HOLISTIC ENERGY MANAGEMENT AND THERMAL WASTE INTEGRA TED SYSTEM FOR ENERGY OPTIMIZATION

HEATWISE

Techno-economic analysis of cooling technologies for edge data centres integrated into tertiary buildings

Project ID: 101138491

Prepared by: Empa

December, 2025

e

Funded by the European Union

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Federal Department of Economic Affairs, Education and Research EAER
State Secretariat for Education, Research and Innovation SERI
EU Framework Programmes

Techno-economic analysis of cooling technologies for edge data centres integrated into tertiary buildings

Project acronym	**HEATWISE**
Project title	Holistic Energy Management and Thermal Waste Integrated System for Energy Optimization
GA Number	101138491
Call identifier	HORIZON-CL5-2023-D4-01
Topic identifier	HORIZON-CL5-2023-D4-01-04
Project duration	36 months
Related work package	WP7
Deliverable no.	D7.2
Dissemination Level	PU
Deliverable type	Report
Due Date of deliverable	31 December 2025
Completion date of deliverable	11 December 2025
Coordinator	H1 Systems
Website	www.heatwise.eu
Lead beneficiary of deliverable	Empa
Author(s)	Humbert, Gabriele Hansson, Jenny Koirala, Binod Prasad
Contributor(s)	Lilach Butchmits Atilla Morotz Seres Tamas Simon Pommerencke Melgraad Sascha Stoller Marco Kreyenbuehl
Reviewer(s)	Seres, Tamás

Keywords | Data Centre, Tertiary buildings, Waste heat recovery, Energy management systems, Optimal Design

Document history:

Revision	Date	Status
V0.1	14/11/2025	Initial draft
V0.8.1	21/11/2025	Draft for review
V0.8.2	01/12/2025	Version after internal review
V0.9	10/12/2025	Final draft after implementing comments, ready for editing
V0.9.1	11/12/2025	Final draft after editing, ready for approval
V1.0	15/12/2025	Final report

Acknowledgment

This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement for Project N° 101138491 and the Swiss Secretariat for Education, Research, and Innovation (SERI) under contract N° 23.00606.

Table of contents

Executive Summary

The increasing deployment of edge data centres (DCs) creates new opportunities for recovering and reusing waste heat within buildings and local energy systems. Within the HEATWISE project, this report contributes to Task 7.2 by analysing synergies and interdependencies between electricity, heating, and cooling vectors, with particular emphasis on how the temperature levels provided by DC cooling systems influence the integration of waste heat into tertiary buildings and building energy management systems. The objective is to assess and compare different cooling technology options numerically and to derive recommendations for their implementation in tertiary buildings.

To address these objectives, a techno-economic optimisation framework based on mixed-integer linear programming was applied to a building-scale multi-energy system integrating data-centre cooling, space heating and domestic hot water provision, chillers and heat pumps, thermal energy storage at multiple temperature levels, on-site photovoltaics, and connections to the electricity grid and district heating network. The optimisation considers system design and hourly operation over a full year under a perfect-foresight assumption. Performance is evaluated in terms of Total Annualised Cost, CO_2 emissions, and additional KPIs defined within the HEATWISE project.

Three cooling and integration concepts are assessed: (i) air cooling without waste heat recovery, representing the reference case where DC heat is fully rejected to the ambient environment; (ii) air cooling with waste heat recovery (≈25°C), where low-temperature heat is upgraded using an air-to-water heat pump for use in the building (35 °C); and (iii) two-phase direct-to-chip liquid cooling, which delivers high-temperature waste heat (65 °C) suitable for direct integration into medium- and high-temperature heating networks with limited additional conversion. These options allow the influence of cooling-system temperature levels on waste heat utilisation, system design, and cost-effectiveness to be systematically evaluated.

The results demonstrate that waste heat recovery from edge DCs is economically viable at the building scale. All configurations featuring waste heat recovery outperform the reference case without recovery, achieving lower total annualised costs while simultaneously reducing imported heat and associated emissions. Among the assessed options, liquid cooling consistently yields the best overall performance due to its high outlet temperature and low auxiliary electricity demand. Importantly, liquid cooling remains cost-effective even under conservative assumptions of doubled specific capital

expenditure per unit of capacity (up to 2474 CHF/kW), indicating robustness against investment uncertainty. Air cooling with heat recovery also proves economically attractive when evaluated at system level, despite requiring temperature upgrading via a heat pump.

The analysis further shows that high-temperature waste heat is not strictly required for cost-effective integration. In buildings with predominantly medium-temperature heating demand, such as space heating around 35 °C, lower-temperature waste heat can still be utilised efficiently, with only limited additional imports needed for domestic hot water preparation. For a building like NEST, recovering heat from a direct-to-chip technology at temperatures of ≈35 °C was still predicted as economically viable. Beyond cooling technology choice, system performance is strongly influenced by the matching between DC size and building heat demand. Increasing DC capacity raises the fraction of demand that can theoretically be covered by waste heat but reduces the fraction that can be effectively recovered, resulting in diminishing returns and, under certain conditions, higher total emissions due to increased electricity consumption. Thermal energy storage is systematically selected as a cost-effective means to mitigate short-term mismatches between waste heat availability and demand, particularly in mid-season periods, but it cannot address structural seasonal imbalances at the scale of a single building.

Location-specific factors play a decisive role in shaping optimal system configurations. Comparisons between pilot sites with different climates, heating demand profiles, and energy price assumptions show that colder locations with more stable and higher heating demand can achieve higher recovery fractions, while milder climates with stronger seasonal variability rely more on medium-temperature storage and interaction with district heating networks. These results confirm that waste heat recovery performance cannot be generalised and must be assessed within its local climatic and economic context.

Overall, the study confirms the central premise of Task 7.2: cooling technologies for high-power-density data centres must be evaluated as integral components of multi-energy systems rather than in isolation. The findings provide quantitative guidance on how cooling-system temperature levels, costs, and efficiencies translate into system-level benefits and identify liquid cooling as a particularly promising solution for future edge DC deployments in tertiary buildings. The optimised configurations presented here represent idealised benchmarks under perfect foresight and provide a foundation for subsequent HEATWISE activities focusing on dynamic modelling, energy management system

implementation (T7.3), and real-world operational constraints, where forecast uncertainty and control limitations will further shape achievable waste heat recovery.

List of abbreviations

Acronym	Meaning
AI	Artificial Intelligence
AC	Air-cooling
BEMS	Building Energy Management System
DC	Data Centre
DHN	District Heating Network
DtC	Direct to Chip
EHub	Energy Hub
EMS	Energy Management System
HDD	Heating Degree Day
HPC	High-Performing Computing
IT	Information Technology
KPI	Key Performance Indicator
LC	Liquid Cooling
NEC	Normalized Energy Cost
NOE	Normalized Total Emissions
NTC	Normalized Total Cost
OE	Operational Emissions
PBT	Pay Back Time
RES	Renewable Energy Sources
REU	Renewable Energy Utilization
TAC	Total Annualized Cost
TE	Total Emissions
TES	Thermal Energy Storage
TESU	Thermal Energy Storage Utilization
WHR	Waste Heat Recovery

1. Introduction

Driven by the rapid growth of digitalization and cloud-based services, electricity consumption by European data centres is expected to rise sharply, increasing from about 96 TWh in 2024 to around 168 TWh by 2030 and reaching roughly 236 TWh by 2035 [1]. This corresponds to an almost 150 % increase, making data centres one of the fastest-growing electricity consumers in Europe, surpassing electric vehicles and approaching the growth pace of electrified industry [1].

Much of the electricity consumed by data centres is converted into heat [2], and if this thermal energy is appropriately recovered and utilised, it represents a substantial waste-heat recovery potential. For instance, recent reviews highlight extensive opportunities for recovering low-grade waste heat from data centres [3], [4]. Incentivised by regulatory frameworks [5] promoting energy efficiency and circular heat use, waste-heat recovery from large-scale data centre facilities is increasingly adopted, particularly in locations proximate to district-heating systems [6]. As a result, the placement of large-scale data-centres often gravitates toward urban or peri-urban zones with access to heating infrastructure. However, this concentrated growth also contributes to electricity-grid congestion: according to [1], some new European data-centre developments face grid-connection wait times of up to 13 years.

The rapid expansion of large-scale data centres, while enabling digitalisation and cloud services, also introduces challenges related to electricity grid congestion, site availability, and infrastructure scaling [7]. In parallel, increasing requirements for low-latency communication, enhanced data security, and reduced bandwidth usage are driving interest in more decentralised computing architectures. This has led to the emergence of *edge computing*, in which data processing is performed closer to the location where data is generated or consumed rather than in remote hyperscale facilities [8], [9]. Edge computing typically relies on *small-scale data centres*, often in the order of tens to hundreds of kW, installed within or near end-user facilities. Although hyperscale data centres continue to dominate the market, the edge segment is expanding rapidly, with an estimated compound annual growth rate of approximately 25.7 % between 2025 and 2033 [10].

In edge computing, small-scale data centres are deployed within or adjacent to buildings, thereby creating new opportunities for integration with local energy systems. Building integrated edge data centres could also increase the self-consumption of locally generated renewables. The close physical proximity between waste-heat sources and

potential thermal end-uses significantly reduces distribution losses and can lower the required supply temperature [11] as fewer heat-transfer steps are involved compared with large-scale data centre recovery schemes. This proximity enhances the technical feasibility of direct heat reuse for space heating or domestic hot water production. However, the temporal mismatch between waste-heat availability and building heat demand remains a major challenge. When considering the building as an integrated energy system, where electricity consumption, IT workload, and thermal demand interact dynamically, coordinating these energy flows to ensure reliable operation and minimal overall cost becomes a complex optimisation problem.

In this report, the optimal design and operation of energy conversion and storage technologies installed within tertiary buildings hosting edge-integrated data centres are studied. Different data-centre cooling systems and heat-recovery solutions are analysed across several use cases and scenarios. In the following subsections, the cooling technologies applicable to edge data centres are reviewed and the technologies considered in the study are outlined. Furthermore, an overview of modelling approaches for waste-heat recovery from data centres is provided, along with a brief description of typical heating systems in the built environment. Finally, the existing knowledge gaps and the specific aims of this report are presented.

1.1. Cooling technologies for data centers

Heat management in data centres is a critical aspect of ensuring the safe and reliable operation of IT equipment and extending its service life. As almost all the electricity consumed by servers is ultimately converted into heat, effective cooling systems are essential to maintain allowable operating temperatures and prevent thermal stress [12] [13]. Cooling design has a strong influence on both energy efficiency and the potential for waste-heat recovery.

Figure 1 summarises the main families of data-centre cooling technologies, classified as:

- **Air cooling** relies on chilled air to remove heat from servers, typically using computer-room air conditioners (CRAC), air handlers (CRAH), or rear-door heat exchangers (RDHx).
- **Direct-liquid cooling** circulates dielectric or water-based fluids in direct contact with heat-generating components, achieving higher heat-transfer efficiency.
- **Indirect-liquid cooling** systems employ intermediate loops such as heat pipes or cold plates, which separate the coolant from the IT components.

