal power of 30 kW and a nominal cooling power of 25-30 kW, Figure 10.

The refrigerant used is R134a. The heat pump has a compact design with a built-in control unit. The multi-compressor design ensures a stepwise regulation of power.

The heat pump is particularly suitable for applications with a relatively high heat source temperature (over 50 °C) and the need for high heating water outlet temperatures (up to 82 °C). These features make it well suited for the Aalborg case study: the heat source temperature is related to the operating temperature of the servers (approx. 60 °C) and the required outlet temperature on the condenser side is related to one of the district heating network (order of 80 °C).

Figure 10 – Heat pump installed with pumping station.

Throughout the project a great effort has been made to find this important part of the system. It was not easy to find a manufacturer for a high-temperature heat pump of the considered size. The selected manufacturer was first contacted by EURAC, thanks to a former collaboration in the EU REWARDHeat project. HEATFLOW finalized the market investigation and purchase process.

3.5 Control

The development of the skid planned two Phases, with a second optional phase for a future possible expansion.

In the first phase, implemented in LIFE4HeatRecovery, only heart recovery from the liquid cooling of server racks was implemented.

The skid can be divided into two different sections: the liquid cooling circuit of server racks and the waste heat recovery section including the heat pump. The two are separated by a thermal energy storage, which decouples their operation.

The skid must ensure that the outlet temperature is high enough to recover the waste heat in the district heating network at the required temperature (as well as to cover local heat demand). This is a controllable parameter (as the desired network supply temperature changes along the year).

This set point temperature also corresponds to the temperature set in the high-temperature heat pump HTHP (set point temperature at the condenser outlet). Modulation of circulation pumps ensures that this set point temperature is met.

The tank is endowed with temperature sensors which estimate its “state of charge”, starting the heat pump operation when the temperature is high enough and stopping it on the contrary. Thermal stratification in the tank plus proper recirculation circuits ensure both that a suitable return temperature to the rack liquid cooling system is available and that the desired temperature difference on the rack circuits is achieved.

For the possible future Phase 2, an air-cooling circuit would also be installed in the system. Phase 2 would involve the installation of a low-temperature heat pump LTHP to serve the air heat exchanger in the server room. Another thermal storage (at a lower temperature than the one of Phase 1) would be installed as a heat sink for the LTHP. The control logic described above should then be properly integrated.

For the current installation, ENISYST provided the implementation of the switchboards (see Figure 11) and of the needed control system.

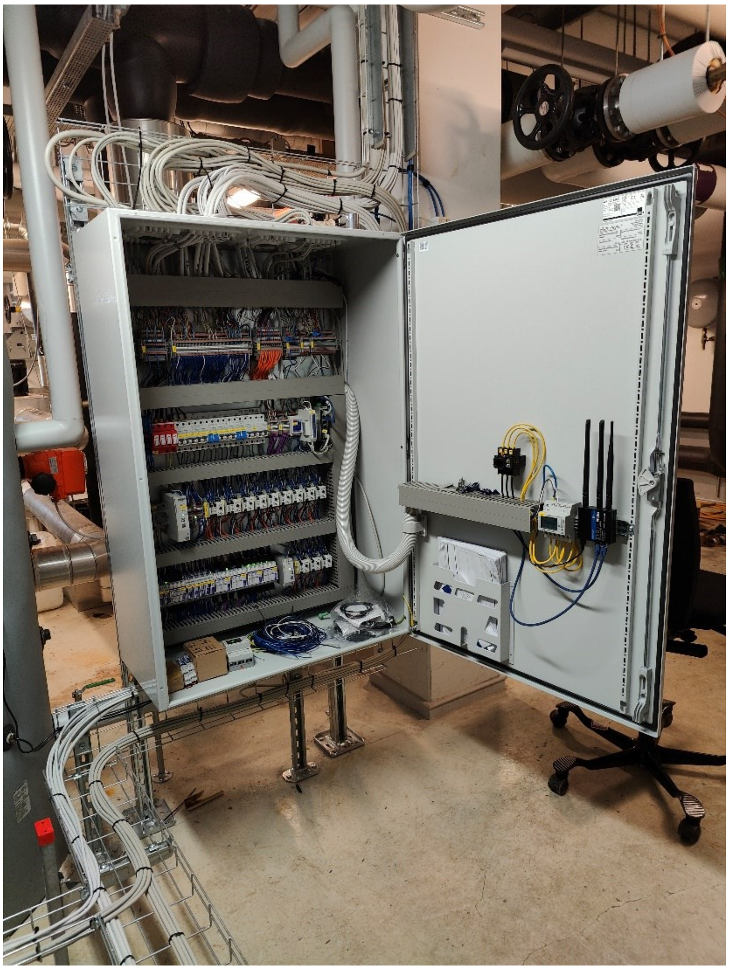

Figure 11. A picture of the main switchboard provided by ENISYST.

