Degree project in Industrial Engineering

First cycle, 15 credits

Study on how data centers can actively contribute to the energy system

DIEGO VICEDO CANDELA

Bachelor of Science Thesis EGI-2025

TRITA-ITM-EX 2025:167

Study on how data centers can actively contribute to the energy system

Diego Vicedo Candela

Approved	Examiner	Supervisor
30/05/2025	Per Lundqvist	Viktor La Torre Rapp
	External supervisor	Contact person
	Kent Thorén	Diego Vicedo Candela

Sammanfattning

Med framväxten av teknologier såsom artificiell intelligens, blockkedjeteknik och maskininlärning har datacenter blivit en kritisk infrastruktur som stödjer högintensiv databehandling i global skala. Som ett resultat har de utvecklats till betydande energikonsumenter, och nya sätt undersöks för att göra deras drift mer hållbar och effektiv. Denna avhandling undersöker hur datacenter kan gå från att enbart vara energikonsumenter till att aktivt delta i energimarknaden genom återvinning och återanvändning av spillvärme samt tillhandahållande av tjänster för nätstabilisering. Genom analys av kvalitativ data insamlad från intervjuer och publicerad forskning besvarar uppsatsen specifika forskningsfrågor som leder till slutsatser om hur detta deltagande möjliggjorts. Avhandlingen baseras på två fallstudier av PDC:s och T-Loops sätt att bidra till energisystemet. Studien undersöker kylsystem, typer av nätstabiliseringstjänster, synergier med samarbetspartners samt affärsmodeller för att förstå hur ett aktivt deltagande i energimarknaden blev möjligt. En avslutande jämförelse av fallstudierna diskuterar båda företagens affärsmodeller, med fokus på återanvändning av värme och tjänster för nätstabilisering, för att ge tydliga slutsatser som underlag för diskussionen. Resultaten visar att befintliga implementationer redan påvisar genomförbarheten hos dessa teknologier, samtidigt som de förbättrar hållbarheten i verksamheten och skapar nya intäktsströmmar. Integrering av värmeåtervinning och nätstabiliseringstjänster i deras drift hjälper datacenter att uppfylla nya europeiska regler och ger dem ett konkurrensmässigt övertag gentemot sina konkurrenter genom att göra deras verksamhet mer miljöansvarig och effektiv.

Abstract

With the rise of technologies such as artificial intelligence, blockchain and machine learning, data centers have become critical infrastructure, supporting high-intensity computing on a global scale. As a result, they have emerged as significant energy consumers, and new ways are being explored to make their operation more sustainable and efficient. This thesis investigates how data centers can transition from merely being energy consumers to actively participating in the energy market through waste heat recovery and reutilization, and providing grid stabilisation services. Through the analysis of qualitative data gathered from interviews and published research, this paper answers specific research questions that lead to conclusions on how this participation was made possible. The thesis is based on two case studies conducted on PDC and T-Loop's way of contributing to the energy system. The study examines cooling systems, types of grid stabilisation services, synergies with partners, and business models to see how active participation in the energy market was made possible. A final case study comparison discusses both companies' business models, focusing on heat reuse and grid stabilisation services, to give clear conclusions on which to base the discussion. Findings indicate that existing implementations already demonstrate the feasibility of these technologies while improving sustainability in their operations and creating new revenue streams. Integrating heat recovery and grid stabilisation services into their operation is helping data centers to comply with new European regulations and giving them a competitive advantage over their competitors by making their operation more environmentally responsible and efficient.

Contents

Figures

Tables

Nomenclature

Abbreviations

4GDH	4th generation district heating
5GDHC	5th generation district heating and cooling
AH	Akademiska Hus
AI	Artificial Intelligence
BESS	Battery energy storage system
BRP	Balance response party
CDU	Coolant distribution unit
CPU	Central processing unit
CRAC	Computer room air conditioning
DC	Data center
DH	District heating
EECS	Electrical Engineering and Computer Science
GPU	Graphics processing unit
IoT	Internet of things
kW	Kilowatt
LTDH	Low-temperature district heating
MW	Megawatt
PDC	Parallel Computer Center
PUE	Power usage effectiveness
SEK	Swedish Krona
TC	Temperature coefficient
ULTDH	Ultra-low-temperature district heating
UPS	Uninterruptible power supply

1 Introduction

With the incremental digitalisation of society, data centers (DCs) are becoming increasingly critical infrastructure as they support cloud computing, artificial intelligence and online services. Providing these services leads to the consumption of vast amounts of electricity and the production of waste heat. Furthermore, data centers must invest in redundant systems to be able to provide uninterrupted operation. Most of these systems are not used during normal operation, but have to be in optimal condition in case a failure in the primary system may occur; redundancy is most needed in cooling and power systems. To make the growth of the data center sector sustainable, providers are looking into innovative solutions to reduce their environmental impact. These solutions include heat reuse systems and providing grid stabilisation services to utility providers, through Battery Energy Storage Systems (BESS).

Heat reuse allows the recovery and repurposing of the heat generated in the server halls; by integrating with district heating networks or other heat consumers, DCs can make revenue from a waste product whilst enhancing their energy efficiency. BESS, on the other hand, enable DCs to collaborate in grid stabilisation, mainly through ancillary services or arbitrage. These technologies open the door for DCs to go from purely being energy consumers to becoming active participants in the energy system.

While examples of data centers with implemented heat reuse and BESS technologies exist, their integration is still not widespread. Existing research and industry reports often focus on theoretical models or isolated case studies, leaving a gap in understanding how these solutions are instituted in the real world. It remains unclear how DCs can effectively collaborate with utility and district heating providers and what business models could be implemented in order to make these synergies possible. This thesis seeks to bridge that gap by analysing existing operational examples and identifying best practices for integrating heat reuse and BESS in DCs.

This research aims to contribute to the growing discourse of sustainable DC operations by mapping out existing solutions and analysing their effectiveness. The insights gained will help inform strategies for enhancing synergies between DCs and district heating providers and improving grid stability through optimised BESS integration. Ultimately, this thesis aims to provide a comprehensive understanding of how DC can transform from being energy-intensive facilities to playing an active role in energy production and distribution.

This thesis is structured as follows:

- **Chapter 2: Background** – Provides an overview of existing research on heat reuse, BESS, cooling systems, co-location, energy usage and district heating.
- **Chapter 3: Objectives and Scope** – Develops the aim and purpose of the thesis and establishes the limitations of the study.
- **Chapter 4: Methodology** – Explains the research approach, data sources, and case study selection.
- **Chapter 5: Case Studies** – Presents real-world examples of data centers implementing heat reuse and BESS.
- **Chapter 6: Conclusion and future work** – Summarises the research and analyses key findings, challenges, and future opportunities.

By following this structure, the paper aims to provide a clear, systematic, and well-supported analysis of sustainable energy solutions in the data center industry.

2 Background

2.1 Data centers

The need for data centres is continuously increasing as society continues to embrace AI and other intense computing activities such as IoT, cloud services, machine learning, and blockchain. Every interaction we make with the internet is processed through a data center (DC), searching on Google, uploading photos to social media, sending emails or streaming your favourite show. DCs are also where our data, images, videos and files are stored. They are a fundamental element in our lives, both at work and at home. For businesses, DCs are essential to keep their operations running; they provide services such as storage, communication, load planning, and, more recently, AI services. A data center is a facility that hosts IT infrastructure for one or multiple businesses.

The data center industry is continuously growing and adapting to the world's needs. Although there are endless types of DCs, most can be classified under four main categories: Enterprise, Hyperscale and Cloud, Colocation, and Edge (Taylor, 2023). Colocation DCs should not be mistaken for co-location of data centers, which will be discussed later in the thesis.

Enterprise DCs are built, run and owned by a company. Usually, they are located within the company's facilities; the servers or rack space are not shared with any other business, and all the costs, initial investment, maintenance and power expenditures are faced by the same company. They constitute the basis for the company's internal network and IT solutions; they mostly use a traditional local server deployment (Howell, 2024).

Hyperscale and cloud DCs offer computing power and data storage as a service. Here, the owner of the DC business is responsible for everything, which is the same as the enterprise variant, but gets revenue from renting the computer systems and resources. Companies can access their data and services through the internet instead of having it locally on their computer or server. This brings an obvious downside: if the connection fails or any of the auxiliary systems of the DC fail, the tenant can't do anything but wait until the service starts working again. This model is rising in popularity as many businesses don't have the investment ability nor the space to have their own DC. Moreover, it reduces costs as companies don't have to buy any hardware, and simplifies the operating process as maintenance is done by the hosting business, all while maintaining excellent flexibility when it comes to scalability because corporations don't have to worry about having excess or insufficient hardware power, they can just increase or decrease their computing needs following demand flows (Pradhan, 2015).

Colocation DCs are the middle ground between enterprise and hyperscale models. A third party owns the facility and rents the space to tenants who bring in their hardware. The colocation agreement usually includes physical security, a certain amount of power availability, internet connectivity and a cooling system from the facility owner's side. Businesses must only worry about purchasing and maintaining their hardware. This is an excellent solution for companies that don't have sufficient space for their computing needs but want to maintain control of their hardware. Costs are reduced not only by not having to own a dedicated space for the DC but also because the maintenance expenditure of the facility is divided among all the tenants, allowing for more competitive cooling and electricity prices (Yllmark, 2024).

Edge data centers are small-scale DCs that enable data to be processed and stored closer to the end user, minimising latency. Unlike traditional centralised data centers, which are often housed in massive facilities, edge data centers operate on a distributed model, enabling real-time data processing for applications requiring high-speed response times, such as mobile connectivity, smart cities, autonomous driving or healthcare. This type of DC is becoming increasingly vital as businesses and consumers demand faster, more reliable digital experiences. They play a crucial role in evolving digital technologies by enabling real-time processing and reducing dependency on centralised cloud infrastructures (Dowds & El-Saghir, 2021).

The need to create, store, and analyse data faster, in more significant volumes, and in a more accessible manner is driving data center innovation and investment. The need for DCs is growing, and so is dependency on them, with the growth and reliance on communication, IoT and AI services.

2.2 Energy usage

The DC business is an energy-intensive sector; even though significant advances have been made regarding energy efficiency, the computing needs of today's society have surpassed the industry's capability to produce more effective components. In 2022, DCs accounted for 1-1,3% of the world's energy consumption, utilising between 240-340 TWh of electricity (Hinterholzer & Hintemann, 2022).

Focusing on Sweden, DC establishments have significantly increased in recent years, as various factors make Sweden an excellent host to this industry. Sweden currently has 91 data centers, most of them in Stockholm, with 42 DCs in its metropolitan area (Map, 2025). Recent studies (Jerléus et al., 2024) show that DCs represent 4.75% of the country's total electricity usage. Although, as I said, great efforts are being made to make data processing more efficient, if the DC industry wants to thrive towards a more sustainable model, greater innovations outside of computing efficiency must be made.

Data centers commonly use the power usage effectiveness (PUE) metric to describe how much of the total energy consumed by the DC is actually consumed by the IT equipment and how much is used by auxiliary systems such as cooling, UPS or transformers. Total energy consumption includes IT equipment usage, cooling system usage, and power distribution losses. Therefore, as PUE values get closer to 1, the DC is utilising fewer resources to power auxiliary systems that don't make any revenue and is maximising the power dedicated to computing and data processing, making for a more efficient system. Nevertheless, it has to be taken into account that PUE doesn't measure how efficient the IT equipment is; it only shows the fraction of total power dedicated to the components of the DC that are not IT-related, as shown in (1). PUE values vary between different facilities; legacy facilities with no optimisation on cooling or power systems can achieve values close to 2, while a low-energy efficiency-focused design can achieve metrics of 1.2 or lower (Al Kez et al., 2025).

$$\text{PUE}=\frac{\text{Total Energy Consumption}}{\text{IT Energy Consumption}} \qquad (1)$$

The hyperscale and cloud DCs are helping the field as they can obtain lower PUE metrics, meaning less power is lost in cooling and infrastructure. Hyperscale DCs are custom-built for performance and energy savings; they use custom hardware, liquid cooling and energy management systems to optimise power usage. Albeit having higher energy consumption due to their immense size and computing capability (Rezaei-Mayahi, 2019), these DCs tend to achieve the lowest PUE, with examples such as Google's with a PUE of 1,1; this means that only 10% of the power used by the facility is wasted on cooling and infrastructure (Google, 2024).

With enterprise DCs, the PUE varies depending on the size of the computer and the company's willingness to invest in newer and more efficient cooling methods, with larger DCs tending to achieve lower PUEs (Lei & Masanet, 2022).

Colocation DCs usually face challenges in achieving optimal PUE due to the variability of customer needs. This type of DC must design its facilities and cooling systems accounting for maximum usage of its server halls, but this rarely happens. Customer needs and load curves are diverse and challenging to predict, making energy optimisation arduous. Furthermore, these centers require extra redundancy to meet uptime and cooling requirements. Underutilisation is also a difficulty these centers face; clients may rent more space than needed, leading to wasted power (Aligned Data Centers, 2018).

To increase energy efficiency in DCs, the industry is focusing on heat reuse and BESS to improve efficiency and sustainability in their facilities (Rehan et al., 2025).

2.3 Heat reuse

Data centers transform electricity into processing power and thermal energy. Heat reuse involves capturing and repurposing the excess thermal energy produced by data centers. This process consists of a heat recovery system that extracts the heat produced by servers and redirects it to other applications. The feasibility of this process depends on different factors such as the operating temperatures of computers, the availability of facilities or processes to absorb the heat and the efficiency of cooling systems. Traditionally, this heat had no value and wasn't recovered, but great efforts are being made to increase the sustainability and energy efficiency of DCs. Sweden leads globally in excess heat recovery for heat supply, with a 9% share, followed by Russia at 6%, and France, Finland, and Denmark, each at approximately 2% (Lygnerud et al., 2022).

Higher computational densities on chips and the evolution of cooling systems towards liquid and water cooling, combined with better thermal management solutions, enable DCs to run at higher temperatures (ASHRAE, 2016). As shown in Figure 1, the server's capability to produce heat from the same input temperature has increased by almost 10ºC in the last 12 years; higher outlet temperatures open a wide range of opportunities for heat to be recovered and reused.

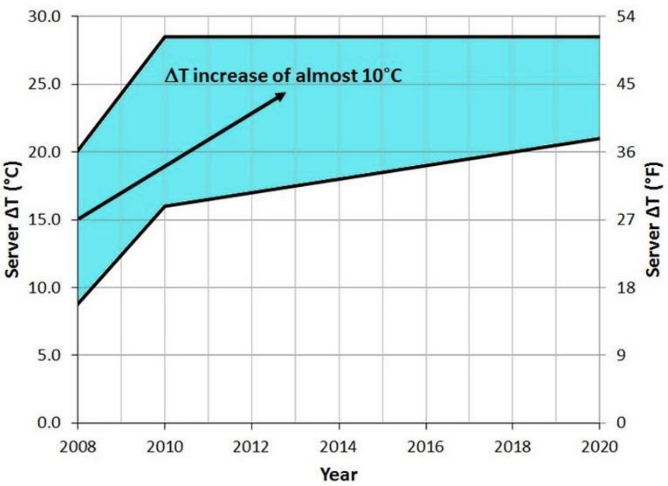

Figure 1 *Server ΔT projection at 25°C server inlet temperature (ASHRAE, 2016)*

2.3.1 District heating and cooling

Sweden has a well-established and extensive district heating system, which plays a significant role in the country's energy landscape. As of 2020, district heating accounted for approximately 58% of Sweden's total energy consumption in residential and non-residential buildings (SBBA, 2021). Nearly all cities and towns in Sweden have adopted district heating systems to provide heating for apartment buildings, single-family homes, and industrial facilities. If we focus on Stockholm, the leading district heating company, Stockholm Exergi, has some 3000 km of piping for its district heating network and around 300 km for district cooling (Stockholm Exergi, 2025).

Traditional district heating networks rely on heat generation plants from which the heat is then distributed to customers. The introduction of 4th generation district heating (4GDH) aimed to incorporate more renewable heat generation processes and low-temperature waste sources into the district heating network (Lund et al., 2014); thus, 4GDH technologies were developed with decarbonisation as the primary goal. 4GDH includes low-temperature district heating (LTDH) and ultra-low-temperature district heating (ULTDH). LTDH works with temperatures around 60°C, which is enough for a direct supply of domestic hot water. In comparison, ULTDH networks operate near 40°C and are designated for a direct supply of space heating (Ommen et al., 2016). Lower grid temperatures reduce temperature differences between the grid and the outside, thus reducing losses (Pans et al., 2023). Temperature levels vary throughout the grid and year; Stockholm Exergi's grid is capable of withstanding temperatures between 5°C and 120°C in the DH network and between 2 °C and 30 °C in the cooling network. Temperatures in the supply line must be at least 68 °C with 0 °C outdoor temperature, linearly rising to 103 °C at -17 °C outdoor temperature. Supply temperatures must never exceed 120 °C and must be below 100 °C if outside temperatures exceed 0°C (Stockholm Exergi, 2025).

Hot water consumption and heat needs are not constant throughout the day and vary with outside temperature and position on the grid, which makes return temperature in DH networks variable. Figure 2 below shows typical values for how the total return temperature in a part of the district heating network varies as a function of outdoor temperature.