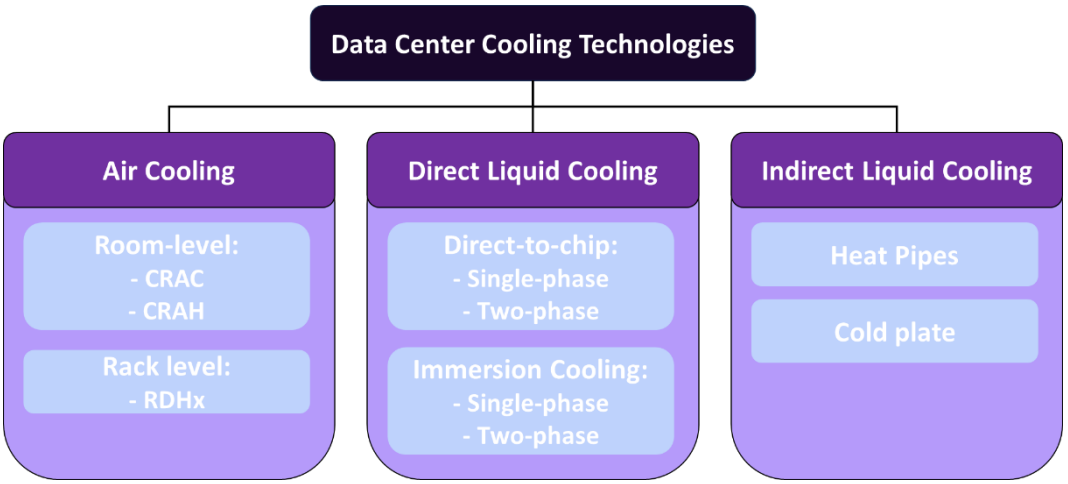

Figure 1 Overview of cooling technologies for data centers, partially adapted from Hao et al. [2].

The key features of the main cooling technologies are summarised in

Table 1, while schematics for the layouts are reported in Figure 2.

Table 1 Summary of the key features for the main cooling systems for data centers.
(*) the outlet temperature reported indicates the maximum delivered temperature temperature, with the reported range defined by sub-technology and manufaturers.

Cooling family	Configuration	Working principle	Typical T_outlet (°C) (*)	Advantages & Limitations	Typical Application	Ref.
Air Cooling	CRAC	Uses direct expansion (DX) refrigeration to cool and circulate room air through raised floors or ducts.	20–25	Simple and reliable but low efficiency and uneven heat exchange.	Small / Medium	[14] [15]
	CRAH	Circulates room air through coils connected to a chilled-water loop.	25–30	Higher efficiency than CRAC, but requires chilled water and large space.	Medium / Large	[14] [15]
	RDHx	Hot air from servers passes through an air-to-water heat exchanger mounted at the rack's rear door.	30–40	Compact and retrofit-friendly; moderate recovery temperature and airflow resistance.	Medium / High density / Edge	[14] [15]
Direct Liquid Cooling	Direct-to-chip (single-phase)	Liquid flows through cold plates mounted on chips (CPU/GPU).	45–55	High efficiency and compact; requires plumbing and careful sealing.	Large / Edge	[16], [17]

	Direct-to-chip (two-phase)	Dielectric fluid evaporates at the chip and condenses in a secondary exchanger.	55–65	High-grade recovery and uniform temperature; complex and costly.	Large / Edge	[16], [17]
	Immersion (single-phase)	Servers immersed in dielectric liquid; heat removed via circulation.	50–60	High power density and uniform cooling; requires hardware adaptation.	Supercomputing / Edge	[16], [17]
	Immersion (two-phase)	Servers immersed in boiling dielectric liquid; vapor condensed by facility water.	60–70	Very high recovery potential; costly and less mature.	Supercomputing	[16], [17]
Indirect Liquid Cooling	Cold plate	Heat is transferred from the chip to the intermediate plate, then to the liquid circuit.	40–55	Reliable and flexible; air still needed for some components.	Large	[16]
	Heat pipe	Phase-change fluid in sealed tubes transfers heat to a remote radiator.	35–45	Passive and low maintenance; limited scalability.	Small / Medium	[16]

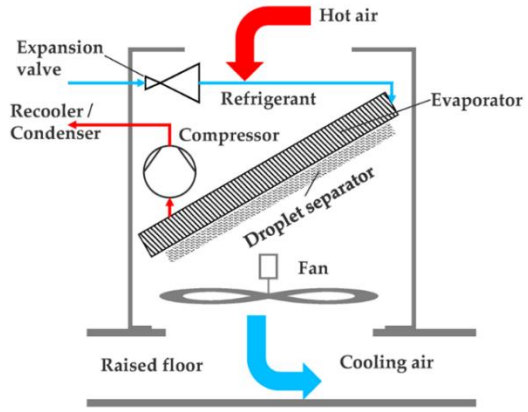

(a) Computer Room Air Conditioner (CRAC).

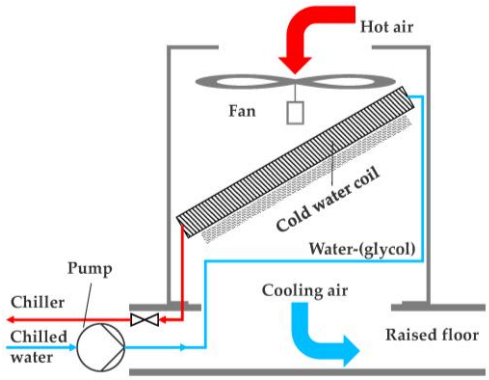

(b) Computer Room Air Handler (CRAH).

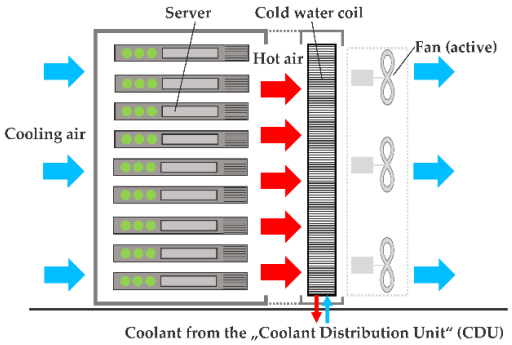

(c) Rear Door Heat Exchanger (RDHX, ILC).

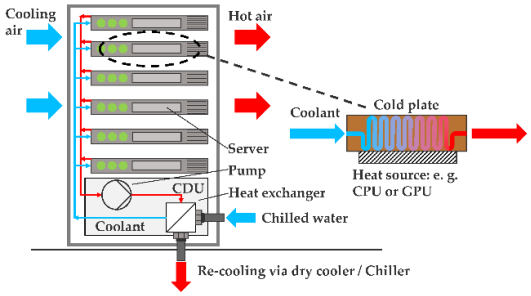

(d) Direct Liquid Cooling (DLC) with cold plates.

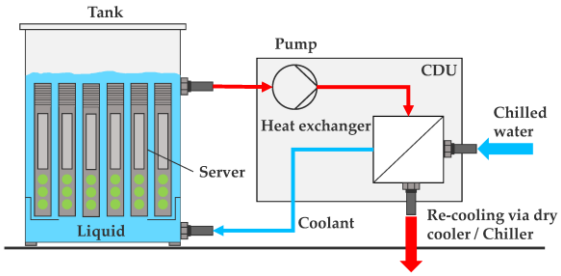

(e) Immersion liquid cooling with single-phase coolant.

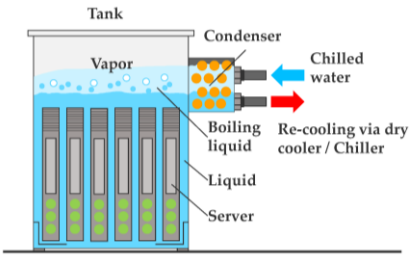

(f) Immersion liquid cooling with two-phase coolant.

Figure 2 Layout of the main cooling technologies for data centers from [18].

Air-based systems remain the predominant cooling approach in current data-centre deployments [16], with industry estimates suggesting that as much as 80 % of facilities still rely mainly on air cooling [19]. However, liquid cooling adoption is rising rapidly, driven by the growth of AI and HPC workloads, increasing rack power densities (Figure 3), and the potential for high-temperature waste-heat recovery. Several market analyses project that liquid cooling will reach a significant market share in new installations before 2030, and more patents related to cooling systems after the year 2000 are related to liquid cooling systems than air cooling ones [18].

Liquid- and air-based systems are frequently deployed in hybrid configurations, as only a subset of components, typically CPUs and GPUs, can be efficiently cooled by liquid-based technologies. Memory modules, storage devices, and power electronics still require airflow, meaning that even liquid-cooled racks generally retain fans and an air-management subsystem [20]. Hybridisation is therefore the norm in practical installations.

Cooling technologies also differ in the working fluids they employ. Single-phase direct-to-chip systems typically use water or water–glycol mixtures, whereas immersion cooling relies on dielectric liquids such as fluoroketones, hydrofluoroolefins, or synthetic oils [17]. Two-phase systems use engineered fluids with tuned boiling points, enabling evaporation at the chip surface and condensation in a secondary exchanger [17] [16]. These choices influence heat-transfer performance, operational temperature levels, and the quality of recoverable waste heat.

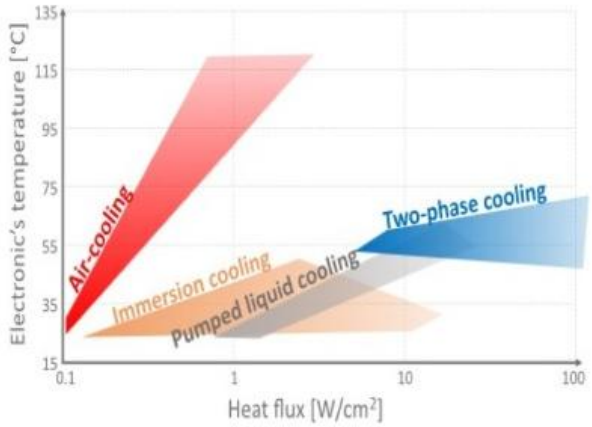

Figure 3 Junction temperature corresponding to different heat flux from Hao et al. [2].

Cooling requirements differ markedly between hyperscale and edge data centres. Hyperscale facilities typically integrate large-scale chiller plants, extensive free-cooling infrastructure, and CRAH-based air systems, benefiting from economies of scale and advanced thermal-management controls. Edge data centres, in contrast, operate within space-constrained environments such as tertiary buildings, often without access to central chilled-water infrastructure. As a result, they rely on compact and modular systems such as CRAC units, rear-door heat exchangers, and direct-to-chip liquid cooling. These technologies are well-suited to heat recovery due to their higher outlet temperatures and physical proximity to potential heat sinks within buildings.

1.1.1. Technologies selected for the study

Edge data centers, as considered in this report, generally operate in the tens to hundreds of kilowatt range. For example, the Eco-Qube project defines edge data centers as a size range up to 250 kW [21]. They are typically located within or adjacent to existing buildings and serve local users or devices. While hyperscale data centers prioritize economies of scale and often employ large chiller plants or cooling towers, edge sites favor compact, modular cooling systems that can be integrated with the building's HVAC and heat-recovery infrastructure.