3.5.1 Software and control system

The control system of the entire hydraulic system, ranging from heat pump to single server cooling circuits, was developed by ENISYST. The figure below shows a screenshots from ENISYST interface, namely the data analysis interface (enabling plots for the various monitored variables).

Figure 12. The SCADA interface developed by ENISYST.

Beside the plant control system developed by ENISYST, a separate control software for the dummy loads was implemented by HEATFLOW. The latter was designed in order to manage separately each power input, providing a great test flexibility. This allows for a more individual control and a better understanding of variants in the performance.

There are two options to set the power load for each "server", either by effect (Watt) or by target maximum temperature (degC). Each dummy "CPU" has a potential of 500 W, so plenty of capacity for future testing.Typical operation occurs in the range 100-200 W per "CPU". This is matching the expected output from the HPe gen8 servers. In this way, the combined monitoring of both real and dummy servers will allow replicating similar stress conditions in a more controllable and diversified way.

4 Conclusions

During the design and integration phase of the Heatflow solutions into the 2U servers provided by the server OEM HPe, it was determined that achieving physical fit and functionality was entirely feasible. Upon installing a Heatflow upgraded server into the rack, no obstacles were encountered with the sliding brackets (rails) on each side of the server. Furthermore, the backside of the rack, where the water manifolds are situated, easily accommodated the water connecting hoses without imposing major constraints on the required power and network cables.

The design process of the thermosiphon system demanded meticulous attention to various components, including mechanical fix-points, manifolds, purpose-built connectors, and the refrigerant filling process. A significant revelation was the substantial impact of brazing quality during assembly on the final thermal performance. To address this challenge, strict process control and a professional manufacturing setup will be implemented for large scale production.

Optimizing the single thermosiphon circuit for maximum thermal performance was a critical consideration, with the conclusion being that a narrow margin of just 5 grams represented the optimal filling ratio. Mechanical design considerations also prioritized ensuring the serviceability of the server and ease of installation.

For future enhancements, the focus will be on improving the supply chain and reducing the cost price of the solution. This may entail material substitution from copper to aluminium and utilizing an HFO refrigerant with lower Global Warming Potential (GWP).

A crucial lesson learned from manufacturing the thermosiphons is the necessity of having all brazing joints done by a professionally controlled atmosphere brazing (CAB) furnace. Despite encountering leakage during increased operating temperatures, resulting in heightened internal pressure, the issues were addressed and will be mitigated in a controlled process.

Regarding the hydraulic system, the system mounting proceeded smoothly, providing a positive experience. The hydraulic system demonstrated efficient thermal capacity and quick, steady-state regulating control, resulting in stable "CPU" temperatures even under dynamic loads. Additionally, the hydraulic design showcased built-in redundancy in pumping capacity, proving valuable for future "tier 3" data center layouts. Successful integration with the district heating network requires early discussion in the process and involvement of all stakeholders.

Regarding the heat pump, while it is a crucial component for achieving the needed winter temperatures of the network, challenges arose due to power mismatches between the machine and the heat source, influenced by limited market availability at the given sizes. Several adjustments were necessary for full commissioning, such as changes in flow rate thresholds. However, the final startup of the system demonstrated the flexibility of the chosen design. Such issues are expected to be reduced in larger scale systems, where better size matching and other modulation options can be achieved.

In terms of control, the system offers excellent monitoring flexibility, online data exchange, and remote-control capabilities, which significantly enhance overall operational efficiency and control. The monitoring flexibility allows for real-time observation of critical system parameters, enabling prompt detection of any deviations or anomalies. This proactive approach to monitoring ensures that potential issues can be identified and addressed swiftly, minimizing downtime and optimizing system performance.

Moreover, the online data exchange functionality facilitates seamless communication and data sharing between various components of the system, as well as with external systems or stakeholders. This enables streamlined data analysis, collaboration, and decision-making processes, ultimately enhancing operational transparency and effectiveness.

The remote-control capabilities further augment operational efficiency by enabling remote management and configuration of the system. This means that system adjustments, optimizations, and troubleshooting can be performed remotely, reducing the need for on-site intervention, and minimizing disruption to operations. Additionally, remote control allows for rapid response to changing operational requirements or conditions, ensuring that the system can adapt quickly and efficiently to evolving needs.