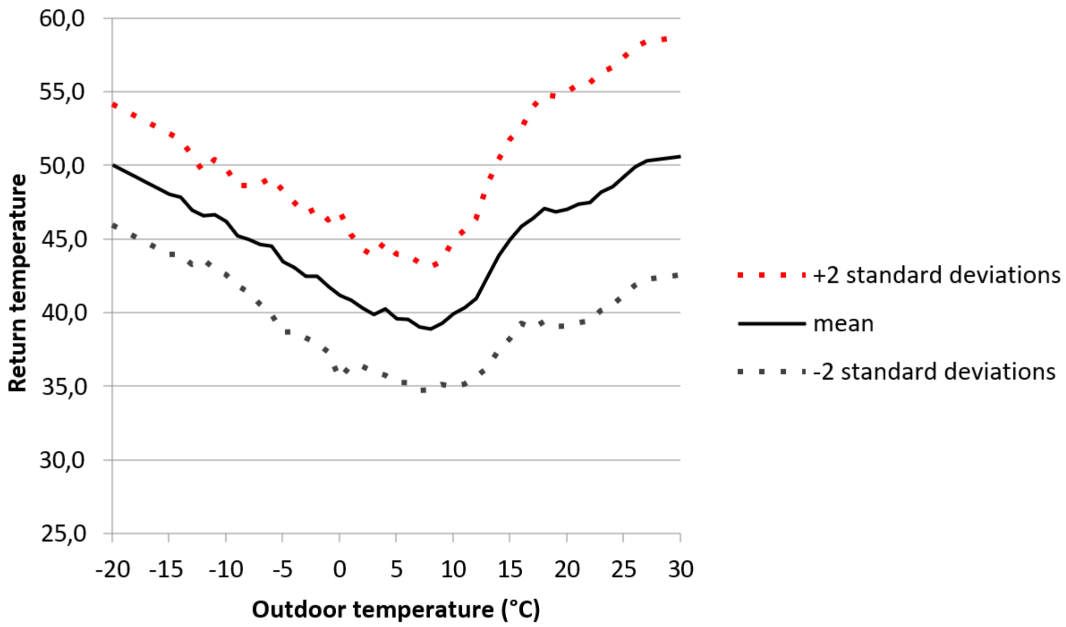

Figure 2 *Variation of return temperature as a function of outdoor temperature for the district heating network (Stockholm Exergi, 2025)*

In recent years, 5th generation district heating and cooling (5GDHC) appeared as a complement to 4GDH; it is described as a district heating grid operating at near-ground temperatures using a bidirectional exchange of heat and cold between connected buildings, facilitated by seasonal thermal storage (Boesten et al., 2019). Reduced grid temperatures require heat pumps at the connected facilities to achieve proper temperatures for domestic hot water. Figure 2 shows the main differences between 4GDH and 5GDHC: lower grid temperatures achieved by 5GDHC and the aim to encourage consumers to become 'prosumers'.

A 'prosumer' is an energy consumer who produces and puts energy into the grid. Prosumers can return excess energy to the grid or share it with a district heating system, contributing to a more decentralised and sustainable energy network. DC are a great example of prosumers in the district heating grid as they can buy cooling and sell excess heat to the network.

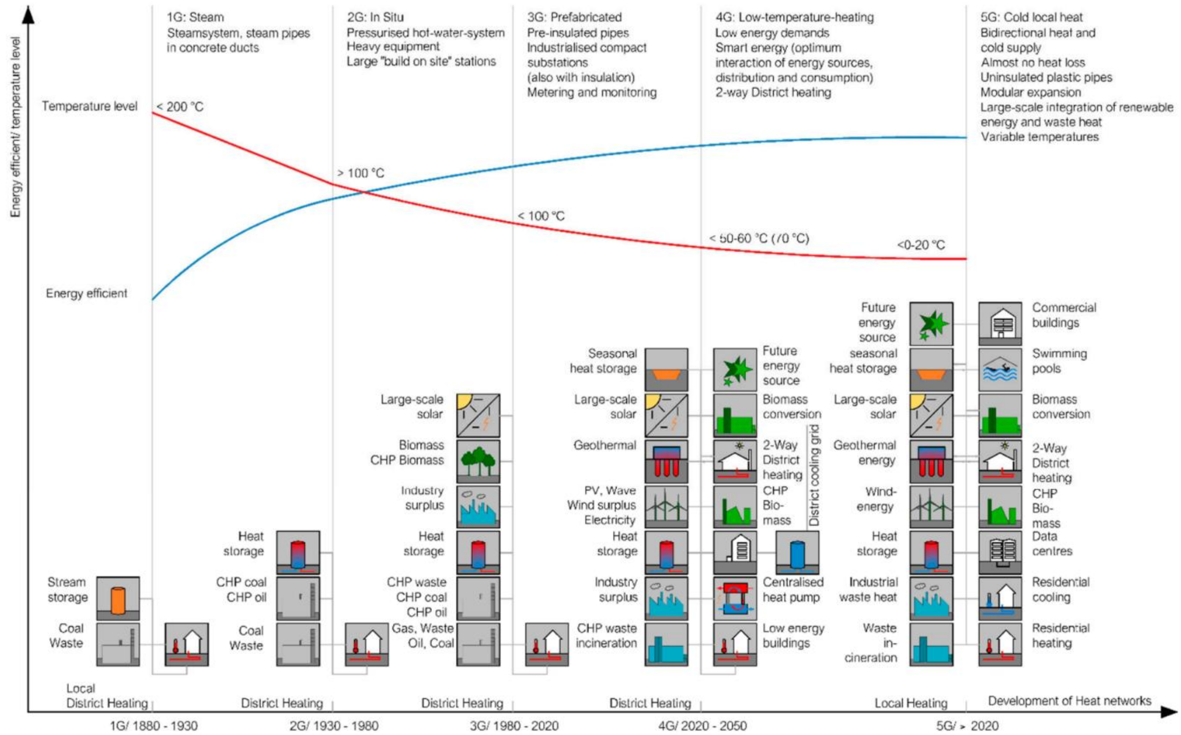

Figure 3 *District heating networks from the first to the 5th generation over the last 150 years (Zeh et al, 2021).*

2.3.1.1 Open district heating

Various studies show the possibility of connecting heat producers such as DC and supermarkets to the district heating network (Zülhsdorf et al., 2018) (Hyvönen et al., 2024). The DC grid provides a reliable source to expel the excess heat generated, as it constantly needs heat all year round. Other heat recipients, such as greenhouses and industrial facilities, tend to have more variable heat necessities and thus may not be able to absorb the required amount of heat reliably.

Forthcoming technologies and initiatives such as Open District Heating aim to bring heat production closer to customers by enabling facilities that generate excess heat, like DCs, supermarkets or ice-skating rings, to introduce their excess heat into the network. This new view of district heating is helping to decentralise DH production, enhancing the grid's stability as DH providers don't have to rely solely on their production plants. Furthermore, suppose the heat source is located close enough to heat consumers. In that case, the supply temperature can decrease as the heat is produced closer to the recipient, so losses are minimised. In addition, this enables heat producers to create a new revenue stream. Indicatively, Stockholm Exergi pays 2 million SEK annually for a heat delivery equivalent to 1 MW (Stockholm Exergi, 2025).

There are various models for connecting to the open district heating system. Stockholm Exergi distributes an instruction manual for connecting to the open district heating system. Taking a look at forms of delivery, this manual offers two.

Option 1: Delivery to the district heating supply line.

A heat exchanger inputs low-grade heat from the district heating return line and outputs high-temperature heat to the district heating supply line. The delivery must meet the temperature level that Stockholm Exergi guarantees to its district heating customers. However, exemptions can be made if an agreement is made between Stockholm Exergi and the heat supplier.

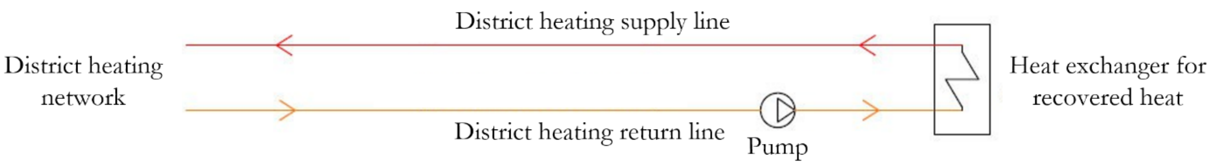

Figure 4 *Delivery to the district heating supply line (Stockholm Exergi, 2025).*

Option 2: Delivery to the district heating supply and return line.

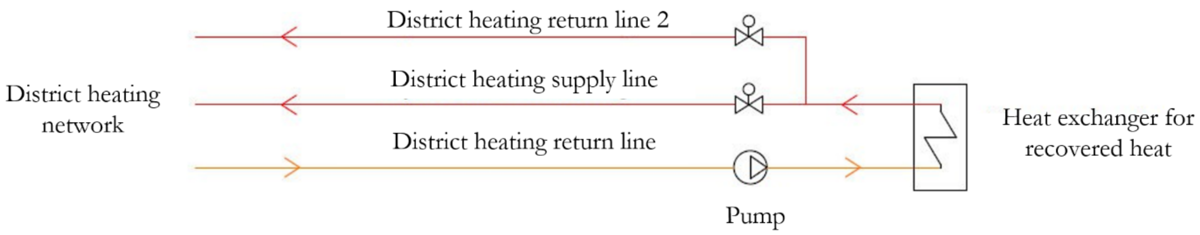

Figure 5 *Delivery to the district heating supply and return line (Stockholm Exergi, 2025).*

In this option, a connection is made to both the supply and return line of the district heating network, thus allowing for heat delivery at different temperatures; a valve is placed at the beginning of the connection to direct the flow to the correct line. A large enough return flow must be achieved at the delivery point in order for heat delivery into the return line to be possible.

Delivery to the supply line of district cooling is also available, but for DCs, that option does not apply as DCs output temperatures are too high for that line. Open District Heating offers multiple connection methods for heating and cooling supply to the network. Selecting one of them depends on factors such as the amount of thermal energy that wants to be supplied, its distribution over the year, the market price of the thermal energy and geographical location. Delivery to the supply line holds the highest value and is least affected by the location of the network connection as long as the required temperature level is maintained.

2.3.2 Other heat recipients

Outside of DH, a direct connection between heat producers and consumers can be made. Research shows that connections between industries can be made in order to utilise waste heat efficiently. (Yang et al., 2024) looked at how industrial waste heat could be used in agriculture to reduce carbon emissions and reliance on fossil fuels. (Oró et al., 2018) proposed the design and economic analysis for reusing heat from a liquid-cooled DC in an indoor swimming pool; this paper shows a reduction in OPEX both in the DC and swimming pool side.

2.4 Cooling systems

IT hardware converts a supply of electricity into heat. In order to keep the hardware running in optimal condition, this heat must be captured and removed from the server halls. This is done by the cooling system. A cooling system extracts warm air from the server halls and, through various processes, depending on the technology used, extracts heat from the air, cooling the inlet air down and heating the outlet air, usually the atmosphere. Cold air is then returned to the DC, and the process starts again.

Continuous evolutions in computing density make IT equipment generate more significant amounts of heat, making cooling an even more critical aspect. IT hardware has been traditionally air-cooled, but that trend is changing as more compact CPUs and GPUs appear. This new generation of hardware components require higher thermal transfer rates to keep them operating at peak performance, which has opened the door to liquid cooling systems. Both these systems are the standard for DC cooling (D. Thorén & O. C. Widell, 2011).

2.4.1 Air cooling

Air-cooled systems are amongst the most common cooling solutions for DCs. Two traditional air-cooling systems are Computer Room Air Conditioning (CRAC) and Computer Room Air Handling (CRAH). These systems ensure that temperature, humidity and airflow remain stable during the regular operation of the DC so that computers can work in a stable environment and maintain high efficiency (Ali Habibi & K. Halgamuge, 2017).

The American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) publishes the Equipment Thermal Guidelines for Data Processing Environments. These guidelines include the recommended temperatures at which data centers should be run, depending on their class. Classes range from A1 to A4, and a special class H1 is reserved for supercomputers. Typical applications and running temperatures of data centers can be seen in Table 1.

As seen in Table 1, with higher classes, for example, A3 and A4, inlet temperatures go as high as 45ºC; such high temperatures enable heat reuse without the need for additional heat pumps. DC operators who employ heat recovery systems try to push these temperatures as high as possible to increase the quality of the recovered thermal energy.

Table 1 *2021 Thermal Guidelines for Air Cooling (ASHRAE, 2024).*

Class	Typical Application	Dry Bulb Temperature (°C)
A1	Enterprise-class servers, storage, and networking equipment (e.g., data centers)	15 - 32°C
A2	Less controlled environments (e.g., small server rooms)	10 - 35°C
A3	IT equipment designed for higher ambient temperatures (e.g., edge computing)	5 - 40°C
A4	Harsh environments (e.g., industrial IT, telecom enclosures)	5 - 45°C
H1	High-performance computing (HPC) environments	5 - 40°C

2.4.1.1 Computer Room Air Conditioning (CRAC).

A CRAC unit is a device that monitors and maintains temperature, humidity, and air distribution in a DC. A CRAC unit operates similarly to a standard air conditioning unit but achieves better air filtration, humidity control and airflow. In the refrigeration cycle, air passes through a cooling coil filled with refrigerant, which is kept cold using a compressor. Newer CRAC units enable airflow control, while older ones can only turn on and off, thus achieving lower energy and cooling efficiency (Nada & Said, 2017).

CRAC systems typically dispense cold air through a raised floor. The cold air flows underneath the floor and rises via perforated sections to reach the servers. Then, it enters the server racks, absorbs the computer's heat, and exits through the backside as hot air. Server racks are positioned to form aisles and aligned so that the hot and cold sides don't face each other, as seen in Figure 7. To allow for increased heat recovery and to avoid the mixing of hot and cold air, hot isles can be contained in a closed cage and connected to the return side of the CRAC unit; this can also be done the opposite way, having the cold aisle contained and connected to the inlet of the CRAC unit through the perforated raised floor, as seen in Figure 6.

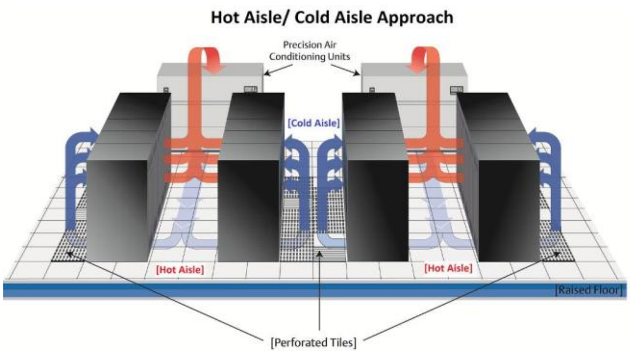

***Figure 7** Air-cooled server hall (Nadjahi et al., 2019)*

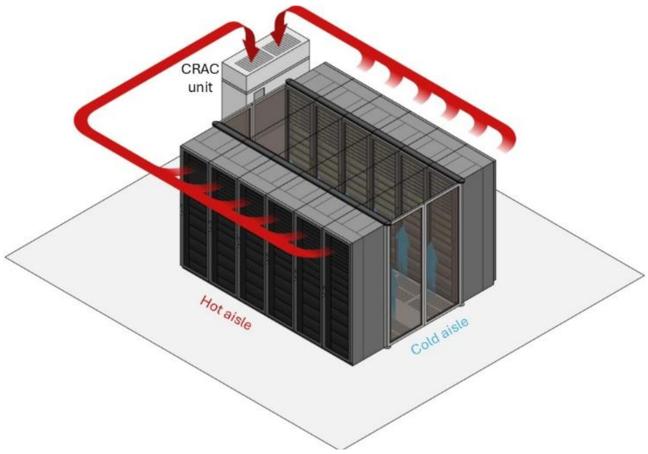

***Figure 6** Cold aisle containment.*

2.4.1.2 Computer Room Air Handling (CRAH)

A CRAH unit cools the air using a cooling coil filled with chilled water. A chilled water plant, also known as a chiller, supplies the water to the unit. This differs from CRAC units, as these use mechanical refrigeration. CRAH units can control fan speed to optimise temperature and humidity levels. Alignment and positioning of the server hall can be done the same way as with CRAC units, with containment options also available with this technology. CRAH units are most efficient when placed in data centers with electrical loads higher than 200 kW. Although the suppression of the compressor reduces the energy usage when compared to CRAC units, CRAH installations are more complex due to the piping that needs to be placed to connect the chiller with the rest of the systems. They also need more space, as part of the floor plan has to be dedicated to fit the chiller (Salih Erden, 2022).

2.4.2 Liquid cooling

With increasingly computationally dense DCs, more efficient cooling methods must be implemented. Liquid-cooled systems have become an increasingly popular solution as they can manage higher heat loads than traditional air-cooled systems while maintaining superior efficiency. Water cooling relies on liquid as a heat transfer medium, allowing for more effective temperature regulation and potential heat reuse. However, they require more complex infrastructure planning to prevent leaks and regular maintenance to ensure reliability. Furthermore, liquid cooling reduces the space needs in DCs as higher-density racks can be used, with direct-to-chip solutions enabling up to 100 kW per rack. (Cool IT systems, 2023).