The cooling technologies for edge data centers selected for this study are based on the HEATWISE pilot installed at the NEST building [22] in Dübendorf, Switzerland, represented in Figure 4. NEST has a typical winter heating load of 50 kW and is also connected to the district heating network of the Empa campus. It hosts an open hardware single rack air-cooled micro data centre from Horizon Europe Eco-Qube project, with 12 kW of IT load. The air-cooled rack incorporates a closed-loop cooling unit with an integrated air-to-water heat recovery coil, operated in accordance with ASHRAE thermal guidelines. Hot return air from the rack first passes through the recovery coil, where low-grade heat is transferred to a water loop, before additional cooling is provided through a downstream cooling coil if required. Adjacent to this, a new rack with 20 kW theoretical peak IT load is installed, equipped with 40 refurbished DELL PowerEdge R640 servers. Both systems are already integrated into the building's thermal and electricity grids. The pilot includes two representative solutions: an air-cooled rack equipped with rear-door heat exchangers, and a direct-to-chip two-phase liquid-cooling system provided by Zutacore [23].

Drawing on technical insights from the installation and commissioning of the pilot, this techno-economic assessment focuses on quantifying the waste-heat recovery potential of:

- **an air-cooled rack with rear door heat exchangers;**
- **a hybrid liquid-cooled rack with a two-phase direct-to-chip system supported by residual air cooling.**

This choice also reflects air-cooling as the most common state-of-the-art cooling technology and liquid cooling as innovative future alternatives. The techno-economic parameters adopted for the investigated cooling technologies and additional components are detailed in Section **Hiba! A hivatkozási forrás nem található.**.

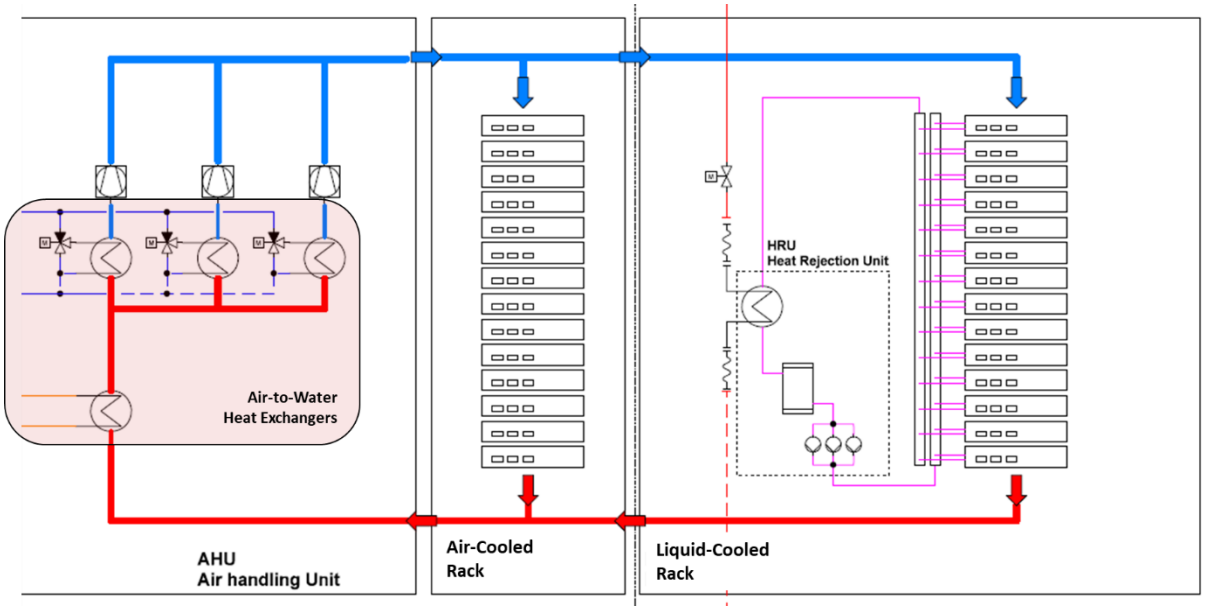

Figure 4 Layout of the cooling technologies installed for the edge data center HEATWISE pilot in the NEST building in Dübendorf, Switzerland. The system consists of two enclosed racks, one air-cooled and the other liquid and air-cooled. Air is circulated within each rack by internal fans, and the resulting warm exhaust air is cooled via an air-handling module equipped with air-to-water heat exchangers connected to water loops operating at different temperature levels. In one of the racks (HEATWISE pilot), a direct-to-chip two-phase liquid-cooling system from Zutacore is also installed [23]; the thermal energy recovered from this loop is transferred to a high-temperature hot-water circuit for efficient heat reuse.

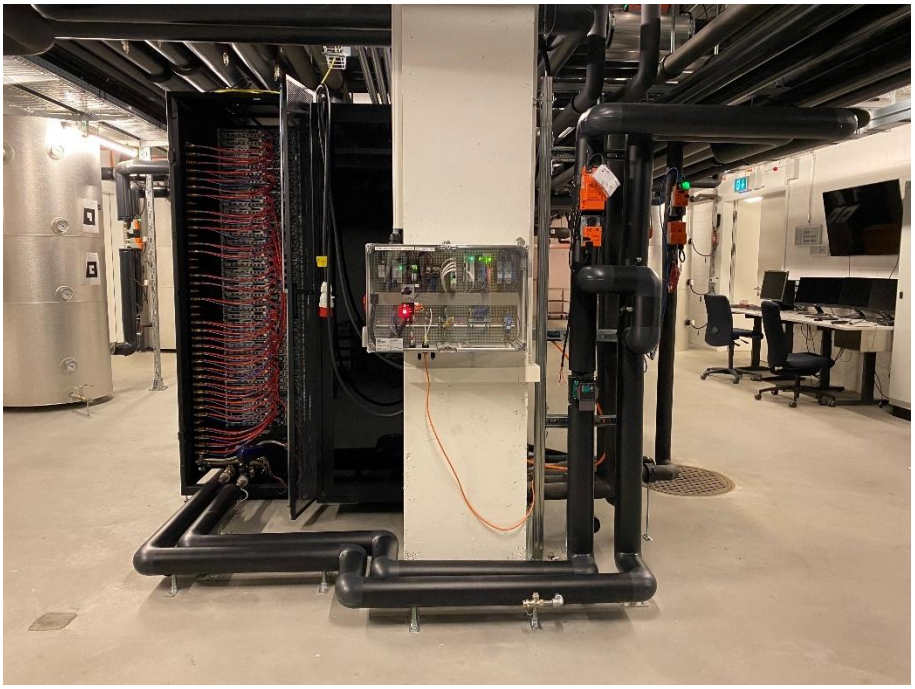

Figure 5 Picture of the liquid-cooled [23] servers integrated into the NEST building [22].

1.2. Heating systems in the built environment

Heating in European buildings is still predominantly supplied by fossil-fuel systems, particularly gas and oil boilers, despite increasing policy pressure to decarbonise the sector. For example, almost 60% of the buildings in Switzerland are still heated by oil and natural gas [24]. According to recent assessments of the EU building stock, combustion-based heating represents the majority of installed systems across both residential and tertiary buildings, while the penetration of heat pumps and low-temperature district heating networks is growing steadily as part of national and EU-level decarbonisation strategies [25].

Several types of heating systems are commonly found in European buildings:

- *Gas and oil boilers*: Widely used in existing buildings, especially older stock. These systems typically operate at high supply temperatures, ranging from 70 to 90 °C, to meet the requirements of radiators and domestic hot water preparation. However, as can be seen in Figure 6 for the instance of Switzerland, new buildings no longer feature gas and oil boilers.
- *Heat pumps (HP):* Including air-to-water, ground-source and water-source technologies, increasingly used in both new and renovated buildings. For the instance of Switzerland, in the year 2023, 70% of the HP sold were air-sourced, 26% were ground-sourced [26] Heat pumps generally supply at lower temperatures, typically 35–65 °C, depending on the emitter system (e.g., underfloor heating or low-temperature radiators).
- *District heating networks (DHN)*: Common in Northern, Central and Eastern Europe, DHNs supply thermal energy to multiple buildings from centralised plants. Conventional networks operate at high temperatures (80–95 °C), but many cities are transitioning toward fourth-generation networks with supply temperatures between 50 and 70 °C to improve integration of renewable and waste-heat sources.

These systems exhibit a wide range of supply–return temperature requirements, which directly affects their compatibility with waste-heat sources. In conventional high-temperature boiler or district-heating systems, recovered heat from data centres, typically 30–40 °C from air-based cooling and 45–65 °C from liquid-based systems, requires temperature lifting via heat pumps [27] In contrast, modern low-temperature heating systems [22] are much better aligned with the temperature levels achievable through direct-liquid-cooling technologies, making them attractive candidates for direct or assisted heat recovery.

As European heating systems evolve toward lower supply temperatures and increasingly integrate renewable and residual heat sources, waste-heat recovery from data centres represents a growing opportunity to reduce fossil-fuel consumption and support local decarbonisation objectives [11].

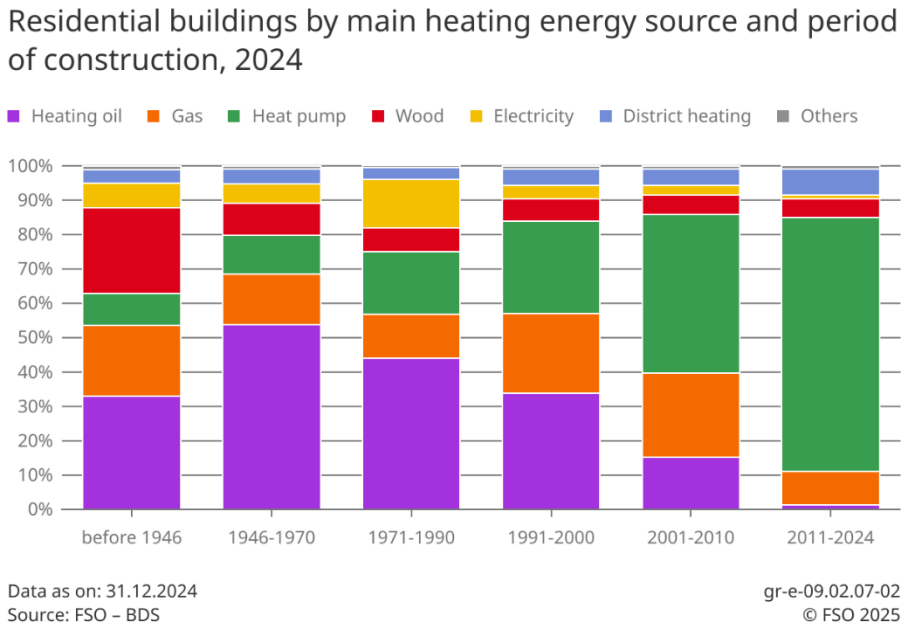

Figure 6 Main heating systems in residential buildings in Switzerland versus the period of construction [24].

1.3. Knowledge gaps and research questions

Despite growing interest in using waste heat from data centres, several gaps remain in the current scientific and technical literature:

- **Limited modelling approaches for integrated energy–ICT systems:** typical energy-planning and building-optimization tools rarely include detailed representations of data-centre cooling technologies or their interaction with building energy systems. Cooling-system performance, waste-heat temperature levels, and operational dynamics are often oversimplified or omitted, preventing realistic assessments of integration opportunities.
- **Lack of holistic techno-economic assessments**: although the choice of cooling technology strongly influences both energy efficiency and the quality of recoverable heat, few studies evaluate these systems within a full building context. The combined effects on electricity consumption, heating demand, heat-pump operation, investment costs, and operating expenditures are rarely assessed together. As a consequence, the cost-effectiveness of advanced cooling solutions,

such as direct-to-chip or two-phase echnologies, remains insufficiently understood, particularly for edge data centres embedded in tertiary buildings.