Liquid-cooled solutions can offer a quieter environment. When compared with air-cooled systems, liquid cooling is predominantly silent, as air doesn't have to be recirculated. The pumps driving the fluids through the circuit are the only noise producers, creating a much quieter atmosphere than the one seen in air-cooled DCs. Many forms of liquid-cooled systems have been and are being developed at the moment; these new cooling solutions aim to help the expansion in power densities that the DC industry is facing. Above all, two major cooling solutions have to be outlined, and they are direct-to-chip and immersion cooling.

Same as with air-cooled systems, ASHRAE has thermal guidelines for water-cooled systems. In this case, classes start with a "W" followed by a number (17, 27, 32, 40, 45, +). This number indicates the maximum temperature at which the supply water enters the system, with a minimum water temperature of 2ºC for all clases. An extract of these recommendations can be seen in Table 2, where classes are paired with their usual cooling technologies and inlet water temperatures.

Table 2 *2021 Thermal Guidelines for Liquid Cooling (ASHRAE, 2024).*

<table>
<tr><th>Class</th><th>Cooling system</th><th>Facility Water Supply Temperature (°C) [a]</th></tr>
<tr><td>W17</td><td rowspan="2">Typically data centers that are traditionally cooled using chillers and a cooling tower, but with an optional water-side economiser to improve energy efficiency, depending on the location of the data center</td><td>17°C</td></tr>
<tr><td>W27</td><td>27°C</td></tr>
<tr><td>W32</td><td rowspan="2">Typically operated without chillers in most locations; however, some locations may still require chillers</td><td>32°C</td></tr>
<tr><td>W40</td><td>40°C</td></tr>
<tr><td>W45</td><td rowspan="2">Typically operated without chillers to maximise energy efficiency and reduce capital expense. Some locations may not be suitable for dry coolers</td><td>45°C</td></tr>
<tr><td>W+</td><td>>45°C</td></tr>
</table>

2.4.2.1 Immersion cooling

Immersion cooling offers an innovative alternative to DC cooling; in these systems, IT equipment is submerged in a dielectric fluid. A dielectric fluid is an electrically insulating fluid that does not conduct electricity and thus can be used for cooling and insulation in electronic systems. Common dielectric fluids include synthetic oils and specially formulated coolants. Having the hardware submerged greatly enhances heat transfer capacity, as liquids have a greater specific heat capacity than air, and the servers are bathed in a continuous flow of coolant.

To implement this solution, three main parts have to be put in place: an enclosure for the servers or tank, a coolant distribution unit (CDU) and a final heat rejector. The cooling process is very simple: servers are submerged in a cooling fluid which is constantly recirculated; this fluid heats up as it absorbs the server's heat. To remove the heat from the coolant, it goes to the CDU, which is a liquid-to-liquid heat exchanger. In the CDU, the coolant relinquishes the heat to another liquid whose temperature is much lower; that liquid is usually water. Finally, the water goes into the final heat rejector, usually an evaporative cooling tower, a dry cooler or an existing chilled water loop, where the heated-up water is cooled down and is ready to go back to the CDU. To summarise, a typical immersion cooling system comprises two hydraulic loops: one between the servers and the CDU, the one done by the coolant, and another one between the CDU and the final heat rejector, typically run by water. This is exemplified in Figure 7 below.

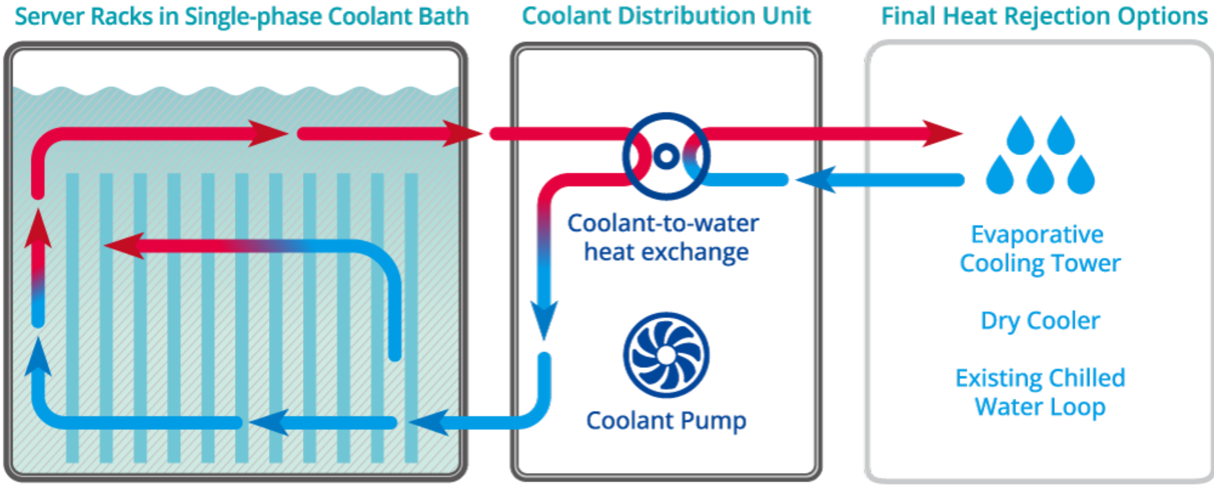

Figure 8 *Immersion cooling loop (Gigabyte, 2023)*

Having the servers submerged provides consistent and precise cooling, preventing hot spots, thermal shutdowns and increasing the hardware's lifespan, as heat transfer is evenly distributed throughout the hardware. In addition, server racks are easily accessible, with server-level components such as CPUs and GPUs being easily accessible. Regarding power usage effectiveness (PUE), immersion-cooled DCs can achieve values of 1.05-1.1 (Mann, 2023). Not only is energy saved through immersion cooling, but space is, too. According to a research article (Haghshenas et al., 2023), immersion-cooled DCs only require about one-third of the space to that of an air-cooled solution. All in all, immersion cooling is set to be one of the biggest trends when it comes to DC cooling, as it is bringing energy and space needs down while increasing power density on the racks and maintaining great scalability and serviceability.

2.5 Battery Energy Storage Systems (BESS)

Data centers include BESS as an uninterruptible power supply (UPS). An UPS is a critical system in DC operations. It ensures the DC can keep running despite power failures, voltage fluctuations, or grid disturbances. Traditionally, UPS systems in data centers rely on lead-acid batteries or flywheels to supply short-term backup power until diesel generators take over. However, with advancements in Battery Energy Storage Systems (BESS), Lithium-ion (Li-ion) batteries are becoming the preferred choice for modern UPS solutions. Lithium-ion batteries provide higher energy efficiency, longer backup duration, lower maintenance costs, smaller space requirements, longer lifespan and less environmental impact than lead-acid batteries (Keshan et al., 2016). Furthermore, this enables Lithium-ion batteries to participate in frequency regulation, ancillary services and grid arbitrage, bringing new income streams to DCs.

BESS enhance the reliability, efficiency and sustainability of data center operations. Unlike conventional UPS solutions, BESS can provide instant power, with milliseconds of response time, support longer backup duration when compared with lead-acid batteries, improve energy efficiency by participating in grid services like frequency regulation, reduce reliance on fossil fuel generators, lowering the DC's carbon footprint and lower operational costs through peak shaving and demand response.

As data centers aim for zero-carbon operations, BESS are expected to replace traditional UPS systems. The integration of renewable energy, AI-driven energy management, and grid services will make BESS a core component of next-generation data centers.

2.5.1 Grid stabilisation

The rise in renewable energy integration into the energy grid proposes a challenge when it comes to grid stability. Renewable power sources like wind and solar are dependent on weather conditions and thus rarely align with energy demands. The integration of battery energy storage systems (BESS) in the grid aims to reduce the instability of these power sources. BESS can help the electrical grid by energy storage and time-shifting, regulation of frequency, voltage support, and enhancement of grid reliability (Raj Suyambu & Kumar Vishwakarma, 2024). The European power grid is designed to work at 50 Hz; when production and demand don't match, the frequency on the grid changes, as seen in Figure 8. For the grid to work correctly, the frequency should be kept within 50 ± 0.1 Hz. Frequencies within this interval are referred to as the "normal frequency band", and frequencies outside of it are referred to as "frequency deviations" (Aasvestad Sandvik, 2024).

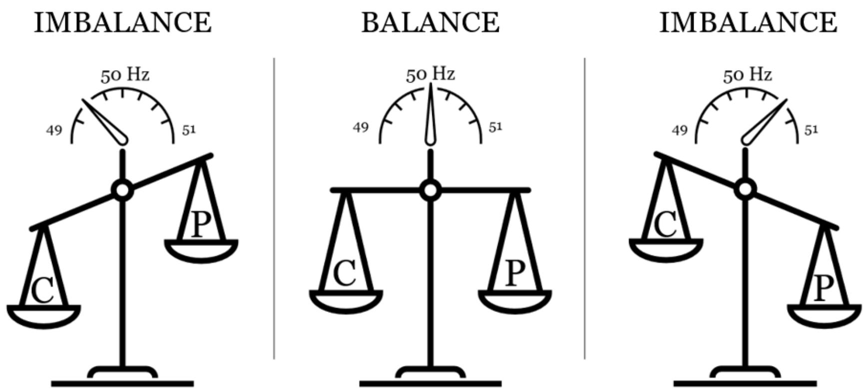

Figure 9 *Imbalance between consumption (C) and production (P) on the grid frequency (Aasvestad Sandvik, 2024)*

BESS enable the storage of electricity when production is high but demand remains low, and then it injects it into the grid when demand rises. This is called arbitrage. It helps both renewable and traditional electricity generation methods, as more production from renewable power sources can be effectively utilised and conventional methods can maintain a continuous output, enhancing efficiency. Another approach can be made by using BESS as an ancillary service, meaning that the main goal is to improve the quality of grid electricity. This is achieved by ensuring the grid's frequency remains within a specific range, maintaining the voltage levels within the grid, providing backup power in case of sudden demand spikes or generation shortfall, and even helping to restart the grid after a complete shutdown.

Svenska Kraftnät, Sweden's largest utility operator, outsources its ancillary services to qualified participants in the electricity market. Svenska Kraftnät requires the following ancillary services in order to balance and manage disturbances in the power system: FFR, FCR-D upward, FCR-D downward, FCR-N, aFFR and mFRR (Svenska kraftnät, 2025).

Fast Frequency Reserve (FFR)

FFR provides an ultra-fast response to stabilise frequency in case of sudden power imbalances, such as generator failures or rapid load changes. Should that happen, FFR is automatically activated within milliseconds to a few seconds to prevent frequency from dropping too low before other reserves can react. An FFR usually stays active for 5-30 seconds until Frequency Containment Reserves (FCR) take over. Traditional FFR providers include BESS, supercapacitors and hydropower. The Swedish grid can require up to 105 MW of FFR.

Frequency Containment Reserve for Normal Operation (FCR-N)

FCR-N handles minor frequency deviations, ± 0.1 Hz of nominal frequency. It is a continuous reserve that operates both upward (increase generation) and downward (reduce generation) to maintain a stable frequency. This reserve is activated within 30 seconds and must endure for 1 hour. Standard providers for FCR-N are hydropower, demand response, and fast-reacting power plants. The Swedish grid requires 224 MW of FCR-N.

Frequency Containment Reserve for Disturbances (FCR-D)

FCR-D is activated during more significant frequency deviations beyond a certain threshold. It is divided into FCR-D upward and downward. FCR-D upward is automatically and linearly activated when the frequency drops to the 49.9 – 49.5 Hz interval. It is used to increase power generation or reduce demand. FCR-D downward is automatically and linearly activated when the frequency rises to the 50.1-50.5 Hz interval. It is utilised to reduce power generation or increase demand. Both FCR-D upward and downward should be activated within 30 seconds of the disturbance and endure for at least 20 minutes. The Swedish grid requires a volume of up to 542 MW of both FCR-D upward and downward.

Automatic Frequency Restoration Reserve (aFRR)

aFRR is used to restore the frequency to its nominal value, 50 Hz, in Europe. Unlike FCR, which stabilises frequency, aFRR corrects deviations over a more extended period. aFRR operates via an automatic control system that adjusts power based on real-time frequency data. It is fully activated within 5 minutes and has to endure for 1 hour. Typical providers for this ancillary service can be flexible generation plants such as hydropower and gas turbines or demand-side response. The Swedish grid requires up to 111 MW of aFRR.

Manual Frequency Restoration Reserve (mFRR)

mFRR is used for long-term frequency balancing after an event when other reserves are insufficient to restore system balance. It is a manual reserve, meaning operators activate it based on system conditions. It has an activation time of 15 minutes and must last for 1 hour. Common providers for this ancillary service can be large generation plants such as coal, gas and hydro or demand-side response. The Swedish grid requires up to 890 MW of mFRR.

BESS are highly suitable for fast and flexible frequency control because of their ability to charge and discharge power rapidly. Although research shows that BESS, especially in Sweden, are best suited and more profitable when used for FCR, FFR and aFRR (Ehrlin, 2024) (Kakabadze, 2024) (Andersen, 2024). mFRR is less suited for BESS as it requires longer activation periods, which degrades and puts excessive strain on the batteries. BESS owners can procure a spot in the ancillary services market and then get paid as cleared. Prices paid for the energy cleared vary throughout the day and service given.

2.6 Co-location

Co-location means placing multiple entities in a single location to share resources and reduce costs. Applying this to a data center that implements heat reuse, co-location aims to place the DC near heat-intense activities or facilities such as greenhouses, swimming pools or the district heating and cooling network so that the recovered heat from the server halls can be repurposed. Smaller edge or enterprise DCs can be placed very near or directly inside the heat recipient's facility. At the same time, larger DCs, such as hyperscale and colocation, have more significant space needs, thus being usually located on the outskirts, leaving fewer possibilities for heat reuse. Nevertheless, Stockholm's vast district heating and cooling network opens an excellent opportunity for this heat to be repurposed, as they require high heat inputs all year round and with initiatives such as Open District Heating, which are open to receiving excess heat from multiple businesses, especially DCs (Stockholm Exergi, 2025).

In addition to heat reuse, co-location strategies can involve energy storage solutions, such as utilising the DC's battery UPS for grid stabilisation. By placing the DC close to a fluctuating grid area, their batteries can contribute to the grid when power needs from other consumers are high and recharge when grid usage is low. This enhances grid resilience and helps integrate renewable power sources such as wind and solar, which are dependent on weather conditions, into the electric power system (Lamp & Samano, 2022).

Co-location can help DC reach superior sustainability levels and create new economic opportunities for both DC operators and local energy providers by simply repurposing what was once rejected. This approach transforms DC from isolated energy consumers into active contributors to a more sustainable energy network (Huang et al., 2020).

2.7 Business models

Changes must be made in the traditional data center business model in order to implement heat reuse and grid stabilisation services into their operation. Integrating heat reuse and grid stabilisation brings new revenue streams to data centers. The business model concept has gained relevance in the academic literature and managerial practice over the past two decades. At its core, a business model articulates how an organisation creates, delivers, and captures value (Teece, 2010). It describes the conceptual framework of how an enterprise operates, including its value proposition, target customer segments, revenue streams, cost structure, and key resources and activities.

Alexander Osterwalder brought a revolutionary look to business models. With his PhD work (Osterwalder, 2004), Osterwalder offers a systematic and visually intuitive method for describing and analysing business models. Breaking them into nine building blocks: customer segments, value propositions, channels, customer relationships, revenue streams, key resources, key activities, key partnerships and cost structure. The structure created by Osterwalder represents the tool to create a rigorous, structured examination of how organisations operate and create value within their respective environments.

On the other hand, synergies must be achieved between data centers, heat receivers and utility providers to enable data centers to contribute to the energy system. A synergy between two companies is the concept by which the value of the whole is greater than the additive value of the individual parts (Tanriverdi, 2006). The book *Strategy, Process, Content, Context* (De Wit & Meyer, 2004) dives into the possible synergies that may appear between companies. The most interesting synergy to the topic of this thesis is the **Synergy by integrating activities**, which is defined in the book as:

"Two or more businesses are related if an integration of their value chains is more efficient and/or more effective than if they were totally separated" (De Wit & Meyer, 2004)

This is then divided into two possible ways of achieving this: sharing value-adding resources and linking value-adding activities. Sharing value-adding resources doesn't align with what is being studied as resources are not shared but reused and repurposed. Linking value-adding activities does explain the relationship between data center providers, heat receivers and utility providers; it is defined in the literature as:

"Business units that are not horizontally but vertically related can have internal customer-supplier relationships" (De Wit & Meyer, 2004)

These vertical relationships are clear in the case of data centers, heat receivers and utility providers, as heat receivers and utility providers can act both as suppliers and customers.

3 Objectives and Scope

3.1 Objectives

DC are a fundamental asset to today's society; as their involvement in everyday life continues to grow, so does their energy consumption. DC operators are striving to make their centers more sustainable and energy efficient. This project aims to map current solutions used by DC operators to make their operations more sustainable while creating new revenue streams by repurposing excess heat or providing ancillary services to utility companies. These new revenue streams will be analysed through the understanding of the business models of the different DC providers.