- **Insufficient analysis of demand–supply interactions**: modern energy systems are increasingly shaped by fluctuating electricity prices, variable renewable generation, and dynamic IT workloads. However, studies typically evaluate data-centre waste heat using annual or coarse time-resolution data. High-resolution temporal modelling is necessary to capture the alignment between waste-heat supply and heating demand, assess storage requirements, and quantify the operational value of heat reuse under realistic market conditions.
- **Limited evidence and studies for small-scale and edge data centres:** most of the existing research focuses on hyperscale and cloud facilities, while edge data centres, often installed in space-constrained buildings and operating at tens to hundreds of kilowatts, remain underrepresented. Their integration potential, constraints, and performance characteristics differ substantially from large-scale infrastructures, yet remain poorly documented.

Together, these gaps highlight the need for systematic, high-resolution, techno-economic analyses capable of capturing the complex interactions between cooling technologies, data-centre operation, and building energy systems. The present report aims to address part of this gap by evaluating cooling configurations and waste-heat recovery strategies for edge data centres integrated into tertiary buildings.

According to the knowledge gaps listed above, the following research questions were formulated and are addressed within this report:

- **RQ1**: what is the techno-economic potential of recovering waste heat from edge data centres integrated into tertiary buildings?
- **RQ2**: how do different cooling technologies and design parameters influence the overall performance of the energy hub?

The overarching goal is to provide guidance for technology developers and energy planners on desirable efficiency levels, waste-heat temperatures, and integration requirements.

2. Methodology

The goal of the modelling framework used in this study is to assess the optimal integration of edge data-centre cooling technologies within building energy systems, with a particular focus on waste-heat recovery and multi-carrier energy interactions. To achieve this, we employ the EhubX tool, a modular techno-economic optimisation tool developed at Empa [28]. The EhubX tool builds on the Energy Hub formulation [29], representing buildings or districts as multi-energy nodes where conversion, storage, and distribution technologies interact to satisfy time-varying energy demands (Figure 7).

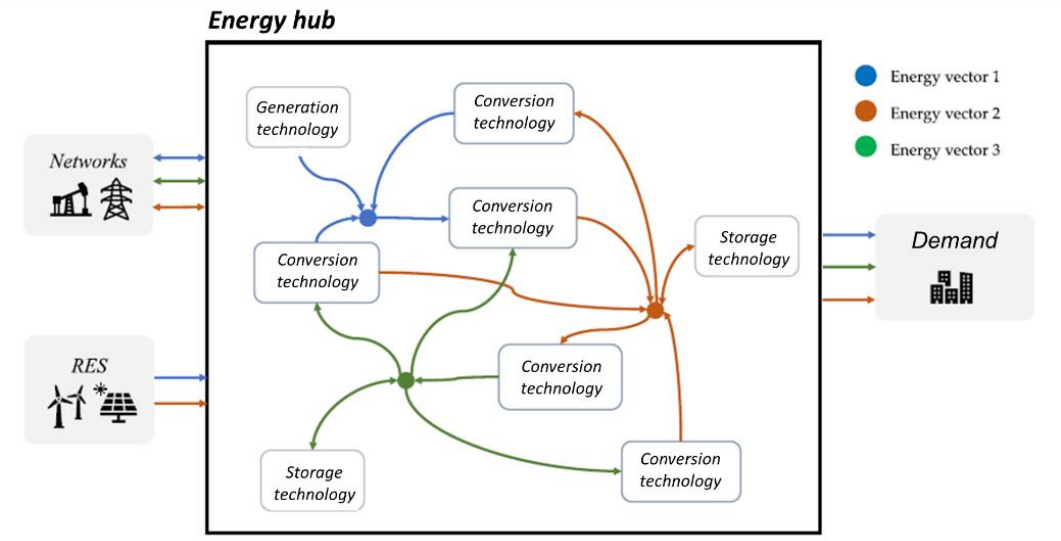

Figure 7 The Energy Hub concept, partially adapted from Marzi et al. [30].

EhubX couples a representation of energy-conversion and storage technologies (e.g., heat pumps, boilers, chillers, photovoltaic systems, batteries, thermal storage, and data-centre cooling units) with a mixed-integer linear programming (MILP) formulation to optimise system operation and design simultaneously. The tool identifies the optimal configuration of technologies and operational schedules for cost or CO_2 minimisation problems and constraints such as equipment capacities, thermal and electrical balances, and technology-specific operating limits.

This approach allows capturing the temporal variability of electricity prices, energy demands, renewable energy availability, and IT workload while explicitly modelling interactions between the data centre and the building's thermal system. Figure 8 provides a high-level representation of the modelling workflow.

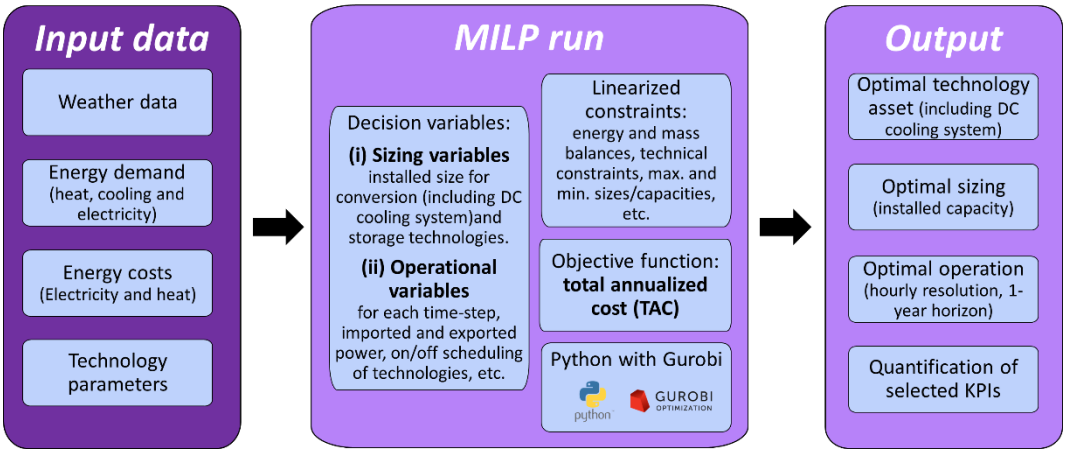

Figure 8 Overview of the MILP design optimization strategy adopted for the techno-economic analysis.

2.1. MILP for the optimal design of energy systems

The design optimisation problem in this study is formulated as a mixed-integer linear program (MILP), which simultaneously determines the optimal sizing and operation of the energy technologies within the building energy hub. MILP is widely used for energy-system planning because it allows capturing discrete design choices (e.g., equipment installation), time-dependent operational decisions, and multi-carrier energy flows while maintaining computational tractability [31].

In the EhubX framework [28], the MILP model optimises the system over a full year with hourly resolution. At each time step, the model must ensure that the energy demands (space heating, domestic hot water, cooling, electricity, and data-centre IT load) are satisfied while respecting equipment limits and energy-balance constraints. The optimisation evaluates both the air-cooled RDHx configuration and the two-phase direct-to-chip liquid cooling configuration [23], enabling a consistent comparison of the waste-heat recovery potential.

The MILP structure consists of:

- a set of **continuous and integer decision variables** describing energy flows, technology outputs, storage operation and investment decisions. These decision variables can refer to both **sizing and operation** of the energy system;
- a **linear objective function**, described in Section 2.2.3, which minimises total annualised system cost or operational CO_2 emissions;
- a set of **linear constraints** ensuring thermodynamic feasibility, capacity limits, and compliance with system-level energy balances, e.g. fulfilment of energy demands.

The following subsections describe the required input data, decision variables, objective formulation, and constraints.

2.1.1. Decision variables

Over the optimisation horizon [$t_0 = 0\ hours, t_{end} = 8760\ hours$] with hourly resolution, the MILP model simultaneously determines the optimal technology portfolio and its operation. The decision variables can be grouped into investment and operational variables.

Investment decision variables

Investment decisions define which technologies are installed and their size. For each candidate technology *i* (PV, heat pumps, storage units, data-centre cooling and heat-recovery components, etc., as detailed in Section 2.6), the model includes:

- **Installed capacity, Cap_i**: a continuous variable representing the installed capacity of technology *i* (in kW or kWh), bounded by user-defined minimum and maximum values. When a minimum capacity value of 0 is selected, the optimizer can decide to exclude a technology in the final optimal design when it is not cost-effective; that is, the optimizer, among a design space of selected technology candidates, identifies the optimal technology asset for a given engineering problem. Furthermore, the installed capacity can also be fixed a priori and, therefore, is not a decision variable of the optimization problem.

These investment variables influence the objective function via annualised CAPEX and capacity-related OPEX.

For the edge data centre, the IT load and rated IT capacity are treated as exogenous inputs, while the capacities of the cooling and heat-recovery components (e.g. rear-door heat exchangers, liquid-cooling loops, building-side heat exchangers) are optimised. Additional details regarding the integration of data centers models into MILP energy planning tools are provided in Section 2.2.

Operational decision variables

Operational variables describe, for each time step *t*, how the installed technologies are dispatched to meet energy demands and how waste heat from the edge data centre is used. They include:

- **Technology dispatch:** for each technology, *i,* and energy carrier, *ec,* (electricity, heat, etc.), $P_{i,e}(t)$ represents the power delivered or consumed at time *t*. In this report, we adopt the letter P for electrical power and the letter Q for heat.

- **Storage operation**: For each storage technology, *s,* (battery and thermal storage in this work), we define:
    - Charging power: $P_{ch,s}(t)$ for battery and $Q_{ch,s}(t)$ for thermal storage
    - Discharging power: $P_{dis,s}(t)$ for battery and $Q_{dis,s}(t)$ for thermal storage
    - State of charge: $SOC_s(t)$
- **Grid and DHN exchanges:** at each time step:
    - Electricity and heat imported from the grid/DHN: $P_{imp}(t)$ and $Q_{imp}(t)$
    - Electricity and heat exported to the grid/DHN: $P_{exp}(t)$ and $Q_{exp}(t)$
- **Waste-heat utilisation from the edge data centre:** the IT workload and resulting heat generation are fixed inputs. The optimisation determines how this heat is handled:
    - Heat recovered and injected into heating networks: Q_{DCtoHP} (t) and $Q_{DCtoHT}(t)$
    - Heat rejected to the cooling network: $Q_{DCtoLT}(t)$

These variables couple the data-centre cooling system with the building's energy hub and allow the model to quantify trade-offs between heat reuse, auxiliary cooling electricity, and overall cost.

The operational variables influence the annualised energy cost and, indirectly, the annualised CAPEX. Efficient operation, for example through smart charging and discharging of storage, can reduce peak loads and therefore lower the required installed capacity of other technologies, such as chillers or heat pumps.

2.1.2. Objective function

The optimisation problem addressed in this section is formulated as a cost-minimisation problem, as this represents the primary focus of the analysis. While this work concentrates on economic optimality, the EhubX tool is also capable of addressing emission-minimisation and multi-objective optimisation problems.