The goal is to analyse how DC can go from energy consumers to energy providers by looking at working examples who have implemented heat reuse and BESS for grid stabilisation. This report pursues to assess how heat reuse systems work in DCs and what synergies can be made between district heating providers or other possible heat recipients and DCs. Furthermore, this paper looks at how DCs can collaborate with utility providers to enhance grid stabilisation and efficiency. These objectives are to be fulfilled through answering the following research questions:

- How is heat recovery done?
- What services are provided to contribute to grid stabilisation?
- What synergies are made between data centers, heat receivers and utility providers?
- What business models are put in place to make these technologies profitable?

These questions aim to give a qualitative background to answer the motivating question of the thesis which is:

- **Can data centers become active contributors to the energy grid through waste heat recovery and grid stabilisation?**

3.2 Scope

This paper focuses on fully operational data center models currently in use at the time of writing, so that real data can be retrieved to analyse their operation. By concentrating on Stockholm, synergies between district heating suppliers and DC can be analysed and some centers could be visited, thus gaining further insights. Stockholm has both a well-established district heating network and clearly defined seasonal temperature variations, making it an ideal location for this study.

4 Method

This section explains the process followed to get data and insights into how data centers contribute to the energy system through heat reuse and BESS for grid stabilisation. This study adopts a qualitative research approach, supplemented by case study analysis and expert interviews.

Given the complexity of integrating heat reuse and Battery Energy Storage Systems (BESS) in data centers, the data obtained in the background follows a qualitative research approach for a deeper understanding of industry trends, operational challenges, and best practices. Additionally, the case studies that will be conducted later in the thesis give a quantitative approach to that background by examining real-world implementations of heat reuse and BESS in data centers.

Various companies were contacted for data collection and insights on the data center and district heating sectors. These companies have proven knowledge and experience in either heat reuse, BESS or district heating systems. To gain further insights, semi-structured interviews were conducted with professionals involved in the data center and district heating industries. The interviewees who have collaborated with the project are:

Luca Manzari – PDC's Systems Manager.

Johan Dalgren – R&D on Distribution and Storage at Stockholm Exergi and PhD student at KTH.

Staffan Styme – Deputy CEO of T-Loop.

Ellinor Grape – Business Development Manager at SDC.

All interviews were transcribed and can be found in the appendix of the thesis. Questions were proposed to answer the research questions stated in the objectives, giving context to the companies, as well as providing information to extend and verify the research done in the background. To do so, the interviews are semi-structured and divided into five main fields: General/Company Context, Economics, Heat Reuse, Grid Stabilisation and Co-Location.

The companies and entities that collaborated with the project are: the Center for High Performance Computing (PDC) from KTH, Scandinavian Data Centers (SDC), T-Loop and Stockholm Exergi. Information was gathered from these companies through interviews and, in some cases, through visits to their centers. Although SDC participated in the data gathering process, and a company representative was interviewed, their data center was not included as a case study, as it was not finished at the time of writing and was in the process of being constructed. Nevertheless, the information provided helped construct the background and define clear objectives for the thesis.

This information was used for the writing of the background chapter as well as for obtaining the necessary data for the case studies. Other sources of data were journal publications, theses, industry reports and technical specifications. All data used in the thesis were analysed using a thematic analysis approach. Thematic analysis was described in *Using thematic analysis in psychology. Qualitative Research in Psychology, as,*

"Thematic analysis is a method for identifying, analysing, and reporting patterns (themes) within

Data" (Braun & Clarke, 2006)

Following that method, interview questions were divided into themes, as stated above, and the theoretical research was organised in a similar way, as seen in Chapter 2. This made a clear structure on which to base the case studies. Business models will be analysed, on one side focusing on synergies between data centers, heat receivers and utility suppliers. And on the other hand, focusing on the differences in the business models presented by the two case studies.

The thesis will study three cases for heat reuse and BESS for grid stabilisation:

- **Case 1: PDC's heat recovery system.**
  A detailed analysis of PDC's heat recovery system and its current relationship with the district heating company. This case study answers the following research questions:
  - How is heat recovery done?
  - What services are provided to contribute to grid stabilisation?
  - What synergies are made between data centers, heat receivers and utility providers?
- **Case 2: T-Loop's Data Energy Center.**
  This case study aims to explain T-Loop's revolutionary Data Energy Center concept through the analysis of their heat recovery system and their contribution to grid stabilisation. This case study answers the following research questions:
  - How is heat recovery done?
  - What services are provided to contribute to grid stabilisation?
  - What synergies are made between data centers, heat receivers and utility providers?

Once the analysis of the two case studies is done, a final case study comparison is presented. This section focuses on the differences in the business models presented by the two companies. Mainly focusing on the heat repurpose process, as each company presents a different model. Although grid stabilisation services are only provided by one of the companies, they will also be taken into account when discussing the business model, as it adds to the company's value proposition.

Through the interviews and visits done in the data gathering process, data centers were asked to provide data on how they make these new systems profitable, but in most cases, no answer was given since these technologies are new to the market, and the sharing of this type of data could compromise their competitiveness. Hence, a comparative case study approach opens a way to analyse differences in their business models. As explained in sub-chapter 2.7, Osterwalder proposes nine blocks which define business models, but as not all of them are applicable or of interest to this paper, six of those blocks will be analysed in this comparison, they are: key resources, key activities, value propositions, key partners, cost structure and revenue streams. This section answers the following research question:

  - What business models are put in place to make these technologies profitable?

A final discussion of the results includes the interpretation of the findings, implications and limitations of the study. Reviewing all the data available on the case studies and outlining the key aspects to draw a conclusion for the thesis.

The conclusion answers the main research question of the thesis, which is:

- **Can data centers become active contributors to the energy grid through waste heat recovery and grid stabilisation?**

This method aims to give a clear structure for the reader to follow, going from the abstract theory to the analysis of real-world concepts and implementations, and concluding with answers to the main problem.

5 Case Studies

5.1 Case Study 1: PDC's heat recovery system

5.1.1 Introduction

PDC, or Parallelldatorcentrum, is a supercomputer located on KTH's campus at the School of Electrical Engineering and Computer Science (EECS). It is the fastest academic supercomputer available for academic research in Sweden. Its goal is to collaborate with academic research and business R&D. PDC also provides data storage facilities, application experts and other services, including HPC training. (KTH, 2025). Although PDC is not a data center by definition, it works similarly to a cloud data center. Clients looking to use PDC's computing power send a request via an online system, asking how many cores and hours they need. The workload manager looks for a gap with the predefined requirements and books it for the client. This is very similar to how cloud DCs work; clients never get to see the computer they use; they just rent a certain amount of computing power. PDC's current supercomputer is named 'Dardel', and it is able to provide 13.5 petaflops of peak performance (Butler, 2023). To put it into context, a new top-of-the-line smartphone can perform 5.6 teraflops; Dardel is 2400 times more powerful.

5.1.2 How is heat recovery done?

PDC uses both air- and liquid-cooled computers, although the majority is liquid-cooled. Liquid-cooled systems account for 800kW, while air-cooled systems account for 35kW. This means that 96 per cent of PDC's total computing power is liquid-cooled, while only 4 per cent is air-cooled. The air-cooled systems are mainly used for the storage of the supercomputer, as hard drives cannot use liquid cooling because they have to be easily replaceable.

The cooling system at PDC is composed of two main loops, the water-cooled loop and a smaller air-cooled loop. The liquid cooling loop uses direct-to-chip technology; some parts of the computer run hotter than others, with the hottest one achieving output temperatures of close to 37ºC, while other parts stay at around 23ºC. These temperature readings can be seen in Figures 10 and 11 below.

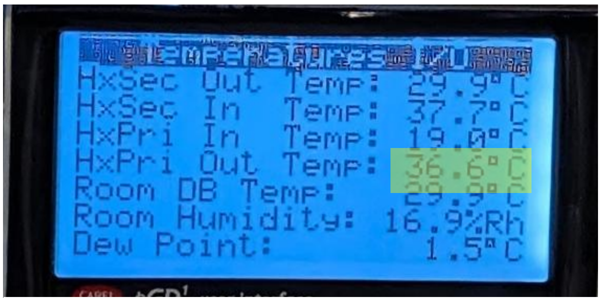

Figure 11 *PDC liquid-cooled temperature readings 1.*

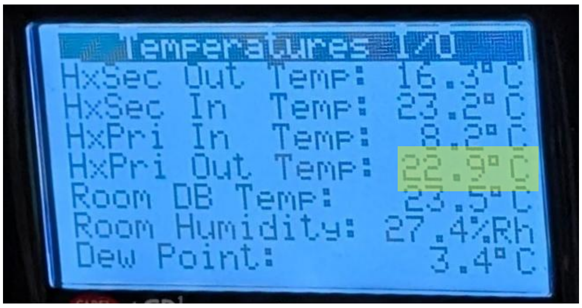

Figure 10 *PDC liquid-cooled temperature readings 2.*

As for the air-cooled machines, although they account for a small percentage of PDC's total power, they still produce considerable and constant amounts of heat. The air- and liquid-cooled systems are placed in different rooms. PDC uses computer room air conditioning (CRAC) units for air cooling; one of them is shown in Figure 13. To maximise cooling efficiency, the servers are set in a hot/cold aisle displacement, and hot aisle containment is not put in place, so hot and cold air is mixed in the room. As air-cooled systems account for such a small portion of PDC's consumption, this cooling system is not nearly as optimised as the liquid-cooled part. In this case, the CRAC units recover hot air at a temperature of around 21ºC, as seen in Figure 14 below.

Figure 14 *PDC air-cooled temperature reading.*

Figure 13 *CRAC unit at PDC.*

As seen above, various systems are used to keep PDC's servers running at optimal temperature. Although these systems differ quite a bit in how they cool the computers, the heat they extract is recovered similarly. Figure 12 shows a schematic drawing of PDC's heat recovery system. PDC buys cooling from Akademiska Hus; this cooling is provided to PDC via the KP Ø1 pipeline. The heat recovery system is divided into two main loops, one connected to Dardel's CPU and another one connected to Dardel's GPU (more systems are connected to these loops, but Dardel's CPU and GPU are the main heat producers); these two loops then have an intermediate cooling loop before the heat is transferred back to the main pipeline (KP Ø1).

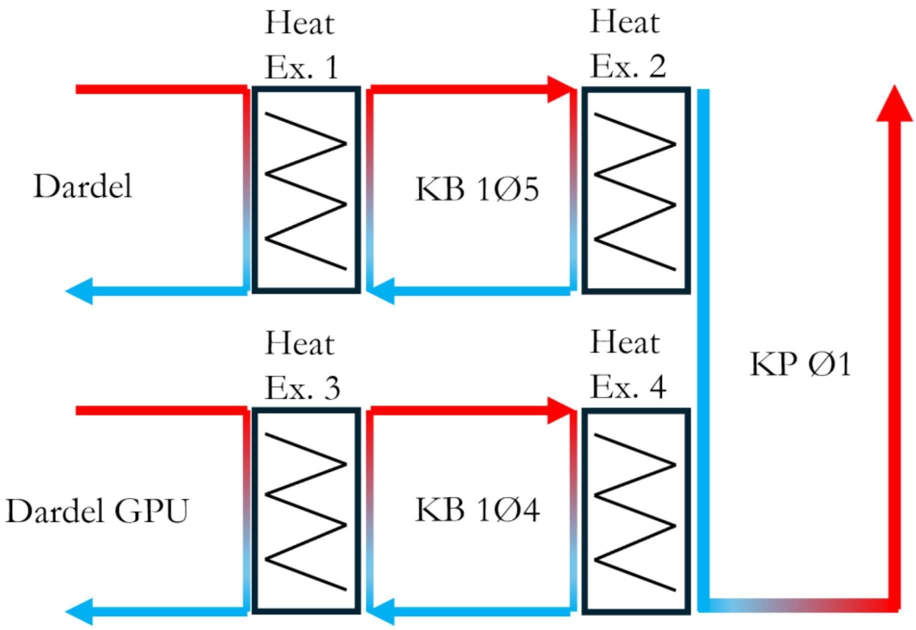

Figure 12 *Schematic of PDC's heat recovery system.*

Starting from the heat source, heat is produced by Dardel's CPU and GPU; this heat is then transferred to the intermediate coolant loops KB 1Ø5 and KB 1Ø4, respectively, through a heat exchanger. Finally, KB 1Ø5 and KB 1Ø4 are cooled through two heat exchangers connected to KP Ø1. The heat absorbed by KP Ø1 is then returned to Akademiska Hus' heat pumps at a temperature of 17ºC, due to the mixing of the high and low temperature flows and the losses in the transfer. PDC's server halls produce around 1MW of heat all year round, making it a major heat source for Akademiska Hus, who can repurpose the heat for space heating and hot water production.

PDC has to guarantee a certain uptime for its users. The system shown above only shows the system running with district cooling, which is the system used the majority of the time. Nevertheless, district cooling may fail; in that case, ordinary tap water is used to keep the cooling loop at an admissible temperature (PDC, 2025).

5.1.3 What services are provided to contribute to grid stabilisation?

Grid stabilisation services are not provided by PDC. They have a battery system as part of their UPS, but it is not prepared to do so, as stated by Luca Manzari, systems manager at PDC, when asked about grid stabilisation:

"We haven't done it for a simple reason: you need special hardware to manage the UPS… We can't do it because our system is from a time when such an application wasn't conceived." (Manzari, 2025)

PDC's batteries are lead-acid; this type of battery has limited deep-cycle capabilities, which limits its possibility of operating in the grid stabilisation market (May et al., 2018). Moreover, to not jeopardise the operation of the supercomputer, the capacity of the batteries would have to be increased to guarantee the same uptime while running on the batteries if grid stabilisation services were provided. The main role of the UPS is to protect against even the briefest power outage; providing grid stabilisation services would mean sacrificing this reliability, as the system is not prepared.

5.1.4 What synergies are made between the data center and the heat receiver?

Akademiska Hus provides cooling to PDC through their district cooling network and then absorbs the heat PDC produces. The collaboration of both entities has helped to develop a more efficient heating system on the campus. Instead of having to buy and run expensive chillers, PDC rejects its heat to Akademiska Hus, who also reduces the running costs of their machinery as they get an inlet of chilled water for their heat pumps. This collaboration increases the value of both assets through efficient planning and resource utilisation. In essence, what would traditionally be a wasted by-product becomes a valuable resource. Turning a high-performance computing center into a climate-positive contributor to the campus ecosystem.

PDC's heat recovery system is an excellent example of how co-location can help increase efficiency in such systems. PDC and Akademiska Hus' heat pumps are located just 100m apart from each other, so their proximity to one another reduces heat losses to a minimum and simplifies the transfer of heat.

As stated above, synergies between Akademiska Hus and PDC seem natural, as sharing their resources benefits both parties. Nevertheless, the heat produced by PDC has a value, and Akademiska Hus considers it when pricing district cooling for PDC. A special contract is put in place to compensate PDC for the heat they are injecting into the district cooling grid.

This contract takes into account the power that PDC is putting into the grid without worrying about temperature requirements, as stated by Luca Manzari, when asked about the requirements set by Akademiska Hus to inject heat into their grid:

"They have flow meters and temperature sensors, but we do not have any requirements to have the temperature higher than x. Of course, we try to put it as high as possible." (Manzari, 2025)

Nevertheless, PDC's heat output cannot be higher than 1.1 MW, as that is the maximum Akademiska Hus can absorb; if they output more heat than that, the auxiliary systems at PDC would have to start, and not all of the heat would be sent to Akademiska Hus. Currently, PDC is staying below that 1.1 MW requirement because, as stated above, they output 1 MW year-round. This makes it possible to recover all the heat they produce.

5.2 Case Study 2: T-Loop's Data Energy Center.

5.2.1 Introduction

T-Loop is a data center provider that is introducing an innovative way of looking at data centers. Introduced in 2020 and operational since 2022, their Data Energy Center concept represents a groundbreaking solution that combines data processing with sustainable energy usage. They combine expertise in the data centre and energy sectors, enabling them to create seamless synergies between them.

"T-Loop stands with one leg in data centres and one leg in the energy field. So, we can translate between these two, and we understand both business, technology and the needs of both sectors," said Staffan Styme, deputy CEO of T-Loop (Styme, 2025).

T-Loop's first data center and proof-of-concept is located in Häggvik, Stockholm, at a business center building. This data center has a total capacity of 0.2 MW, with 30% less energy consumption than that of a traditional data center and with 100% lower climate CO_2 footprint thanks to Sweden's green energy mix and their heat recovery and grid stabilisation systems (T-Loop, 2025). They design, build and operate edge data centers strategically located in buildings with heating needs that match the excess heat produced by the servers. Through collaboration with forward-thinking property owners and energy companies, T-Loop is setting a new standard for how high-density demands of modern technologies like HPC, AI, Iot, and Blockchain can be processed while maintaining a sustainable energy model. As stated by their deputy CEO, the traditional and wasteful model of running data centers was what encouraged the creation of T-Loop:

"That's why we started T-Loop. We were frustrated with the untapped resources in data centres" (Styme, 2025)

T-Loop's energy optimisation doesn't end with waste heat recovery. They implement grid stabilisation services through their UPS system. These services help the national grid integrate renewable energy sources like wind and solar. With the apparition of new regulations from the government bodies, like the energy effective directive put in place by the EU, which introduces an obligation for the monitoring and reporting of the energy performance of data centers, operators will have to look out at more sustainable and efficient models for running their centers and T-Loop's model is a proof of concept that high-density computing demands can be energy efficient and aligned with sustainable energy production.

5.2.2 How is heat recovery done?

T-Loop's proof-of-concept model integrates a cutting-edge cooling system, based on immersion technology. This enables T-Loop to produce a constant flow of heat at high temperatures of around 50-55ºC. As stated by Styme:

"In the proof of concept we have built and showcase that we use immersion cooling. That means that we can produce heat at a temperature of maybe 50-55 degrees" (Styme, 2025)

T-Loop looked at various options for liquid cooling before deciding to use immersion. They saw that direct-to-chip cooling could provide higher output temperatures, of around 70ºC, but the problem is that direct-to-chip only cools one part of the computer, around 70%, the other 30% comes from dissipation from the rest of the computer and has to be dealt with by an air-cooled system. At the end, two heat flows are created, a high-temperature heat flow from the direct-to-chip loop and a lower-temperature flow from the air-cooled. The mix of these two could provide a final output temperature of 40-45 ºC, which is lower than that of the immersion cooling system, which is the reasoning behind T-Loops' choice of cooling system.

After the heat is extracted from the servers by the coolant in the immersion cooling system, it goes to the coolant distribution unit or CDU, where the heat is transferred to a water-based system through a heat exchanger. Finally, there is an additional heat exchanger that transfers the heat to the heating system of the building or the district heating grid, depending on the season or heat demand. The heat exchangers used are plate heat exchangers, which have high efficiency levels and reduce heat losses to 1-2 ºC, as stated by Styme:

"Plate heat exchangers are very efficient, you may be losing maybe one or two degrees in temperature" (Styme, 2025)

T-Loop's system is prepared to recover 100% of the heat that it generates. Even if 100% of the heat is recovered, there needs to be a large and reliable enough source to absorb the heat. T-Loop designs its Data Energy Centers so that at least 70% of the recovered heat is reused (Styme, 2025). To do so, they look for buildings with heating needs that match the DC's heat output to minimise waste. Currently, according to their proof-of-concept model, T-Loop uses direct heat recovery, meaning that the heat outputted by the servers is directly pumped into the building's heating system without the need for heat pumps. This system can also regulate the temperature by managing the flow of liquid on the racks to match the building's needs as accurately as possible. However, as Styme said, this is not always possible:

"We are not saying that we always can use direct heat recovery. It's depending on the heat sink demand for temperature, also the server load and the cooling technology." (Styme, 2025)

For example, storage demands in IT need to be air-cooled for accessibility reasons, so having loads with significant storage needs reduces the possibilities of heat recovery, as heat pumps would be needed to elevate the flow temperature. Even in T-Loop's proof of concept, heat cannot always be directly recovered. During summer, heat demand in the building is low, so T-Loop has to look for another heat sink. At the moment, they use a heat pump to elevate the output temperature of the servers even further so that it can be injected into the district heating grid. This creates a system that not only ensures heat is recovered year-round but also acts as a backup in case the heat demands of the building drop during any other season.

To summarise, T-Loop uses immersion cooling as it provides a constant flow of heat, at a temperature of 50-55ºC so that it can be directly used in the heating system of the building in which the DC is placed. Nevertheless, heat pumps are used during the summer, when heat demand from the building is lower, to inject heat into the district heating grid.

5.2.3 What services are provided to contribute to grid stabilisation?

To contribute to grid stabilisation, T-Loop provides ancillary services by participating in the Frequency Containment Reserve or FCR market. This is done through the utilisation of their UPS's overcapacity. At T-Loop's proof-of-concept, the UPS consists of Lithium-Ion batteries; it is prepared to incorporate diesel generators, but these were not installed at the time of the interview. The use of the UPS to provide ancillary services is carefully managed to ensure that the operational integrity and availability of IT hardware remain uncompromised.

To make the UPS FCR ready, proprietary hardware and software were specifically developed to meet the stringent reliability demands of a data center. Participation in the FCR market also requires batteries with capacities exceeding those typically required solely for backup purposes. Although this oversizing comes with a cost, it not only extends the DC's possibilities regarding grid stabilisation but also enhances the DC's resilience by extending its runtime during potential power outages. As mentioned above, T-Loop plans to incorporate diesel generators into their current UPS. The addition of such generators would enable them to participate in slower balancing markets, thereby expanding the range of ancillary services they provide.

The participation in these markets is subject to approval and prequalification by the Swedish electricity transmission system operator, Svenska Kraftnät. This process involves a series of compliance tests to ensure the provider's capability to respond to grid frequency deviations as specified. Market participation operates on a bidding basis, where providers submit day-ahead and intraday bids through a Balance Responsible Party (BRP), who aggregates these for submission to Svenska Kraftnät. The activation of services is automatic and triggered by real-time frequency deviations in the grid.

In summary, the data centre contributes to grid stabilisation by offering ancillary services, primarily FCR, through an advanced integration of battery systems and planned generator support, under a regulated, automated, and economically incentivised market framework.

5.2.4 What synergies are made between T-Loop, heat receivers and utility providers?

The relationship between T-Loop and its primary heat receiver, which is its host building, goes beyond an exchange of heat. At T-Loop's proof of concept, they are a tenant to the building's owner and at the same time, provide heating for the building. Both parts act as suppliers and customers to one another. On one side, the landlord provides space for T-Loop to house their data center; on the other, T-Loop provides heating back to the building.

As explained above in section 5.2.2, T-Loop looks for buildings that match their DC's heat output so that heat absorption is maximised. Their Data Energy Center provides hot water at the temperatures needed by the building at each moment. It is not only T-Loop that benefits from outputting heat to the building, the building itself also gets a reliable, renewable and sustainable in-house heat source. Heat production from T-Loop's Data Energy Center comes from Sweden's green energy mix, thus carbon emissions are minimised.

When building a new Data Energy Center, T-Loop works closely with the landlord to make an agreement on how the space-heat exchange will work. As stated by Styme, this negotiation varies from building to building, and they are flexible to the property owner's requisites:

"We don't have a standard model for that because it's depending on what the property owner prefers…Sometimes we pay a full rent, and we charge the property owner about the same as they pay for just heating, but if we pay less rent they may pay less for the heat as well." (Styme, 2025)

T-Loop prices its heat at a similar price to the one offered by district heating companies. This makes them competitive in the market and more appealing to customers. As said by Styme, their model is adaptable to the landlord's needs and the heat can either be paid for or exchanged for a deduction in the rental agreement. During periods of low heat demand by the building, T-Loop outputs its heat to a district heating company. These periods of low demand bring heat prices down, making it difficult for heat producers to make revenue, but for data centers, it is still a great option. Exporting the heat out to a district heating grid means not having to start their own chillers to cool the system, even if heat prices are low, it is still an interesting option for DC, as Styme said:

“In Stockholm, you have the base production that you use is energy to waste. And that's a really low production cost. So, that means you need to compete with a cost that is almost zero. So, in that case, it's not profitable, but it might be better to do that than to cool it away because that's a cost as well.” (Styme, 2025)

The synergies between T-Loop and both the property owner and the district heating company make heat recovery possible. Furthermore, they open a new revenue stream for T-Loop, as heat is sold or exchanged for a discount on the rent to the landlord. It is not only tangible revenue that comes with heat recovery. It also differentiates T-Loop from its competitors by making it more environmentally friendly and sustainable. With the new Energy Efficiency Directive, approved by the EU, data centers will have to report key energy performance indicators annually, including energy reuse and heat recovery. This means that data centers that do not already have heat recovery systems in place will have to implement them sooner rather than later, as new EU regulations encourage these technologies. With the approval of such regulations, synergies between data centers and heat receivers are crucial and more important than ever to make their compliance possible.

Synergies between T-Loop and the utility provider are built around the mutual benefits derived from integrating the data center’s UPS into the grid services framework. These synergies are realised through both technical and economic collaboration mechanisms, with a focus on improving grid stability, enhancing energy efficiency, and optimising resource utilisation. The participation of the DC in the ancillary services market contributes to reducing the use of more traditional, less efficient peaking plants and fast-forwarding grid assets. This collaboration with the transmission system operator, Svenska Kraftnät, facilitates enhanced utilisation of the data center’s electrical infrastructure. Overcapacity in battery systems is now dual-purposed for grid services, creating economic value from assets that would otherwise remain idle during normal operations.

The data centre coordinates with a Balance Responsible Party (BRP), who aggregates and submits energy bids to Svenska Kraftnät. The remuneration model has transitioned from a "pay-as-bid" to a "marginal pricing" system, whereby all accepted bids are compensated at the highest cleared market price, thus incentivising participation and enhancing market efficiency. This process allows private infrastructure providers to create functional and strategic alignments with the national grid operator, enabling access to distributed ancillary resources while offering data centers participation in a remunerative energy market.

Collaboration between data centers and electricity providers play a significant role in integrating renewable energy sources into the Swedish power grid. Renewables such as wind and solar introduce greater variability and intermittency into the electricity supply, and the need for fast, flexible, and decentralised balancing resources becomes significant. By transforming the data center’s UPS into a grid-interactive asset, fluctuations in renewable generation can be absorbed and frequency deviations can be smoothed. In doing so, data centers not only support decarbonisation but repurpose an infrastructure to create new revenue sources and increase sustainability in their operation.

5.3 Case Study Comparison

This section compares both case studies from a business model point of view. Utilising Osterwalder's methodology, explained in section 2.7, the business models of both cases would be analysed and then compared. First, PDC's heat recovery system will be described from a business model point of view. If we look at PDC as a company, this comparison only focuses on one of its products, excess heat. T-Loop's Data Energy Center will also be described from a business model's point of view, focusing on the contribution to the energy system, heat repurposing and grid stabilisation.

5.3.1 PDC's Heat recovery

Table 3 *PDC's heat recovery business model.*

Building blocks	Description
Key partners	Akademiska Hus: Provides district cooling and repurposes excess heat.
Key activities	High-performance computing and storage resources: Generate the necessary heat for repurposing.
Key resources	IT, cooling and auxiliary equipment and personnel: Ensure that heat is produced constantly.
Value propositions	Excess heat: The heat generated by the IT equipment is captured.
Cost structure	Hardware, cooling and auxiliary systems acquisition, personnel salary and energy bills: Primary running costs of the supercomputer, necessary to generate heat.
Revenue streams	Excess heat repurposing: Akademiska Hus absorbs the heat generated and repurposes it for heating and hot water distribution.

5.3.2 T-Loop's Data Energy Center

Table 4 *T-Loop's Data Energy Center business model.*

Building blocks	**Description**
Key partners	Property owner, district heating company and Svenska Kräfnat: The property owner provides space in exchange for heating, the district heating company repurposes excess heat in warmer months, and Svenska Kräfnat relies on T-Loop's UPS for grid stabilisation.
Key activities	Colocation and edge computing services and UPS maintenance: Colocation and edge computing generate the necessary heat for repurposing, and UPS maintenance ensures grid stabilisation services can be provided.
Key resources	IT, cooling and auxiliary equipment, UPS and personnel: Ensure that heat is produced constantly and the UPS makes providing grid stabilisation services possible.
Value propositions	Excess heat and grid stabilisation services: The heat generated by the IT equipment is captured, and grid stabilisation services are provided to Svenska Kräfnat.
Cost structure	Hardware, cooling and auxiliary systems acquisition, UPS, R&D for grid connection, personnel salary and energy bills: Primary running costs of the data center, necessary to generate heat. R&D for developing hardware and software to connect the UPS to the grid and the costs of running and acquisition of the UPS.
Revenue streams	Excess heat repurposing and grid stabilisation services: The building where the DC sits uses the excess heat for heating, the property owner either pays for the heat or makes a discount on the rental agreement. Grid stabilisation services are remunerated based on power and energy injected into the grid.

5.3.3 Discussion

From the building blocks described above in sections 5.3.1 and 5.3.2, both business models can now be compared, focusing on how each entity has managed to contribute to the energy system utilising excess heat and grid stabilisation. Both models utilize waste heat from IT equipment to provide heating and hot water, contributing to circular economy principles. Nevertheless, heat produced by PDC arrives at Akademiska Hus at a temperature of around 17ºC. In contrast, the heat generated by T-Loop's Data Energy Center enters the building's heating system at around 50-55ºC. This means that the heat generated at PDC must be increased in temperature by heat pumps, before it can be used for heating or hot water. Meanwhile, T-Loop's heat is already at the temperature the building needs. PDC's heat has to go through a heat-up process after it exits its facilities, which implies higher costs for Akademiska Hus. Hence, the price of PDC's heat is lower than that produced by T-Loop, as it is of a lower quality.

Regarding cost structure, both entities have to make elevated investments in hardware, cooling, auxiliary systems and personnel. However, there are key differences; T-Loop's cooling system is more traditional and standardised in the DC sector, as it uses cooling methods that have been used for some time now in the industry. This makes for a less efficient cooling system, which is part of the reason for the lower heat output temperatures, but comes at a reduced price when it comes to installation. Meanwhile, T-Loop implements a cutting-edge cooling system, which enables them to output such high temperatures. Designing and implementing such a system increases investment costs as it is less standardised in the industry. T-Loop invested in a battery system and in developing hardware and software to make connecting their UPS to the grid possible. On the other hand, PDC has a battery system that cannot be connected to the grid, as they do not have the necessary hardware or software to do so. This prevents PDC from providing any type of grid stabilisation, leading to unused resources and a weaker value proposition, as they only produce excess heat. Either way, both companies still have to face significant energy bills to keep all systems running.

Partnerships are a fundamental part of both business models, while PDC only relies on Akademiska Hus for its heat recovery, T-Loop has two possible heat recipients; the heating system in the building on which the DC resides and a district heating company, when the heat demand from the building is not high enough. Although both companies' models enable 100% heat recovery and repurposing, T-Loop's model has a backup in case one of the heat receivers were to fail, whereas PDC entirely relies on Akademiska Hus. Finally, the partnership between T-Loop and Svenska Kräfnat ensures grid stabilisation services can be provided.

(Dowds & El-Saghir, 2021) described the potential of excess heat in edge data centers, but working models were not analysed. This study shows that their theoretical work could be made possible, as T-Loop's heat reuse model significantly aligns with one of the case studies they conducted. Furthermore, PDC's heat exchange model also aligns with another of the case studies (Dowds & El-Saghir, 2021) conducted in their work, although PDC is not an edge data center, their heat recovery model is very close to one of the ones analysed in their work.

In conclusion, while both models interact with the energy system, PDC's model seems to have come from the adaptation of existing infrastructure to repurpose heat. Meanwhile, T-Loop's model looks at it from another point of view, starting from the design process of the DC. Contributing to the energy sector was always one of their goals. The result is a holistic design that enables them to create new revenue streams from otherwise wasted resources. The combination of grid stabilisation services with heat reutilization makes their operations not only more efficient but also more environmentally conscious. Nevertheless, PDC's system still makes good use of the existing infrastructure and allows for 100% heat recovery, even if it is at lower temperatures.

6 Conclusions and future work

This thesis set out to explore whether data centers could become active contributors to the energy system. By analysing the data center's systems, two possible mechanisms were found to do so: heat recovery and repurposing and grid stabilisation services. The results show that data centers, traditionally seen as large energy consumers, can also act as valuable assets by repurposing otherwise wasted or unused resources like excess heat and BESS. Findings suggest that integrating data centers into energy planning is not only possible but can benefit both parties while contributing to sustainability and energy resilience. This chapter shows the conclusions drawn from the study.

- Heat reuse can be done in a number of ways, as the variety of cooling and heat recovery systems and heat recipients is wide. Two different heat recovery systems were analysed in the case studies; results show that even though output temperatures from both systems differ, both of them contribute to a heat recipient, who, regardless of the temperature, can make use of the heat.
- Agreements between data centers and heat recipients vary and are adapted to each case. Factors like temperature, distance, need for heat and space availability are taken into account when putting a price on the heat. As seen in T-Loop's case, heat can also be exchanged for a discount in rent. Regardless of the type of agreement, heat reuse brings a new revenue stream for data centers from what would otherwise be a wasted byproduct.
- Heat recipients can be diverse. Although injecting heat into a district heating grid may sound like the best option, T-Loop's system shows that there are other possible solutions that reduce heat losses and bring computing power closer to the end user.
- The European Union's Energy Efficiency Directive encourages the implementation of heat recovery systems in data centers, thus putting data centers that have already incorporated them ahead of their competitors.
- Providing grid stabilisation services complements data center operations, as it makes better use of their BESS while helping the grid to integrate more renewable power sources, like wind and solar.
- Designing data centers to integrate heat recovery and provide grid stabilisation services, creates new revenue streams, makes operations more sustainable and efficient and makes better use of existing resources.

6.1 Future work

While this thesis demonstrates the potential of data centers to contribute to the energy system, several areas remain for exploration. Firstly, although data and insights were gathered from the interviews and visits, economic type data could not be gathered, and specific investment and operational costs were not included in the thesis. Data center operators confirmed that using these technologies improved their business model, but without specific metrics, it is unknown to what extent. The focus of this thesis was on data centers situated in the Stockholm area, due to weather, energy mix and availability of heat recipients, it makes for an unbeatable spot for placing data centers, but it is still to be seen if these technologies could be as useful in other parts of Sweden and Europe. Only two cooling solutions were analysed in this paper for heat recovery; new technologies arise daily, and the exploration of other possible solutions remains open. Finally, other possible heat reuse methods could be analysed, outside of heating and hot water.