The model determines the set of technologies to install, their capacities Cap_i and their hourly operation to minimise the Total Annualised Cost (TAC) of the energy hub over one representative year:

$$TAC = C_{inv,a} + C_{maint} + C_{energy}$$

Where:

- $C_{inv,a}$ is the annualised investment cost;
- $C_{maint,a}$is the annual maintenance cost;
- C_{energy} is the annual energy-related cost (imports minus revenues from exports or heat sales).

For each technology, i, the annualised investment cost is computed as:

$$C_{inv,a} = \sum_{i}^{\#\,tech} \left(C_{inv,fix,i} + C_{inv,var,i} \cdot Cap_i\right) \cdot CRF_i$$

where $C_{invfix,i}$ is the fixed cost, $C_{inv,var,i}$ is the specific cost per unit of installed capacity Cap_i, and CRF_i is the capital recovery factor of technology *i*. The capital recovery factor is given by:

$$CRF_i = \frac{r \cdot (1+r)^{N_i}}{(1+r)^{N_i} - 1}$$

With *r* the interest rate and N_i the economic lifetime (in years) of technology *i*.

This term reflects the contribution of all conversions (including data center cooling systems) and storage technologies to the annual cost.

Annual maintenance costs are represented as

$$C_{maint} = \sum_{i}^{\#\,tech} \left(C_{maint,fix,i} + C_{maint,var,i} \cdot Cap_i\right)$$

where $C_{maint,fix,i}$ is a fixed yearly charge and $C_{maint,var,i}$ is a capacity-dependent maintenance cost per unit of installed capacity. If required, additional energy-dependent maintenance terms (e.g. per MWh operated) can be added in the same way as in the original EhubX formulation.

Finally, the yearly energy-related cost C_energy accounts for all imported energy carriers and possible revenues from exported electricity or heat:

$$C_{energy} = \sum_{t}^{8760\ hours} \sum_{ec}^{\#\,ecs} \left(cost_{imp,ec}(t) \cdot E_{imp,ec}(t) - cost_{exp,ec}(t) \cdot E_{exp,ec}(t)\right)$$

Where $E_{imp,ec}(t)$ and $E_{exp,ec}(t)$ are, respectively, the imported and exported energy of carrier ec (electricity or district heat in this work), while $cost_{imp,ec}(t)$ and $cost_{exp,ec}(t)$ are the corresponding prices or tariffs.

The specific costs and technology parameters selected for the techno-economic investigations are detailed in Section 2.6.

Overall, this objective function ensures that investment decisions, Cap_i, and operational decisions (P (t), Q (t), SOC (t), waste-heat utilization, etc.) are co-optimised to minimise the annualised total cost of supplying all building and data-centre demands, while exploiting on-site generation, storage, and waste-heat recovery whenever economically favourable.

2.1.3. Constraints

The optimisation problem is subject to a set of linearized constraints that ensure physical feasibility of the energy hub and consistency with technology limits. They can be grouped into: (i) energy balance constraints, (ii) capacity and operating constraints, (iii) storage dynamics, and (iv) additional constraints for peak power and waste-heat handling.

For each energy carrier ec (electricity, LT heat, MT heat, HT heat, cooling) and each time step *t*, supply must match demand and internal use. Energy balances are enforced as follows:

$$\sum_{i}^{\#\ conv\ tech} P_{out,i,ec}(t) + \sum_{s}^{\#\ stor\ tech} P_{dis,s,ec}(t) + P_{imp,ec}(t)$$
$$= Dem_{ec}\,(t) + \sum_{i}^{\#\ conv\ tech} P_{in,i,ec}(t) + \sum_{s}^{\#\ stor\ tech} P_{ch,s,ec}(t) + P_{exp,ec}(t)$$

Where $P_{imp,ec}$ and $P_{exp,ec}$ are the imported and exported powers for energy carrier ec at time t, while $Dem_{ec}\,(t)$ is the exogenous demand profile for energy carrier ec. These balance constraints are written separately for electricity, LT/MT/HT heat and cooling.

The installed capacity are bounded by:

$$Cap_{i,min} \le Cap_{\mathrm{i}} \le Cap_{i,max}$$

While the operating limits are based on the installed capacity and conversion efficiency:

$$0 \le P_{out,i,ec}(t) \le \eta_{i,ec} \cdot Cap_{\mathrm{i}}$$
$$0 \le P_{in,i,ec}(t)(t) \le Cap_{\mathrm{i}}$$

Constraints related to minimum part-load or on/off behaviour were not considered in this study.

Concerning storage technologies, the state-of-charge (SOC) evolution for storage *s* is represented by:

$$SOC_{\mathrm{s}}(t) = (1 - \lambda_{\mathrm{s}}) \cdot SOC_{\mathrm{s}}(t-1) + \eta_{ch,s} \cdot P_{ch,s}(t) - (1/\eta_{dis,s}) \cdot P_{dis,s}(t)$$

Where λ_{s} is the standby efficiency and $\eta_{ch,s}$ and $\eta_{dis,s}$ are, respectively, charging and discharging efficiency. The state-fo-charge limits are imposed as follows:

$SOC_{min,s} \cdot Cap_s \leq SOC_s(t) \leq SOC_{max,s} \cdot Cap_s$While charging and discharging power limits are modelled as:

$$0 \leq P_{ch,s}(t) \leq \rho_{ch,s} \cdot Cap_s$$

$$0 \leq P_{dis,s}(t) \leq \rho_{dis,s} \cdot Cap_s$$

Cyclic boundary conditions are enforced for storage technologies:

$$SOC_s(t_0) = SOC_s(t_{end})$$

unless otherwise specified.

To enable a holistic assessment of tertiary buildings with embedded data centres, the data-centre load and its cooling systems are represented as conversion technologies within the energy-hub framework. Their detailed mathematical formulation is provided in Section 2.2.

2.2. Modelling of edge-data centers and cooling systems

The representation of an edge data center and connected cooling system is reported in Figure 9.

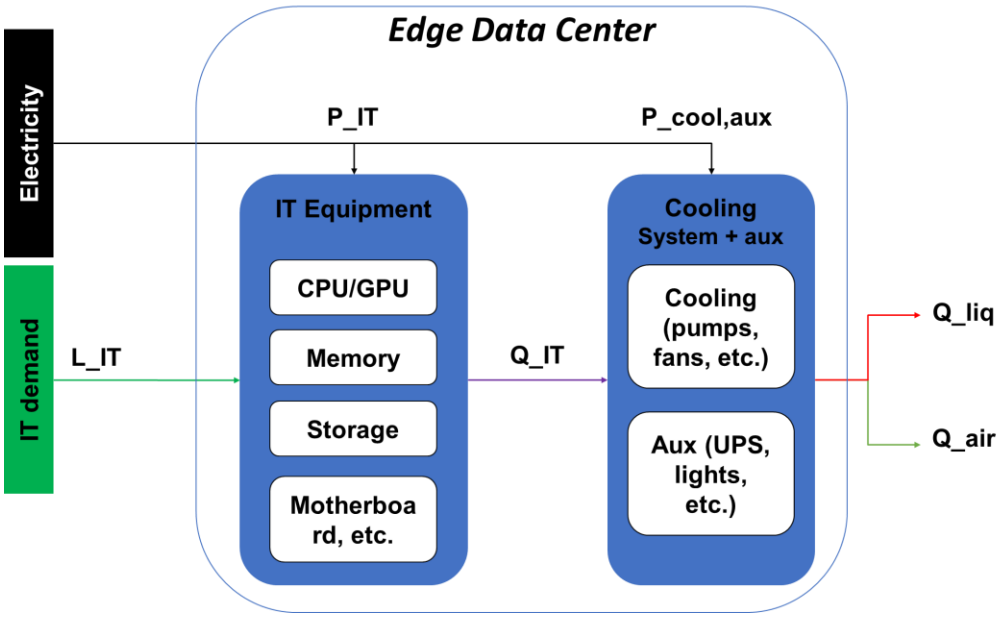

Figure 9 Schematics for the modelling of data centers and cooling systems as conversion blocks.

The edge data center is composed of two main blocks: (i) the IT equipment and (ii) the cooling system and auxiliary block. The IT equipment block needs to satisfy an imposed IT workload, L_{IT}, which dictates an electrical power consumption P_{IT}. The consumed electrical power by the IT equipment is converted to heat. That is:

$$L_{IT}(t) = P_{IT}(t)$$

$$P_{IT}(t) = Q_{IT}(t)$$

The size of the data center (kW_{el}) is defined by the peak of electrical power demanded:

$$L_{IT}(t) \leq Cap_{IT}$$

The heat from the IT equipment, Q_{IT}, needs to be processed by the cooling system block. Here, depending on the type of cooling system considered, different conversion factors are adopted, as detailed in Table 2.

The size of the cooling system is dictated by the maximum amount of thermal power over the investigated horizon:

- For AC cooling: $Q_{air}(t) \leq Cap_{AC}$
- For LC cooling: $Q_{liq}(t) \leq Cap_{LC}$

The amount of electrical power consumed by the cooling system and auxiliary is defined as:

$$P_{cool,aux}(t) = (1 - PUE)\, P_{IT}(t)$$

DTC liquid cooling systems are connected only to CPU components, which represent around 70% to 80% of P_{IT}. This fraction is considered within the parameter η_{LC}. Concerning the air cooling system, the amount of recovered heat also accounts for auxiliary systems, and thus, consider the PUE term within the energy balance equation of Table 2. Nevertheless, the term η_{AC} is included to account for heat losses towards the room. For the considered liquid cooling system, the components that are not cooled by DTC technology are considered to be cooled down by an RDHX.

The way the generated heat is recovered or dissipated depends on the choice made at the building level. The integration of the edge data center models as part of a building's energy hub is detailed in the next section.

Table 2 Energy balances considered for different cooling system types.
The values selected for the investigation are reported in Section 2.4 together with the technology parameters for the other conversion and storage technologies.

Cooling system type	Air outlet (≈25 °C)	Liquid outlet (≈65 °C)
Air Cooling – RDHX	$Q_{air}(t) \leq \eta_{AC} PUE\, Q_{IT}$	0
Liquid Cooling – DTC With air cooling RDHX	$Q_{air}(t) \leq \eta_{AC}(PUE * P_{IT} - \eta_{LC} Q_{IT})$	$Q_{liq}(t) \leq \eta_{LC} Q_{IT}$

2.3. Modelling of the building

This section presents the choices made for the modelling of the energy hub and technology candidates for tertiary buildings. The models for edge data centers and cooling systems presented in the previous section are integrated into the building model.

Two building-level scenarios were considered:

I. Building connected to DHN, Figure 10;

II. Building disconnected from DHN and heated by air-sourced heat pump, Figure 11.

The type of building considered is based on the NEST building [22], with some modification to improve the generability of the results.

The building relies on three hydronic networks operating at distinct temperature levels, which allow flexible distribution of heating and cooling:

- **Low-Temperature (LT) network:** *Supply 7°C, return 14°C.* Used to satisfy cooling demand (high in summer, low in winter).
- **Medium-Temperature (MT) network:** *Supply 35°C, return 25°C.* Used to meet space heating demand (high in winter, negligible in summer).
- **High-Temperature (HT) network:** *Supply 65°C, return 45°C.* Used to cover domestic hot-water (DHW) demand (relatively constant year-round).

These three temperature levels are natively compatible with the heat-recovery potential of the two cooling technologies assessed in this study (rear-door air cooling and direct-to-chip liquid cooling).

When the building is connected to the district heating network (Figure 10), the DHN acts as the main heat supply for the HT network. The DHN is assumed to deliver heat at a temperature compatible with the building's HT requirements. Within the energy hub, water-to-water heat exchangers allow temperature downgrading from the DHN/HT level. This enables the DHN to supply thermal services also for space heating, depending on the operational strategy.

In addition to heat exchangers, water-to-water heat pump is considered to upgrade heat from MT to HT loop when this option is cost-effective and air-to-water heat pump to upgrade low-grade heat recovered from the air-cooling system to MT-level temperature. Such a heat pump, as well as all conversion and storage technologies in the energy hub, is installed only if it is cost-effective. If installing the ASHP for heat recovery is not economically optimal, the heat extracted from the DC via the air loop is rejected to the LT

network, where it appears as an additional cooling load. The cooling load is satisfied by an air-sourced chiller.

PV generation, battery storage, and electricity exchange with the grid are available as technology candidates in both scenarios.

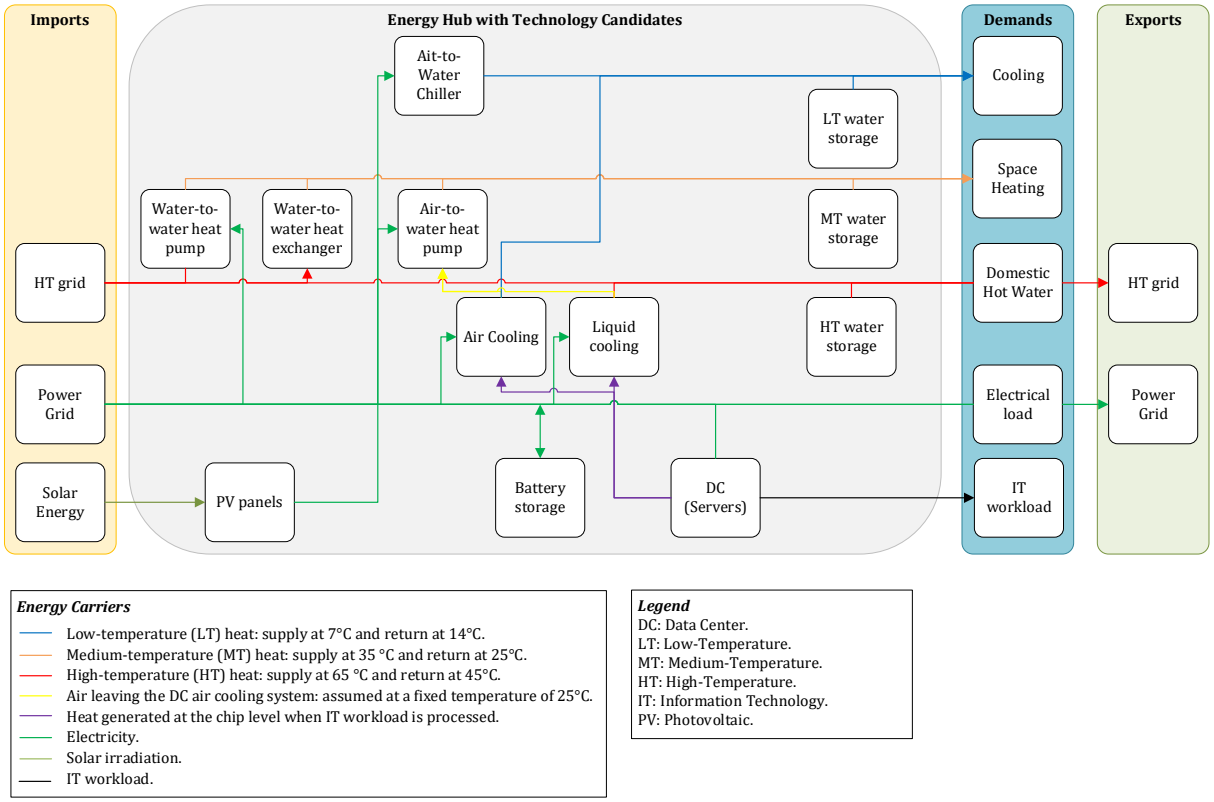

Figure 10 Schematic of the energy hub and technology candidates for a tertiary building connected to a district heating network.

When the building is not connected to a DHN (Figure 11), the air-source heat pump (ASHP) becomes the central heating technology. The ASHP extracts heat from Ambient air and from the Recovered heat from the data centre cooling system.

This configuration leads to a nonlinear COP, as the heat pump's source temperature depends on the operational choice between ambient heat and recovered DC heat. To maintain linearity in the MILP formulation, this nonlinear behaviour is treated through pre-processing:

- The heat pump is assumed to operate in parallel with the space-heating demand.
- Recovered DC heat is always prioritised when available (based on the fixed IT workload).
- A heat-weighted average source temperature is computed for each time step.
- The estimated temperature is used to determine a dynamic COP for the heat pump.

The resulting COP time series is passed to the MILP as an external input, preserving accuracy while keeping the optimisation linear and tractable.

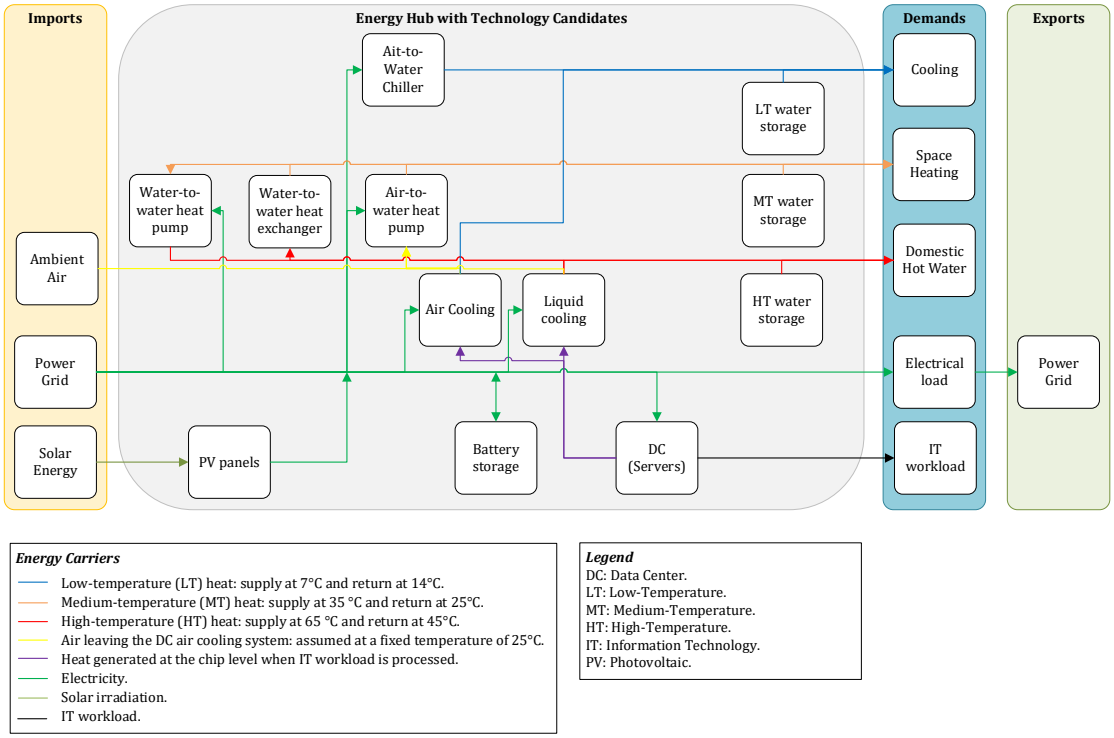

Figure 11 Schematic of the energy hub and technology candidates for a tertiary building not connected to a district heating network and using an air-sourced heat pump as the main heating system.

If not otherwise specified, the results presented in this report refer to the DHN-connected building scenario (Figure 11).

2.4. Input data

The input data can be categorized into five main groups: energy demand, weather data, IT workload, energy costs, and technology parameters.

Energy demand, weather data, and IT workload

The energy demand data used in this study includes hourly profiles for electricity, space heating, domestic hot water, cooling, and IT demands. For space heating, a supply temperature of 35 °C was assumed, consistent with low-temperature distribution systems such as underfloor heating, and aligned with the design specifications of the NEST demonstrator [22]. The adopted demands are based on historical data and integrated effects from the occupants' behaviour and heat gains. Examples of the time-varying input data are shown in Figure 12, although these inputs vary depending on the use case analyzed (Section 2.6). The represented IT demand is converted into electrical power. The

IT demand was extracted from for a single week of operation and repeated over the year [32].

The developed model also requires irradiance data to analyze the performance of PV systems, and ambient temperature data to infer the coefficient of performance (COP) of heat pumps and chillers.

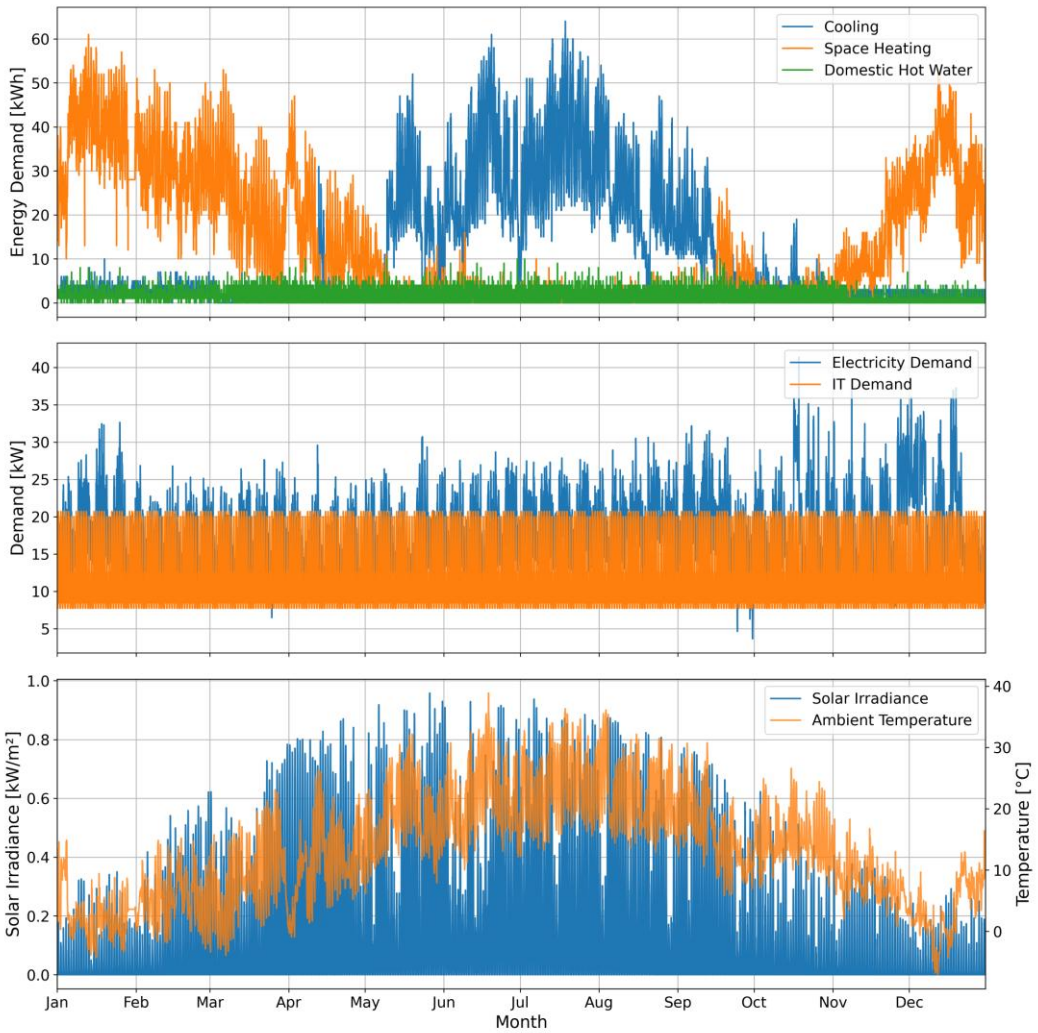

Figure 12 Energy demands and weather data for the NEST building in Dübendorf.