References

ASHRAE. (2024). *Equipment Thermal Guidelines for Data Processing Environments.* Peachtree Corners.

Aasvestad Sandvik, S. (2024). *Comparative analysis of the Swedish.* Norwegian University of Science and Technology.

Al Kez et al., D. (2025). *Ai-Driven Cooling Technologies for High-Performance Data Centers: State-of-The-Art Review and Future Directions.* Manchester: SSRN.

Ali Habibi, K., & K. Halgamuge, S. (2017). A Review on efficient thermal management of air- and liquid-cooled data centers: From chip to the cooling system. *Applied Energy*, 1165-1188.

Aligned Data Centers. (2018). *COMPARING PUE BETWEEN COLOCATION DATA CENTERS – AND WHY IT MATTERS.* Aligned Data Centers.

Andersen, M. (2024). *Hydrogen and Battery Energy Storage Systems: A Techno-economic modelling of Ancillary Services and Additional Income flows.* Lund: Lund University.

ASHRAE. (2016). *Data Center Power Equipment Thermal Guidelines and Best Practices .* ASHRAE.

Boesten et al., S. (2019). 5th generation district heating and cooling systems as a solution for renewable urban thermal energy supply. *Adv. Geosci.*, 129-136.

Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 77-101.

Butler, G. (2023, January 23). Sweden's KTH Royal Institute of Technology to get new supercomputer. *Data center dynamics.* Retrieved from https://www.datacenterdynamics.com/en/news/swedens-kth-royal-institute-of-technology-to-get-new-supercomputer

Cool IT systems. (2023, February). *News Archive › 10 Considerations for Implementing Direct Liquid Cooling in Data Centers.* Retrieved from Cool IT systems: https://www.coolitsystems.com/10-considerations-for-implementing-direct-liquid-cooling-in-data-centers/

D. Thorén, T., & O. C. Widell, A. (2011). *Development of Liquid Cooling for PCs Component Analysis for Factory Sealed Systems.* Göteborg: Chalmers University of Technology.

De Wit, B., & Meyer, R. (2004). *Strategy, Process, Content, Context.* Thomson.

Dowds, E., & El-Saghir, F. (2021). *Utilising waste heat from Edge-computing Micro Data Centres Financial and Environmental synergies, Opportunities, and Business Models.* Stockholm: KTH Royal Institute of Technology.

Ehrlin, E. (2024). *Profit Optimization of Battery Energy Storage Systems.* Luleå: Luleå University of Technology.

Gigabyte. (2023). *Single-phase Immersion Cooling with GIGABYTE.* Retrieved from GIGABYTE: https://www.gigabyte.com/Solutions/gigabyte-single-phase

Google. (2024). *Google data center PUE performance.* Retrieved from Google Data Centers: https://datacenters.google/efficiency/

Haghshenas et al., K. (2023). Enough hot air: the role of immersion cooling. *Energy inform.*

Hinterholzer, S., & Hintemann, R. (2022). Data centers 2021. Cloud computing drives the growth of the data center industry and its energy consumption. *Borderstep Institute.*

Howell, J. (2024, 10 23). *Enterprise Data Center – The Complete Guide [2024]*. Retrieved from Encor Advisors: https://encoradvisors.com/enterprise-data-center/

Huang et al., P. (2020). A review of data centers as prosumers in district energy systems: Renewable energy integration and waste heat reuse for district heating. *Applied Energy*, Volume 15.

Hyvönen et al., J. (2024). Potential of solar photovoltaics and waste heat utilization in cold climate data centers. Case study: Finland and northern Japan. *Renewable and Sustainable Energy Reviews.*

Jerléus et al., K. (2024). Environmental footprints of the data center service sector in Sweden. *Heliyon*, Volume 10, Issue 11.

Kakabadze, S. (2024). *Integrating Battery Storage in Ancillary Service Markets: Market Analysis and Comparative Study of Design and Regulatory Trade-offs in Great Britain, Texas, and Sweden.* Stockholm: KTH Royal Institute of Technology.

Keshan et al., H. (2016). Comparison of lead-acid and lithium ion batteries for stationary storage in off-grid energy systems. *4th IET Clean Energy and Technology Conference (CEAT 2016).*

KTH. (2025, March). *PDC Center for High Performance Computing*. Retrieved from KTH: https://www.pdc.kth.se/

Lamp, S., & Samano, M. (2022). Large-scale battery storage, short-term market outcomes, and arbitrage. *Energy Economics*, Volume 107.

Lei, N., & Masanet, E. (2022). Climate- and technology-specific PUE and WUE estimations for U.S. data centers using a hybrid statistical and thermodynamics-based approach. *Resources, Conservation and Recycling*, Volume 182.

Lund et al., H. (2014). 4th Generation District Heating (4GDH): Integrating smart thermal grids into future sustainable energy systems. *Energy*, 1-11.

Lygnerud et al., K. (2022). Risk assessment of industrial excess heat collaborations – Empirical data from new and ongoing installations. *Energy*, Volume 255.

Mann, T. (2023, April 14). LiquidStack CEO on why you shouldn't ignore immersion cooling. *The Register.* Retrieved from https://www.theregister.com/2023/04/14/liquidstack_immersion_cooling

Manzari, L. (2025, February). (D. Vicedo Candela, Interviewer)

Map, D. C. (2025). *Data Center Map*. Retrieved from https://www.datacentermap.com/sweden/

May et al., G. J. (2018). Lead batteries for utility energy storage: A review. *Journal of Energy Storage*, 145-157.

Nada, S., & Said, M. (2017). Effect of CRAC units layout on thermal management of data center. *Applied Thermal Engineering*, 339-344.

Nadjahi et al., C. (2019). A review of thermal management and innovative cooling strategies for data center. *Sustainable Computing: Informatics and Systems*, 14-28.

Ommen et al., T. (2016). Lowering district heating temperatures – Impact to system performance in current and future Danish energy scenarios. *Energy*, 273-291.

Oró et al., E. (2018). Design and economic analysis of liquid cooled data centres for waste heat recovery: A case study for an indoor swimming pool. *Sustainable Cities and Society*, 185-203.

Osterwalder, A. (2004). *THE BUSINESS MODEL ONTOLOGY A PROPOSITION IN A DESIGN SCIENCE APPROACH.* Lausanne: Université de Lausanne.

Pans et al., M. (2023). Modelling of 4th generation district heating systems integrated with different thermal energy storage technologies – Methodology. *Energy Conversion and Management*, Volume 276.

PDC. (2025, January 7). *Cooling*. Retrieved from PDC Center for High Performance Computing: https://www.pdc.kth.se/hpc-services/infrastructure/cooling-1.737778

Pradhan, R. (2015). *Software defined infrastructures : implications for technology and business strategies for competing in the era of hyper-scale computing.* Massachusetts Institute of Technology.

Raj Suyambu, M., & Kumar Vishwakarma, P. (2024). Improving grid reliability with grid-scale Battery Energy Storage Systems (BESS) . *International Journal of Science and Research Archive*, 776-789.

Rehan et al., K. (2025). Data Center Waste Heat Reuse: An Investment Analysis. *Engineering for Sustainable Buildings and Cities*.

Rezaei-Mayahi, M. (2019). Temperature-aware power consumption modeling in Hyperscale cloud data centers. *Future Generation Computer Systems*, 130-139.

Salih Erden, H. (2022). Comprehensive energy and economic assessment of CRAH bypass method in air-cooled data centers. *Sustainable Energy Technologies and Assessments*, Volume 52, Part B.

SBBA. (2021). *Heating Market Report 2021*. SBBA.

Stockholm Exergi. (2025). *Heat recovery*. Retrieved from Stockholm Exergi: https://www.stockholmexergi.se/foretag/fjarrvarme-for-foretag/varmeatervinning/

Stockholm Exergi. (2025, 2 20). *Installationsanvisning Öppen Fjärrvärme.* Retrieved from Stockholm Exergi: stockholmexergi.se

Styme, S. (2025, March). (D. Vicedo Candela, Interviewer)

Svenska kraftnät. (2025, 1 15). *Information on different ancillary services*. Retrieved from Svenska kraftnät: https://www.svk.se/en/stakeholders-portal/electricity-market/provision-of-ancillary-services/information-on-different-ancillary-services/#:~:text=FFR%20%E2%80%93%20Fast%20Frequency%20Reserve,in%20the%20Nordic%20power%20system.

Tanriverdi, H. (2006). Performance Effects of Information Technology Synergies in Multibusiness Firms. *MIS Quarterly 30*, 57-77.

Taylor, H. (2023). *The Edge Data Center*. Business Expert Press.

Teece, D. J. (2010). Business Models, Business Strategy and Innovation. *Long Range Planning*, 172-194.

T-Loop. (2025). *DEC001*. Retrieved from T-Loop: https://www.tloop.se/dec001-haggvik

Yang et al., E. (2024). Integrated Energy Model for Carbon Neutral Agriculture: Utilizing Waste Heat and Renewable Energy. *2024 IEEE Asia-Pacific Conference on Applied Electromagnetics (APACE)* (pp. 293-296). Langkawi, Kedah: IEEE.

Yllmark, V. (2024). *Sustainable Business Models in the Context of Sustainability Transitions The Case of Co-location Data Centres* . Stockholm: KTH Royal Institute of Technology.

Zeh et al, R. (2021). Large-Scale Geothermal Collector Systems for 5th Generation District Heating and Cooling Networks. *ResearchGate*.

Zülhsdorf et al., B. (2018). Analysis of possibilities to utilize excess heat of supermarkets as heat source for district heating. *Energy Procedia*, 276-285.

Appendix A: Interview with PDC representative.

This appendix shows the transcription of an interview held with PDC's systems manager, Luca Manzari, on the 10th of February 2025 at PDC's facilities. The interview is divided into subjects which are of interest to the project.

A.1 Heat reuse

What cooling technologies do you use at PDC?

We have a number of different things. The supercomputer itself is directly liquid-cooled, which means that there are no fans, there is no air-to-air cooling. There are heat sinks and copper heat sinks on the CPUs. The coolant runs directly into this device and is remarkably close to the system and chips. Then, this cooling loop goes into a heat exchanger. So, you have no mass transfer but heat transfer between one coolant and another coolant on a bigger loop.

Most machines are directly cooled, so water to water and then eventually it goes into district cooling, so it's out in a cooling tower somewhere and released into the air. But that is not us anymore. It is like some energy provider in Stockholm.

We have a little bit of air cooling. As with big fans and radiators, there is not so much anymore; I would say we have around 800 kW, which are directly water-cooled.

Are the water- and air-cooled systems both recovering heat for DH?

Everything eventually goes into district cooling, so everything is reused anyway. Everything is going to coolant and liquid in the end from our side. The question is whether you cool the electronic components by having air go through them, which warms up, and then you cool down this air by heating up the coolant and the cooling loop. Or if you directly cool down this component with some coolant and then you put coolant against coolant, not coolant against air. It's more efficient to do direct cooling (water-cooled).

Will it be possible to get the air-cooled systems to liquid?

You cannot turn everything to direct cooling because some things are just inconvenient to direct cool. So, what happens typically, and I would say this is common in most computer halls, is that the storage systems tend to be air-cooled because you want to have easy access to the hard drives. For example, think you have a bay that's got like 108 drives or something, and you need to get access to it; imagine if this was fully immersed in oil; it would be madness to maintain the hardware. But also the other thing is that they're not so dense. That's why it doesn't make a lot of sense to directly cool them, in comparison with the computer nodes.

On a system like ours, you have computation in one part, and then you store the data in another part, so you have separate machines. The part that does the computing consumes so much electricity, and they heat up greatly in such a little space, while the hard drives don't heat as much. For us, the compute part is around 800 KW, and the storage is 35 kW, so it is negligible when compared to the computing side. So, you can put the storage system with an air-cooling system.

Does the load curve fluctuate? How does it affect the heat recovery?

No, because we run with a scheduler, which means that a request for computing is not granted instantaneously; it goes in a queue. You have a number of users who want to run their jobs and applications. They have to send a request asking for these many cores for these many hours, and a special software called scheduler or workload manager finds a gap that fulfils the user's specifications and books it for him. The point is that you try to keep utilisation constant and maximise it because you have to keep running the cooling loop. We have many circulation pumps, and whatever workload is on the computer, they are all still running at all times, so the cost to keep the machine idling is very high. It's in our best interest to keep working at 100%. The heat recovery stays constant the majority of the time.

What percentage of heat is recovered?

In theory, 100% in the sense that everything that we produce goes back into the primary cooling loop, and the primary cooling goes to Akademiska Hus, just 100m from here. All of our heat goes there and into the radiator and sanitary system. So, in theory, 100% of our heat is reused.

Is that your primary cooling system?

Correct, as in the main cooling system. Then, we have a number of backups. Which, of course, are wasteful, but by definition, they do not run. Unless the primary cooling system fails.

Do you still recover your heat during the warmer months, when the heat demand is lower?

Yes, people still need water to wash their hands. But we have limitations. We, as PDC, are a division of the EECS school. The limitation is not on PDC but on Akademiska Hus, which is the owner of the premises because KTH rents from Akademiska Hus; the cooling pumps are theirs. We buy cooling and heating from Akademiska Hus, and we give back part of our heat, all of the heat that comes from the computer hall, so that it can be reused with an increase in quality with heat pumps.

Akademiska Hus eventually connects to Stockholm district cooling; they can't put out more than 1.1 MW of heat during the summer months, which means that we can't put out more than that; we put almost 1 MW out year-round. So we sit exactly on the limit. People still need hot water to wash their hands during summer, but it's not nearly as much as you need in the winter. So, in that case, I would say most of the heat goes wasted. From my perspective (PDC), it's reused in the sense that I give it to them (AH), and then they do whatever they want with it.

Do you have any requirements on the minimum temperature output set by Akademiska Hus?

No, they measure the power we put back. They have flow meters and temperature sensors, but we don't have any requirements to have the temperature higher than x. Of course, we try to put it as high as possible.

For me, I would like to run things as hot as possible, of course, because the entropy is higher. The quality of your heat is increased, and it is easier to reuse. But the problem is that you always have to fight with manufacturers, such as hardware manufacturers because they won't run their components as hot as possible. If you wanted to run a processor at 130°C, the manufacturer would say that the TC is 70°C. We know that processes can run hotter than that, but then they void the warranty, or they would not offer it to you with 5 years of service.

To give you some perspective, the previous flagship supercomputer, Beskow, it was sitting on a cooling loop where you had 16°C water in, and the return was no higher than 26-27°C. Now we have between 22°C and 25°C in and a return of almost 40°C hot water, which is not super good, but it's better than 25 °C. If we went higher than this, we would start to get issues from the manufacturer. At the fan blade, you can give 30°C water to the computer we have now, which is quite good because that's quite warm already, and we get it back at 40°C or so. We got this computer in August 2021, quite some time ago. It's almost on the way out now, and hopefully, the new one can run even hotter. But as I said, it is always a fight with the manufacturer because when you buy it, you open a procurement and ask, say, Dell, HP and IBM, what's your best supercomputer?

In this procurement, you'll set your requirements: I wish for this machine to have this much computing capacity, this much network, I want to take water at 40°, etc. You start negotiating because manufacturers always want to run things as cold as possible, and we want the opposite until we reach an agreement.

If you change the supercomputer, can you keep the same cooling system, or does that change as well?

It depends on how much heat you can get out, of course, you need to change stuff, for example, the previous machine we had was air cooled. So, of course, you need to redirect stuff. The bottleneck is always how much heat we can get rid of; we can't give more than 1.1 MW, so it doesn't matter whether it's air-cooled or liquid-cooled. We can't get out more than 1.1 MW. We could put an air cooler on top of the roof and just throw part of the heat away and reduce our heat recovery to 40%, but that is not how we intend to work.

We build new systems when a new machine comes in; sometimes, we even build stuff midway just because we figure out that there's a better way of cooling things down, as in it costs less.

When we had a lot of air-cooled machines, we used to have what's called CRAC; they're big air conditioners, they work a bit like a car radiator. With these, you regulate temperature usually by opening or closing the valve that regulates how much flow of coolant should run through, but the fans always spin the same speed, and that's a 6 kW motor just spinning nonstop, so if you need to cool anything, you're still wasting 6 kW to pump a certain amount of l^3 /s of air around the room, is that effective? No, it's not; of course, it would be better to put in some extra money and get a frequency converter for the motor so it can spin up, on or down accordingly.

Another example would be encapsulating the machines by doing hot islands. You organise the racks in such a way that you separate the hot side from the cold side with a hard wall, and then you force a circulation pattern such that you don't have mixing because that's inefficient as well.

We try to optimise our systems all the time, and it's always a fight because you have to make investors believe that it's convenient to do this and that they will save energy in the long run. We do rebuild. What you need to change generally is as you get more dense machines, you want to get rid of more heat, so more often than not, what you just need is bigger pipes.

Are your growing capabilities, such as expanding the supercomputer, capped by how much heat you can output to the DC network?