Energy costs

In this work, energy costs are related to electricity prices from the grid and heat prices from district heating networks.

Electricity prices are based on the 2025 Glattwerk business tariff [33], which applies to tertiary buildings. The tariff includes energy supply, grid usage, system services, surcharges, and VAT, and has a seasonal structure. Glattwerk defines three seasonal windows, with different prices for each, as summarized in Table 3. Each period has peak

(high tariff) and off-peak (low tariff) prices. Peak hours refer to Monday to Friday, between 07:00 and 20:00, and Saturday, between 07:00 and 13:00. Feed-in tariffs were instead considered constant and equal to 0.03 CHF/kWh.

These values are converted into an hourly time series (8760 points) following the seasonal and day-type rules. The fixed monthly charges ("Grundpreis") and optional power-based tariff ("Leistungspreis") are excluded because they represent a minor share of the total cost.

Table 3 Variation of the electricity price over the year from [33].

Period of the year	High tariff (Rp/kWh)	Low tariff (Rp/kWh)
Winter 1 (Oct–Dec)	26.76	25.35
Winter 2 (Jan–Mar)	25.49	24.08
Summer (Apr–Sep)	23.00	21.60

When the building is connected to DHN, heat prices are based on the ewz waste-incineration and wood-based district-heating tariff applicable to the Zurich region [34].

The DHN energy price is defined by ewz, leading to a tariff of 0.0806 CHF/kWh. No return-temperature surcharge is applied, as the goal of this study is to isolate the effect of heat recovery rather than analyse contract penalties. In addition to the DHN energy price, district-heating customers also pay an annual capacity charge based on their subscribed thermal capacity. For a representative subscribed capacity of 100 kW, this yields ≈ 5763 CHF/year in the instance of the NEST building.

To obtain an effective heat cost comparable to the electricity tariff, the capacity charge is allocated over the annual thermal demand. For the considered heating demands, the effective heat price was calculated to be approximately 0.126 CHF/kWh. Empa pays a higher price for heat instead, up to 0.21 CHF/kWh due to the installed heating supply systems, where gas boiler systems are still largely used in winter. Such a price will also be considered in the analysis to generate insights related to the Empa pilot.

No revenues were attributed to heat exported to the DHN, reflecting the typical situation for small-scale buildings. However, the DHN was assumed to be continuously available for heat injection in the presence of surplus production, and was therefore modelled as a heat sink for the building.

The hourly carbon intensity for the Swiss electrical grid was taken from [35], while a fixed carbon intensity of 0.210 kg_{CO2}/kWh is assumed for the DHN. The carbon intensity of the DHN was determined based on the assumption that the heat was produced from gas-

based sources [36], using emission factors of 0.130 kg_{CO2}/kWh for biogas and 0.227 kg_{CO2}/kWh for natural gas [37]. This value depends strongly on the specific heating technologies supplying the DHN.

Technology parameters

Several technology parameters were specified for the considered conversion and storage energy technologies, with the full list of parameters reported in Table 4 and

Table 5.

While we acknowledge that CAPEX per capacity does not scale linearly with the installed size [38], we selected the most plausible values in the size range of interest to comply with the linear input requirements of MILP approaches. We chose not to apply piecewise affine linearization to limit the computational burden of the model, with the finality of exploring a vast range of scenarios. Nonetheless, the most uncertain technology parameters were included in a sensitivity analysis.

For most technologies, we assumed a fixed installation cost corresponding to approximately 3 person-days of professional work (≈ 5 kCHF).

Table 4 Technology parameters for conversion technologies.
(*) A maximum installation capacity of 400 m2 was set for PV.
(**) The IT equipment costs were not included in the analysis as they do not impact the energy hub optimization.
(***) The PUE adopted for air-cooled data centers accounts for auxiliary equipment, compressor, and the air-to-water rear-door heat exchanger. The energy consumption of the chiller is not included in the PUE and is modelled within the optimization framework. The selected PUE value was derived from measurements.
HE: Heat Exchanger; HP: Heat Pump; LT: Low-Temperature; MT: Medium-Tempearture; HT: High-Temperature; WW: Water-to-Water; AW: Air-to-Water.

Technology	Lifetime [years]	Capex fixed [€]	Capex per capacity [€/kW]	Opex per capacity [€/kW/year]	CO_2 per capacity	PUE	eta_th	COP	Applicable temp. level	Ref
IT Equipment (**)	-	–	-	-	-	-	1	-	–	-
Air Cooling	10	5000	825	21.45	30	1.2 (***)	0.8	-	25°C	[39], [40], [41], [42], [43]
Liquid Cooling	10	8000	1237	33	60	1.2 Air 1.1 Liquid	0.8 Air + 0.8 Liquid	-	25 °C Air 65°C Liquid	[40], [43], [44]
HP_WW	20	5000	900	20	30	-	-	4.5	35→65°C	[45]
HE_WW	20	5000	608	6	20	-	0.95	-	65→35°C	[45]
HP_AW	20	5000	1100	22	30	-	-	5.0	25→35°C	[45]
Chiller	20	5000	670	17	20	-	-	COP (T_amb)	7°C	[45]
HE_LT	20	5000	608	6	20	-	0.95	-	35→7°C	[45]
PV(*)	25	1075	280	2.6 per m2	252 per m^2	-	0.17	–	-	

Table 5 Technology parameters for storage technologies from the EhubX database and [45].
(*) While batteries are typically operated in an SOC range from 0.2 to 0.8, this constraint was lifted in the analysis to reduce computational cost, and given the heating system focus of the study.
(**) Given the similar temperature differences for the TES, we assumed the temperature difference doesn't affect the Capex per cap.

Technology	Lifetime [years]	Capex fixed [CHF]	Capex per cap [CHF/kWh]	Opex per cap [CHF/kWh/yr]	CO2 per cap [kg/kWh]	η_in	η_out	Standby loss [1/h]	SOC_min	SOC_max
TES_MT	30	1000	90 (**)	0.9	10	0.95	0.95	0.001	0	1
TES_HT	30	1000	90	0.9	10	0.95	0.95	0.001	0	1
TES_LT	30	1000	90	0.9	10	0.95	0.95	0.001	0	1
Li-ion Battery (*)	12	1600	500	7.5	157	0.98	0.97	0.001	0	1

2.5. KPIs definition

The mathematical definitions of the key performance indicators (KPIs) used in this study are reported in Table 6. The chosen KPIs are aligned with those defined in previous HEATWISE deliverables and were adapted here to reflect the modelling choices, boundary conditions, and technology configurations used in this work. The Power Usage Effectiveness (PUE) of the data center is provided as model input and remains fixed throughout the optimization. The adopted value for PUE can be found in

Table 5.

The Heating Degree Days (HDD) are defined as:

$$HDD = \sum_{days}^{365} \max\left(T_{base} - T_{amb}(day), 0\right)$$

With $T_{base} = 18°C$.

Payback time (PBT) indicators are not included in the present KPI set, as all case studies are formulated as greenfield optimization problems. This modeling choice allows the assessment of the holistic techno-economic and environmental performance of optimized energy hub configurations, without constraining the solution space to existing system states. Payback time analyses, which require explicit reference baselines and investment phasing, will be addressed in Task T7.4, where retrofit-oriented scenarios will be considered.

Table 6 Definition of the KPIs calculated in this report.
Given the focus on the NEST pilot, most of the economic results presented in this report are expressed in CHF.

Name	Definition	*Ref*
Energy Reuse Factor: ERF [-]	$\frac{\sum_t^{8760} Q_{DCtoBuilding}(t)}{\sum_t^{8760} P_{DC}(t)} = \frac{\sum_t^{8760} Q_{DCtoHT}(t) + Q_{DCtoHPAW}(t) - Q_{exp}(t)}{\sum_t^{8760} P_{DC}(t)}$	[5] [46]
Renewable Energy Factor: REF [-]	$\frac{1}{\mathrm{N}}\sum_t^{N=8760} min\{\frac{P_{PV}(t)}{P_{DC}(t)}, 1\}$	[5] [46]
Degree of Coverage: DOC [-]	$\frac{1}{\mathrm{N}}\sum_t^{N=8760} min\{\frac{Q_{DCtoBuilding}(t)}{Q_{dem}(t)}, 1\}$	[11]
Potential waste heat coverage: PWHC [-]	$\frac{1}{\mathrm{N}}\sum_t^{N=8760} min\{\frac{Q_{DC}(t)}{Q_{dem}(t)}, 1\}$	[11]

Normalized Total Cost per Heating Degree Day: NTC-HDD [CHF/(m²·K·day)]	$\frac{TAC}{Heating\ Area\ \cdot HDD}$	[47]
Normalized Total Emissions per heating degree day: NTE-HDD [g_{CO2}/(m²·K·day)]	$\frac{TE}{Heating\ Area\ \cdot HDD}$	[47]
Renewable Energy Utilization: REU [-]	$\frac{1}{\mathrm{N}}\sum_{t}^{N=8760} min\{\frac{P_{PV}(t)}{P_{PV}(t)+P_{imp}(t)+P_{exp}(t)}, 1\}$	[47]
Thermal Energy Storage Utilization: TESU [-]	$\frac{1}{\mathrm{N}}\sum_{t}^{N=8760} min\{\frac{Q_{TES,out}(t)}{Q_{dem}(t)+Q_{exp}(\mathrm{t})}, 1\}$	[47]

2.6. Multi-vector interdependencies and EMS implications

The analysis presented in this deliverable adopts a greenfield, system-level optimisation approach to investigate the optimal cooperation between multiple energy vectors within tertiary buildings hosting edge data centres. Rather than analysing technologies in isolation, the energy hub formulation implemented in EhubX considers electricity, heating, IT tasks, cooling, waste heat, and thermal and electrical storage simultaneously, allowing interdependencies and trade-offs to emerge endogenously from the optimisation.

A key aspect of this approach is the explicit modelling of waste-heat recovery from data centres as a conversion process that couples the electricity and thermal domains. Electricity consumed by IT equipment is converted into thermal energy at different temperature levels depending on the selected cooling technology. This recovered heat can then substitute other heat sources, be upgraded through heat pumps, stored, exported to the district heating network, or rejected when surplus occurs. As a result, electricity consumption, thermal demand fulfilment, cooling requirements, and heat exports are tightly interlinked.