As I said, we are capped at 1.1 MW in summertime. In the future, it depends on which kind of service you want to guarantee, not just cooling. If you offer a service, you usually offer it with a certain quality as well; for example, if you offer an uptime of 99%, you need to have the machines working 99 days out of 100. This means that you need to be able to cool 1.1 MW, relying on district cooling. If that fails, you need to be able to get rid of 1.1 MW of heat relying on your emergency cooling.

It doesn't end there; this means that if you just want to run the machines, if you lose the electrical network, you need to have UPSs and diesel for 1.1 MW, and that's not trivial. If one of your transformers catches on fire, you need another one which can still produce that power output. If it's summertime and, the outside temperature is hot, and you use compressors for your emergency cooling, you need to count that the best-case scenario is getting 25 l/s of water at 13°C. How much can you cool with that? Turns out you're somewhere around 1.2 MW. So we are capped, not only by how much heat we can get out but also by a number of things.

My answer to your question would be, here, as things are now, that is not the only limitation for us, to get more machines and more power, we would need to change more things about emergency cooling, about electricity availability or we should accept to offer a poorer quality of service. We could give 2.5 MW probably max out, but then we won't be able to provide power redundancy everywhere, and if there's a cooling emergency, then we would have to just shut down brutally, and in summertime, you can only use 35% of this amount because we still can't put out more than 1.1 MW.

A.2 Applications of BESS as an UPS

Have you thought about using your BESS as a grid stabilisation system as well as a UPS? Do you have other UPS systems?

Yes. We even collaborated with some people at the energy division about trying a project proposal; we just said that we would be available, but the project didn't come to fruition. We haven't done it for a simple reason: you need special hardware to manage the UPS. This hardware charges and discharges the batteries, and it sits between the load of the grid and the batteries. The batteries need to be able to back feed into the grid. You can't take a UPS and get it to back feed into the grid as long as you don't stabilise the frequency of your UPS. We can't do it because our system is from a time when such an application wasn't conceived.

We have these generators as well. So, if we wanted to, we could backfeed the network with these generators. We do not do it. We do not plan to do it either. There are many safety concerns regarding backfeeding the network, you need some special authorisations. Imagine that Stockholm Municipality sends a guy to work on an electric line, and they check that the line is disconnected, but we're back feeding on the upstream, which could cause a big accident, so you need some special authorisations if you want to put more than a certain amount of kilowatts. We don't do it at the moment; we thought about it, but it never happened.

A.3 Co-location

How relevant is the placement of the data centre in relation to the cooling facility?

That's a question I don't know if I can answer because I started to work here in February 2021, and obviously, the facility was already here. We're 100 metres away from where the next heat exchanges are. Of course, having it close to where the heat can get reused is an asset.

I sometimes give tours to students or curious people like you who want to have a look and learn a bit more about PDC. Also, the Nobel Foundation has asked us to come during the prize week and to put into perspective how much energy the computer uses I give this example: a hair dryer is 2 kW, and this is 800 kW, and people still don't get it. So, I say me and my girlfriend at home use the same energy in a year that the computer uses in 3,5 hours. But if you look at the biggest supercomputer in the world, called Frontier. The energy that me and my girlfriend use in one year at home, it uses that in 4 minutes! That's how much electricity goes into supercomputing resources. That supercomputer is placed in the middle of nothing; there's nowhere to introduce the heat whatsoever, so it gets blasted into the air.

In principle, you can dump heat wherever you want. If you want to reuse it, of course, it makes a lot of sense to have it close to a facility where it can be successfully utilised. You could build it in the middle of nowhere, boil some water with it, and have turbines turn it into electrical energy. But to start with, you're not going to boil water with a 40° C return. You need to put extra energy into it.

What I often say, especially to students, is to look at it from a Carnot cycle perspective. From a thermodynamics perspective, the efficiency of a computer is zero because it doesn't move anything. Our placement is very convenient, being this close means that nothing leaks out. It is quite important to have it close. I don't know if it was chosen for this reason, and honestly, I don't think so, because nowadays you can build whatever and you bring water wherever. But it surely is convenient.

A.4 Economics

What's the primary cost of the operation of the data centre?

That's a difficult question to answer. The machine itself costs a lot. When you buy it, you don't really buy the machine; you buy the machine and the service. For the purchase of the machine, we spent around 200 million SEK; this doesn't mean that we spent 20 million € on just hardware. A part of it is, but it also includes five years of people who will come and fix and replace parts for you. We also spend around 3 million € on energy per year. Think of it this way: after 5 to 8 years, it's inconvenient to keep running a supercomputer because it costs too much in terms of electricity.

It gets old, so it's better to get a new one. Otherwise, you would start losing money when compared to just buying a new one. After five years, the calculations per euro, per dollar, or per crown you get out of it start to get too little because it's just not energy efficient anymore.

Appendix B: Interview with Stockholm Exergi representative.

This appendix shows the transcription of an interview held with Stockholm Exergi's representative and PhD student Johan Dalgren on the 20th of February 2025 at his office on the KTH campus. The interview is divided into subjects which are of interest to the project.

B.1 Heat reuse

Who are the biggest heat contributors to the Open District Heating?

It's not a separate network. It's the common distributing network of heating and cooling. I think the largest the largest producer is Jästbolaget in Sollentuna. It's a factory that has a lot of excess heat. It's hard to distinguish what's Öppen Fjärrvärme and what's a normal production plant. If you just say everything that is not owned by Stockholm Exergi and producing heat into the system, that's Öppen Fjärrvärmet because they are using the business model, then we have a few large prosers or production plans. Crematories are also a source of a heat recovery,

If we talk about low temperature heat producers, we have supermarkets and data centers and ice rings. I think one of the largest customers there is Bahnhof, they have a few data centers. I think you have some references on this at Stockholm data parks. Also, ICA and Coop have a few supermarkets giving us the excess heat from the chillers.

Does Stockholm Exergi put restrictions on the amount of heat a proser can inject into the system?

Yes and no. I mean it depends on the conditions where they are at, if they are close to major pipes, it's easier to deliver a lot and if they're not, it's more difficult. But it is the infrastructure that puts the restraint on how much a proser can output.

Does the heat absorption capability vary during the year?

It depends on where you are and in summer, we have a lot more heat than we need, as heat is, of course, abundant. So, you don't get paid for excess heat in the summer. Since we have waste heat from incineration, you have waste heat from cooling production and so on.

What if efficiency levels are achieved when transferring heat from a DC to the grid?

That's quite a tricky question, because it depends on what your option, or your alternative is, right? So, if you only look within the walls, you could expect a COP of a heat pp of four or five.

But of course, as we talked about earlier, if you don't deliver heat to the district heating system, and use the heat source yourself, compared to if you would use the district cooling network and make that excess available for centralized heat recovery. If you have a heat source from a data center of 20-30 °C and you want to elevate that to district heating spot price delivery, as we talked about the heat that could be injected directly into supply needs to be at 65-80 °C so the efficiency would be around a COP of four or five.

Do district heating temperatures vary depending on the area?

It depends on the customers that you are supplying to, you need to fulfil the third-party agreement. So, if we at Stockholm Exergi promise our customer 65 °C in summer, of course, a heat supplier must fulfil those criteria as well.

But you could also mix temperature, so if we have an area where we have excess temperature that we can utilize, if a data center is connected to major pipe where we have excess temperature, the data center could supply a temperature of maybe 50 °C and the mixing temperature would still be above 70-80 °C.

Does your entire grid support Open District Heating?

We cannot inject heat at any point, but the problem is, is it meaningful and is it profitable to do it? if we have for instance a few transit pipes where we really want to move a lot of energy from one base plant, a major plant, to an area and where the distribution pipes in between set a limitation, a bottleneck. We are forced to have quite high temperatures to have a large temperature difference to move all this energy, here the supply temperature is close to 100 °C, just to get the transport, and it's not good there to mix in 50 °C water, that would destroy the distribution capacity.

But if you downstream this bottleneck, then on the other hand you have quite a big benefit from introducing these heat sources because then you could normalize the temperature without destroying the capacity. So, you can mix your 40 °C water with 100 °C water and utilize this excess temperature, while lowering the distribution losses.

How relevant is the location of the heat producer in relation to the DH grid?

If it's placed before a bottleneck, it's really bad for the system, but if it's on the downstream of a bottleneck, then it could be really good for the system.

How do you find suitable heat suppliers?

I think that there are a lot of studies, if you look at Heat Road Map Europe, you have Peta4, which is a web platform, where they have mapped potential excess heat producers and heat demand densities. So, you can map where you have a lot of excess heat and the waste heat potentials in terms of sewage water treatment plants, data centers and other sources. Of course, a lot of industry factories and so on. We also did a study I was part of where in Stockholm we tried this same approach, but in a higher resolution, and then we looked more in depth at what kind of data centers we have available, subway stations, ice rinks, and supermarkets. That project was called High-resolution mapping of the clean heat sources for district heating in Stockholm City.

Has Open District Heating helped or deteriorated grid stability?

In theory, it should help to have decentralized production, so that if you have a failure in a large plant, you would still have some kind of operation. So, in theory, it should help, in practice, I don't really know because there's a lot more technical things that could go wrong, you could have leakage with the refrigerant directly into the DH network of water. It's quite sensible to connect the condenser in the heat pump directly to the distributing network and it's not problem free.

If you we will have thousands of prosumers in the future, it will be very difficult to demand good enough leakage detection on the refrigerant side and on all the potential technical errors that could come from all these installations. It's a lot easier to monitor a few plants. So, I don't know, in theory, it should help, in practice, I don't know

B.2 Economics

How much energy is saved through implementing heat reuse?

It depends on, if you are taking you away the heat source from the district cooling network or are you like otherwise just wasting it. If you supply the cooling to the data centers, with district cooling, then the heat would be in the return of the district cooling network and come to the centralized plants where the heat would be recovered. So, it wouldn't be wasted. But if you have a data center that isn't connected to the district cooling, and then suddenly it recovers its own heat, then it's a huge benefit, of course.

So, it really depends on what the alternative is. Which I think is super important to understand that if you have wasted that energy before and now, you're doing a heat recovery, then it's great, but if you're removing the heat recovery from one part, and adding it through a new investment, maybe it's not that great from the system point of view.

But from a DC owners' point of view is better, because they get paid more, no?

I mean if they from a business model point of view can generate that heat at a lower cost than what they paid for the district heating, then it makes sense from their point of view. But giving it to the district cooling, makes that heat recovery is available for the entire city.

It depends on the CAPEX. I mean, it's very easy to just look at the OPEX. So, if you have a system where you more or less have no CAPEX yourself, and but you're paying both for the district heating and cooling then, of course, you have high operational cost, but virtually no CAPEX and then if you invest in your own heat recovery, then the OPEX becomes maybe negative. You get paid for delivering heat to the district heating network and you have free cooling, but you have made a huge investment.

What are the investment costs related to connecting to Open District Heating?

Normally the connections are even there, you have used district heating for some heating purpose, and then there is very little cost. But if you need to connect to a pipe that is far away, of course it could be super costly.

Who covers the infrastructure costs related to the connection e.g. piping?

Normally that would be covered by the prosumer so if you would like to produce and inject, then you need to do the investment that's needed.

Do you have to be a user of DH in order to become a prosumer?

No

How the does the compensation work?

It is hourly rated so there's a spot market so you have for every hour of the day, you see the marginal price in the system. It could be for this hour; it's 300 SEK and for next hour it's 350 and so on. So, it depends on the electricity prices, heat demands, is like the marginal price of the system and the district heating system. So, in summer, you won't get paid anything and, in the winter, when it's super cold outside, you could get plenty.

Appendix C: Interview with SDC representative.

This appendix shows the transcription of an interview held with SDC's sustainability and communication responsible, Ellinor Grape, on the 14th of February 2025. The interview is divided into subjects which are of interest to the project.

C.1 General/Company context

How does SDC differentiate from other data centre providers regarding service offerings, sustainability, and efficiency?

The biggest difference is the heat reuse; if you compare it to a lot of other data centres being built right now, especially for big companies, they don't have any heat reuse at all. Data centres use a lot of power, so a lot of heat is generated. So, I think that that that is our biggest advantage. Also, being close to the batteries, for now, we have the batteries up and running, and the data centre is going to start being built this year, but in the long run, also being able to use the batteries as backup power, for example, for the data centre and then hopefully not having to use the backup diesel generators as much, if it's ever needed. But you can also connect the batteries to the data centre. For now, while the data centre is being built, the batteries are used to balance the network instead because we have the allocated power. So, I think that is our most significant advantage when compared to others.

Is ScandiDC I your first working model?
It is our first site, that is correct.

What is the current state of ScandiDC I?

The BESS has been built and operational since November last year (2024), but we are going to start building the data centre this year. Due to the heavy investments needed for the DC, we are building it in steps, so you need to get clients and build it at the same pace. So, you don't build the entire thing at the same time.

What is the size of the first phase going to be?

In total, we're going to have around 10 megawatts of IT load. But how big the steps are, we have to see.

C.2 Economics

Do you think implementing heat reuse and BESS for grid stabilisation will reduce your OPEX?

It’s hard to say right now. We really hope so. By implementing heat reuse, the chillers won’t be necessary for regular operation of the DC. Of course, that is a win because we don't have that much energy needed for the data centre’s cooling. It's not quantifiable yet because it also depends on where you draw the line for what is owned by the DC or not. If the data centre owns the heat reuse heat pumps. So, I think that is a system question as well: is the heat reuse owned by the data centre or is it owned by the district heating network? So, I think it's a question of how you box it in in the end. But of course, the goal is that it should be cheaper, and it should be a win, win, where the data centre can sell the heat because you're also replacing something that costs money for the district heating company, for example, pellets or oil, there is a cost associated with that. The heat also has a price tag to it, but it should be cheaper, so it should be a win for the data centre and the district heating company, of course. From the operational side, in the end, we should be able to turn off the chillers, which is everyone's goal, because it would be cheaper for the distribution network to receive residual heat from the data centre. But it should also be cheaper for us, the DC, to just forward it to the district heating system instead of using the chillers. So, the answer would be yes, but it’s not quantifiable at the moment.

C.3 Heat reuse

What cooling technologies are you planning to use?

We are planning to use both air-cooled and water-cooled, but in the first step, it's just going to be air. But to make heat reuse profitable or worth it, it’s better to use liquid cooling; if it's just air-cooled, you get lower temperatures out. So, you want to use water-cooled systems. We're going to do it in steps. The first step will just be air, but then we will include water increasingly when building up the steps.

Once everything is finished, do you plan to change your air-cooled halls to water-cooled?
The ones that have been built with air will remain like that, you don't rebuild anything. But then also, like I said, it will be built in steps, so it's really hard to say exactly what will happen within three or five years. Also, the development of cooling systems within data centres is happening very fast, and we have to keep up with that as well. So, if we don't think that customers want air-cooled anymore, then we can't sell it, of course, but that will probably be further ahead.

As a collocation data centre, I'm assuming that the load on the IT equipment is going to vary depending on the client; how do you think this is going to affect heat recovery?

What we will target and want to have in the data centre is clients with high and steady use of the data centre. We usually take Högskolerådet, the ones that you do your application for university, as a bad example because they have a super low IT load and then when you do the applications, their load goes up and then down again. So that is a customer we don't want. We're targeting customers who have steadily high usage. Then, of course, we also know that there's a reality where we can't get everything we wish for, but that is what we're targeting because that will affect heat reuse. So, we're hoping to get a very high usage to get as much heat out as possible.

Where are you planning to repurpose the recovered heat?

The district heating network in Eskilstuna, so the site is located right next to the district heating line. What we're looking for, in general, is district heating because we want to have someone who can take all the heat that we produce. And sometimes, if you look at other things outside of district heating, they can take about one MW, but then we have to air out the rest. So, district heating is an exceptionally good heat reuse partner for us, but we also looked at greenhouses. Of course, it must be kind of big greenhouses.

What connection are you planning to make to the district heating network?

In order to supply to the district heating grid, you have to have heat pumps in between. The district heating network has a temperature we have to reach. That's why we have to have heat pumps in between, and then, depending on which time of the year, they will work more or less.

Does the district heating company have a limit on the amount of heat they can absorb?

I don't know if I can share the exact numbers from this case, but in general, towns have a heat demand that goes up and down. What we get is a curve that they need to fulfil, but it varies a lot between different cities.

Will you be able to repurpose the heat even in the warmer months?

Yes, so from our side, if they can receive the heat, we will send it to them. It will never be a limitation from us, but also, what we do when we evaluate locations for the data centre is that we look at the district heating or the greenhouse that we will cooperate with to see how much heat they can take, so that we don't have a big difference. That is very important for us. If it is in the middle of the warmest part of the summer, that differs a bit, but in general, we want to send all our heat today to the district heating system and in Eskilstuna, that is basically the case that we can do it, they will always be able to receive except in really special cases.

Are you open to connecting to other heat receivers, such as greenhouses?

Yes, greenhouses would be great. It will be another challenge, but we are open to anyone who can receive the heat basically. We know others who have done this with wood-drying facilities, which is also something that we have looked at, but we haven't had any real discussions about that. What is good with Sweden is that most cities in Sweden have district heating, so the playing field, if you want to work with that, is really big, so it's a good business case working with district heating.

C.4 Grid stabilisation

What are the advantages and disadvantages of using a BESS as an UPS when compared to more traditional systems like diesel generators?

The upside is that it's a more renewable source. You can load them with renewable energy, so it's renewable energy instead of burning fuel. That is the upside that we see. Then, one more thing is it could be a backup for the backup; you have two different systems that you can use, which I would say would be the biggest advantage. We still have to have a diesel generator because the batteries will run out if a long blackout were to occur. But for this generator, you can. You can fill that up in another way, so you have to have that as well. Our goal is to use the batteries in case of a smaller blackout.

If you use the batteries for grid stabilisation, how are you planning to maintain enough energy on them in case a blackout where to occur?

There has to be an agreement between the battery part of the company and the data centre because, as of today, the batteries are being used in the ancillary market. I don't think you want to take the battery completely out of the ancillary market or arbitrage, for example. I think there has to be a middle ground to keep the battery powered enough to run the servers for a certain amount of time. But then, of course, that comes with some kind of business case that is not 100% sorted out, how that could actually work. Also, the power needs of the data centre in the beginning will be low as it will be built in steps, so it will not be 10 megawatts, so the battery doesn't have to have that much power.

Is the BESS being built in steps as well?
No, in the case we have, they will stay as they are.

How does the grid stabilisation work? Do you do arbitrage?

That is one of the services, but then also. We support the grid with ancillary services for Svenska kraftnät. They are buying these services to balance the grid because the grid always has to stay at 50 Hz, and when it's not, someone is taking out a lot of energy or putting in a lot of energy, for example, then that has to be balanced out and that is where batteries can help. So, it is a very big energy market where you trade. That is what the batteries are doing if they're not supporting the data centre.

Do you get revenue from that?

Yes, it's like a stock market, and you can sell one MW in one hour because Svenska kraftnät needs that, and then you get paid for that. You can both load your battery, or you can put power into the grid.

Appendix D: Interview with T-Loop representative.

This appendix shows the transcription of an interview held with T-Loop's co-founder and deputy CEO, Staffan Styme, on the 21st of March 2025. The interview is divided into subjects which are of interest to the project.

D.1 General/Company context

What would you say are T-Loop's key competitive advantages?

I think TLOOP's competitive advantage is that we are looking at data centres from a different perspective. So, we look upon a data centre as both a way of providing computing but also of being an energy plant actually. And it maybe sounds easy, but it's not really, because it's two different sectors we are working in. This sector integration could be quite complex in terms of both engineering and design. You have different demands of how things should work in a data centre. You have an SLA that you need to fulfil and so on.

And you have almost always a playbook of how to build a data centre. And that may not always be the fit to produce energy. So, that's a technical part, but it's also a part that is about business culture and business model. To really get the full match between the energy sector and the data centre sector, you need to understand what the needs are and what the technical demands are from the energy side, from the distributing company, from the property owner, or from the power grid company. And you need to understand both the business model from an energy company point of view, what's the driving force for them, what's important for them. And you need to understand how they are building up their value in their business to get a really good match in an agreement.

I think, that T-Loop has both, we stand with one leg in data centres and one leg in the energy field. So, we talk, we can translate between these two and we understand both business and technology. And the needs of both sectors. So, I think that's the key competitive advantage that we have compared to other data centre operators, I would say.

How do you see the demand for sustainable data centres evolving in the near future?

I think it's growing all the time, quite fast. That's why we started T-Loop. We were frustrated of the untapped resources in data centres. And the data centre industry is quite young. It's about 20 years or so. And it's been under the radar for a long time from politicians and also from different kinds of authorities. And there's been a lack of transparency in that sector.

If I compare it to the energy sector where I come from, there we had so much more regulations and demands of transparency. I think that's one thing we see now that the data centre industry is growing so fast, it's an eye-opener for politicians and authorities that they need to do something. And so, now you have, for example, the regulation of the energy efficiency directive. That's a new directive that started this year actually, that you need to report key figures from data centres, actual energy use, actual how much do you reuse heat and so on. So, I think that's just the beginning to start knowing, really know, because it's really hard to get statistics today from data centres. But that's the first step, and then the next step will probably be that there will be different kinds of regulations of how you can build data centres and operate data centres.

And we already see that in Germany, for example, you need to have, if you build a new data centre, a PUE lower than, I think it's 1.3 or something. And you also need to prepare it for heat reuse and so on. I think that's a strong driving force, but I also see that it's, if you look on the customer side of the IT users, they are also beginning to get more and more aware about the problems that data centres can create.

Now bigger companies also are obligated to report according to CSRD, which means that you need to involve scope 3, and that means also that you need to report the carbon emission from the data you use, and that means also they need to ask the data centres operators and so on. Also, it's not just about carbon dioxide, I think, it's a new important infrastructure that needs to be built. And that requires more electricity, more capacity, and it requires more competence, more skills, more labour.

So, there’s also a kind of social perspective in this as well, how to build the data centres, and I think that data sets could be a piece in the puzzle for the green energy transition. Where you can use the untapped resources in batteries and backup power investments that you already are doing in data centres to help stabilise the grid, and therefore also make it easier to deploy more solar power and wind power and so on.

Do you only operate co-location data centres or do you work with other technologies?

Co-location, that's our main business, but it could also not necessarily be co-location. It could be a purpose-fit data centre for one certain customer, so it looks like on-premises. But we are more like a data centre as a service. It could also mean a wholesale or a single tenant, I mean, that we can sell, provide a data centre to another co-location operator. So, I would say we are a co-location provider, but it depends on what the customer wants, actually.

D.2 Economics

Do you think customers are willing to pay a premium for this, for having a sustainable data centre, or do you think it should come at the same price?

No, I don't think they are willing to pay more. Some customers are, maybe, but I think the key here is not to pay more. And I think what's really interesting in this case is that if you get more efficient, you can be more competitive, at the same time you are more sustainable. So, it goes hand in hand here. I think you will probably, if you are doing this in the right way, find a business model where you can actually earn money on selling heat and grid balancing. You can actually be more competitive rather than more expensive. But, of course, it's an added value with the sustainability aspect.

It's a premium product if you sell, and that could mean that you could be able to take some more payment, but I don't think that's possible today, actually. But maybe in the future, when it comes. I mean, we don't know. It could be some future regulation that says that you need to pay for the carbon dioxide you use, or something like that. And then, of course, it will be possible to pay more for a computing service that uses less energy and also less carbon dioxide.

Do you think that implementing heat reuse and grid stabilisation can reduce your operating costs?

Not the operating costs, I think, but definitely it will have new revenue streams. So, I think the operating costs would be similar, but the big upside is that you can create new revenue streams. Maybe even that the operating costs could be lower, if you compare it to a hyperscaler.

Our concept means that if you want to sell heat, then you need to be nearby a heat sink, of course. And that's urban areas, and usually you can't deploy several hundred megawatts data centre in an urban area. That means that we need to think smaller and more decentralised. And that means that maybe the operational cost would be actually a bit higher than if you build a very large-scale data centre somewhere in the outer region. But that's by far compensated by the increased revenue streams.

How does your relationship with the heat recipient work? What kind of agreement do you put in place?

It's different but it's a kind of negotiation because we are a customer to the property owner as well, we are a tenant and we pay a rent for placing our data centre there, but at the same time the property owner is a customer to us because we sell the heat to them. We don't have a standard model for that because it's depending on what the property owner prefers, I would say.

Sometimes we pay a full rent, and we charge the property owner about the same as they pay for just heating but if we pay less rent they may pay less for the heat as well, so it's quite a flexible model we have.

D.3 Heat reuse

What types of building or facilities benefit most from your heat reuse system?

Yeah, it's different aspects, I would say. Of course, we want to have a perfect match between the size of the data centre and the heating demand. So, that's one factor that you need to take into account. You don't want a very large data centre in a very small building. And the other part is if you want a really efficient heat recovery, of course, the building internal heating system is quite important. So, if they can use lower temperature, that's more efficient, of course. And then they also need, of course, the necessary infrastructure for building a data centre in the property. You need to have the power capacity, you need to have the fibre connections, and you need to have the space and so on. But if you have that and you have a large campus or a large commercial building, I think that's a perfect match.

When it comes to residential houses, we see that it's a bit more complicated because you need to also consider the distance to these kinds of installations. For example, when it comes to backup power, if you are going to use diesel generators, you need to consider if that could be a problem for residential areas. Or also security, for some customers, they wouldn't accept to have a data centre in a residential area.

The way your system works is that you mostly reuse heat in the same building that the data centre is based on?

Yeah, we reuse heat in the same building. That's our priority. But it's not always possible, for example, during the summer when the heat demand in the building is too low. The second best is to export the heat to a district heating network. If there isn't a district heating network nearby, you don't have so many options. You need to cool it away to outdoor air. But it's quite interesting if you could use some kind of thermal storage as well, seasonal storage, for example, geothermal storage. You can combine it with heat pumps and store the surplus during the summer and then use it for peak production during the winter.

During the summer, are you still able to export your heat to the district heating network?

Yeah, it depends on, of course, how much heat you produce and how large the district heating system is. Here in Stockholm, it's no problem because you have such a big heating need. It's a quite big system. But if you are in a smaller area, it could be a problem. So, it's the same thing that you need to match the data synthesised with the district heating system.

Is it still profitable, exporting the heat during the summer? I would say no. It could be in some places. It depends on how the energy production looks like. I mean, in smaller grids, it actually could be quite profitable because they often use a CHP plant, biomass plant, during autumn, winter and spring. But then they can't use it during the summer because the heating demand is too low. They need to shut it off. And then maybe use a pellet boiler. And pellets during the summer, that's quite expensive. So, in that case, it's profitable probably because you get more paid for the heat in the summer than during the winter.

But, for instance, here in Stockholm, you have the base production that you use is energy to waste. And that's a really low production cost. So, that means you need to compete with the cost that is almost zero. So, in that case, it's not profitable, but it might be better to do that than to cool it away because that's a cost as well.

Do you export the heat directly from the servers or do you heat it up even more with heat pumps?

It depends on what kind of IT load we have and what kind of cooling system we have and what kind of heat sink we have. So, for instance, in the proof of concept we have built and showcase that we use immersion cooling. That means that we can produce heat at a temperature of maybe 50-55 degrees. And that's about the same what the property energy internal heating system is designed for. So, there we have a perfect match and we actually can regulate the liquid in the racks. We can regulate that temperature so we have a perfect match at all time with the internal heating system. So, in that case, we can use it directly. But, it requires also that you have these kind of servers that can produce 55 degrees.

For example, HPC or AI servers that are also adapted to immersion cooling. But, for example, if you have a storage demand in IT, it's an overkill to use liquid cooling, you use air cooling. And then you get maybe 25 degrees or something and then you need a heat pump to lift it up. We are not saying that we always can use direct heat recovery. It's depending on the heat sink demand of the temperature and also the server load, I would say, and the cooling technology. But, we always try to optimise so we choose the most efficient way of recovering the heat.

When it comes to district heating, you need, 90% of the cases you need to have a heat pump to upgrade the heat. But, with future generations of district heating will probably be possible to do direct heat recovery.

Would you say that by having your heat recipient very near or in the same building as the data centre, you can get rid of the heat pump?

Yeah, you could, under certain circumstances, like in our showcase in Häggvik Solentuna, it's directly connected. But, during the summer there, we have a heat pump too, but then we export it to the district heating system.

With what cooling technologies can you get the highest output temperatures?

It's an interesting question because we thought for a long time it's immersion cooling that produces the highest temperatures. But, it's actually on-chip, with on-chip cooling you can get higher temperatures. But, with on-chip you're applying cooling specifically on the CPU, on the processor, and the processor temperature is about 80 degrees. So, you can probably get up, in theory, you can come up to maybe 70 degrees, but the hardware suppliers won't support that. So, but in theory, you could go up to maybe 60 or 70 degrees, I would say, with on-chip. But, that's only about 70 percent of the heat that is produced by the processor. Then you have heat dissipation from the rest of the server as well, and that's about 30 percent. And that must be cooled with air. So, if you mix these two temperatures, you will probably end up with maybe 40 or 45 degrees or something like that. But, if you can split these streams, you can have a high temperature of maybe up to, I would say, in theory, it could come up to 70 degrees. If you have a processor temperature of 80. But, in practice, you will have maybe 40 degrees, and that's because of the limitations that the hardware supplier sets on what temperature you need to have. And then, that's a high temperature stream, and then you will have a low temperature stream with a rest of 30 percent that is air cooled. On immersion, on the other hand, you get a mixed temperature from the beginning. You can reach maybe 55 degrees or something directly. So, I think that's a more elegant solution because you don't need any air cooling at all. You just put the server down in this liquid, and it's really quite simple installation. It is quite similar to what I'm used in the district heating business.

So, to run that immersion system, you should

Do you have a hydraulic system to run the immersion cooling system?

You have a CDU. You circulate. It's like a bathtub. You put the servers in, and then you have a circulation of liquid and you need to spread the circulation in, so you get a good heat dissipation of all the service and then you have a heat exchanger in a CDU and exchange it to a water-based system and then you have a heat exchanger to the internal system or to the distributing system.

Is heat exchange liquid-to-liquid more efficient than liquid-to-air?

Yeah, much more efficient, you get to have all these plate heat exchangers that are very efficient, and you may be losing maybe one or two degrees in temperature but not a lot more, so it's very efficient and they are quite small exchangers as well.

What percentage of heat is recovered from the servers?

It depends. From the servers, we can recover a hundred percent, I mean, but it's really about the receiver. How much, I mean, if you're connected to this district heating system, it's a hundred percent most of the times, I would say. But if you're only connected to a single building, we try to design it so it's minimum 70 percent.

The size of the data centre could be, could be limited by the amount of heat we can recover. But most of the cases, I would say it's the power connection that is limited, that is the bottleneck. It's not the heat sink.

D.4 Grid stabilisation

How do you contribute to grid stabilisation? Do you provide ancillary services or arbitrage?

We are not doing arbitrage at the moment, maybe we will do that in the future but it's ancillary services we have started with and we have developed our own software and hardware to do this, because to do this in a data centre that's mission critical, you need to be very sure about what you are doing so that you don't jeopardise the IT, so you need always to prioritise that, so that's very, very important. But when you have secured that you can do the ancillary services and use the overcapacity, so to speak, in the UPS system and batteries to provide these services.

Are your batteries in the UPS larger than you would need just to keep the service running?

It depends on what kind of ancillary services you want to provide, if it's very fast, FFR, for example, you don't need to oversize them but if you want to be in the FCR market you need to oversize them a bit, so yes, we are into the FCR market and we need to oversize them but it depends also if what you mean with oversizing because what you actually are doing is extending also the possibility of running, in case of an outage you will extend the time that you can run the data centre on, so you are actually extending also the redundancy so to speak, in the data centre by doing this, that's another benefit because it's very unusual that when you are providing and you are activated for FCR, like 20 minutes, it has never happened, it's usual is that it's maybe a few seconds that you are activated on the market, so in most cases this is a benefit for the data centre as well, but to guarantee the 7 or 10 minutes that it needs for the IT load, you need to oversize it.

What battery technology do you use?

Lithium.

Besides the batteries, do you have any other UPSs?

Yeah, in the showcase we are prepared for generators, but we don't have them. But we also will use the diesel generators, the gensets for performing, with balancing your more capacity service I would say, it's a slower service, it's another market, but we will be able to do that as well.

With the diesel generators?

Yes.

What conditions are set by the electricity provider for you to be able to inject electricity into the grid?

Here in Sweden it's different, a bit different depending on what market you are in, if it's the Swedish market we are talking about, you need to be approved, that means that you need to do an application, prequalification at a Swedish grid authority, Svenska Kraftnät, so you do need to do some tests that you are capable of responding in the way they specify, it's quite a complex test programme that you need to go through. You do the application and when you are approved, you are approved for participating in the market.

How does it work, if Svenska Kraftnät needs an X amount of energy to stabilise the grid, does he ask you for that energy?

It's a bidding process actually, so first you lay a bid two days before, that is based on a forecast, so you lay a bid that you can provide your service these hours, so you have 24 hours, 24 bids per day and to what price that you can provide this and it could be accepted and then you have another round the next day, so it's a bid process too and you do the same thing and if you are accepted, you are in the market and you don't know if you will be activated because you are activated automatically, you need to measure the frequency and if the frequency drops or goes up too high, the system kicks in, so it's an automatic activation.

And is the bidding also done automatically?

Yeah, it's automatic, I mean it's files that you send, but you need to, you can do it manually as well, put in the bids and send a file, but that's not practical, so what you do is that you set a bid and then we need to go through to BRP, Balance Responsible Party, so we are sending our bids to them and then they are aggregating all the bids for their customers and they send an aggregated bid to Svenska Kraftnät. And now, until a year ago, it was pay as bid, you get paid what you did buy for, but now, since last year, you get paid on the marginal price, so even if you are going on a very low bid, you get always paid for the highest accepted bid.

D.5 Co-location

And when it comes to the location of the data centre, what factors do you consider when choosing the placement?

It's both technical and market fit, I would say. It's both that you need to have power, you need to have access of land, of space, you need to have connectivity. That's the most important part. Also, we look at, can we apply our concept, is there a heat sink we can use? Is this location, is this attractive for customers, for IT customers? I would say that is the basic criteria, there are numbers of factors that you need to consider. I think we have a list of hundred different things you need to check before you can really be able to match different type of customers and so on.

TRITA – ITM-EX 2025:167
Stockholm, Sweden 2025

www.kth.se