The results show that integrating data-centre waste heat alters the optimal cooperation between energy carriers and technologies. For instance, waste-heat recovery reduces the need for electrically driven chillers and changes the role and sizing of heat pumps, thermal energy storage, and electricity imports. Conversely, auxiliary electricity consumption for cooling and heat upgrading affects PV sizing and electricity flows. These interactions demonstrate that the benefits of waste-heat recovery cannot be assessed through single-technology metrics, but rather through a holistic, multi-vector analysis.

The resulting solutions provide insights into which interdependencies are most critical from an operational perspective, highlighted in Section 5 and that directly informs the

development of energy management and control strategies addressed in Task 7.3 of the HEATWISE project.

2.7. Analyzed use cases and scenarios

In this report, we utilize design optimization tools for energy systems to assess the cost-effectiveness of various cooling and heat recovery strategies for edge data centers integrated into tertiary buildings, such as schools, hospitals, restaurants, etc. To assess the robustness of the results, multiple scenarios and use cases are explored, analyzing the influence of building type, technology options, location, data center size and IT workload on the economic viability of heat recovery. The use cases explored are summarized in Figure 13.

A benchmark case is defined for each building type by assuming no heat recovery from the data center and a conventional cooling strategy (air cooling with full rejection to ambient). This benchmark provides a consistent baseline against which the techno-economic value of heat recovery and advanced cooling technologies is measured.

The energy-demand profiles for space heating, cooling, domestic hot water, and electricity are derived from measurements for the NEST building in Dübendorf, Switzerland, while modelled demand profiles for representative Swiss tertiary buildings are also analyzed [48].

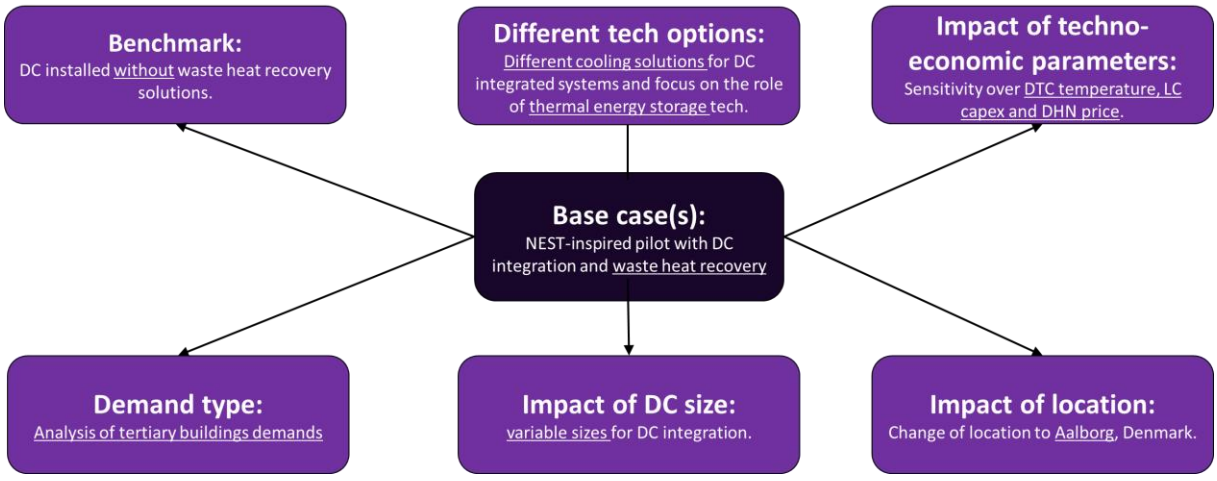

Figure 13 Scenarios and use cases explored in this study.

3. Results

Figure 14 summarises the optimal technology portfolios and key performance indicators for the three scenarios analysed: liquid cooling, air cooling with heat recovery, and air cooling without heat recovery, i.e. no air-to-water heat pump within the design space. **The results show that recovering waste heat from the data centre is economically advantageous for the building**. Both liquid cooling and air cooling with heat recovery lead to a lower total annualised cost compared with the case without heat recovery. This cost reduction is achieved because the recovered heat partially replaces heat that would otherwise be supplied by electricity-driven systems, reducing annual energy costs, as well as it allows for a reduction in the sizing of other heating system components.

Waste-heat recovery also influences the optimal energy hub configuration. When heat recovery is enabled, the required chiller capacity decreases because a portion of the cooling demand can be met through heat extraction from the data centre. In the optimisation framework, this leads to smaller chiller sizing and associated CAPEX reductions, although in a real installation, redundancy or safety considerations might still require keeping the full chiller capacity.

Thermal energy storage is also affected: medium-temperature storage is installed in most cases, since the building's heating needs are mostly at this temperature level. The liquid-cooling configuration with heat recovery features the largest storage capacity to better exploit its waste heat potential at higher temperatures.

The water-to-water heat pump is never installed in any scenario, as the majority of the demand can already be met at medium temperature and the high-temperature load is relatively small. The air-to-water heat pump appears only when upgrading air cooling waste heat to medium temperature is beneficial; it is not needed for liquid cooling, where most of the heat is recovered directly at high temperature.

Photovoltaic capacity is not affected by the cooling choice within the maximum allowed capacity (400 m2). When higher capacities are allowed, we noticed that the liquid cooling scenario installs a slightly smaller PV system because of the absence of an air-to-water heat pump, thus reducing the incentive for on-site electricity generation. This behaviour highlights how the cooling technology indirectly reshapes the optimal electricity supply configuration.

The ERF varies substantially across scenarios. Liquid cooling achieves the highest recovery, with values around 55%, because it recovers high-temperature heat suitable for both space heating and domestic hot water. Air cooling with heat recovery reaches around 40%, reflecting the lower temperature of the recovered heat and the need to

upgrade it through a heat pump. Increasing the energy recovery fraction further would require addressing seasonal mismatches between the availability of waste heat and the building's heating demand, for example, through long-term thermal storage. The renewable energy fraction remains essentially unchanged across cases, since battery storage is never installed and PV capacity differs only marginally.

At the considered scales, the air-to-water heat pump required for the liquid-cooled system is relatively small but still economically viable. However, its cost-effectiveness is strongly influenced by fixed installation costs, which therefore need to be carefully considered when assessing small-scale installations.

Overall, the results indicate that heat recovery from edge data centres is cost-effective and significantly influences the optimal design of the building's energy system. Waste-heat recovery reduces chiller capacity, increases the use of medium-temperature thermal storage, and improves the utilisation of recovered heat, while the choice of cooling technology strongly affects the achievable recovery levels.

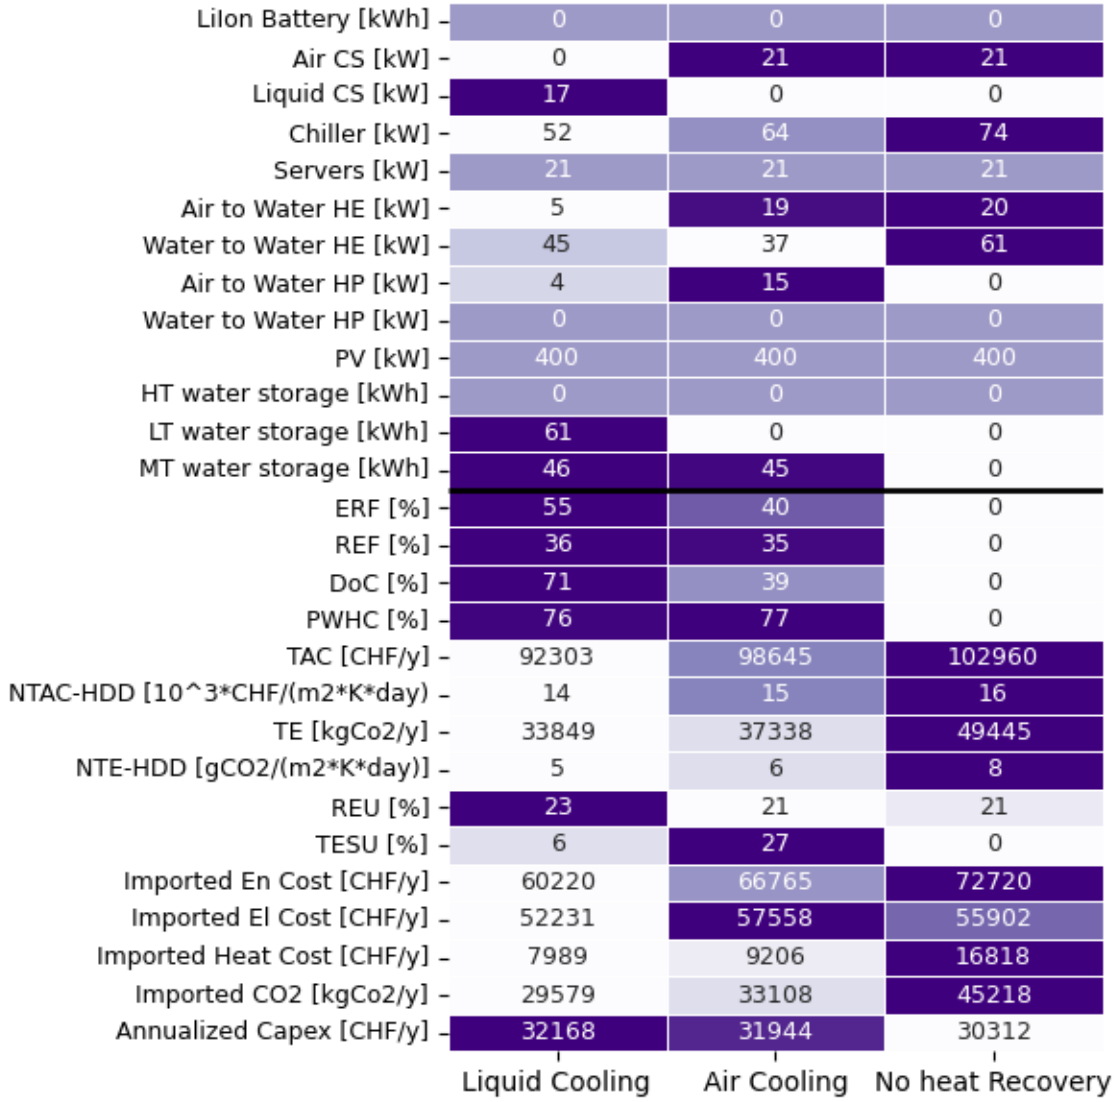

Figure 14 Impact of edge data center cooling system on the energy hub.

3.1. The role of thermal energy storage

Figure 15 illustrates the annual operational profile of medium- and high-temperature heat demand and supply for the air-cooling configuration. The stacked curves clearly highlight a strong temporal mismatch between the availability of waste heat from the data centre and the building's thermal demand. During summer months, roughly between hours 2000 and 6000, the medium-temperature demand is very low, while the data centre continues producing a nearly constant heat output. Because demand is insufficient and storage capacity is limited, a significant share of the available waste heat cannot be absorbed by the building and is therefore rejected (shown as Lost Heat in the Figure).

This seasonal imbalance directly constrains the ERF: even though waste heat is consistently available, only a portion can be utilised when the thermal demand is low. By contrast, in winter, most of the available heat can be absorbed, either directly meeting the medium-temperature demand or being upgraded through the heat pump for high-temperature domestic hot water. The figure therefore makes evident that ERF is fundamentally limited by the seasonal structure of the building's thermal loads rather than by the performance of the cooling technology itself.

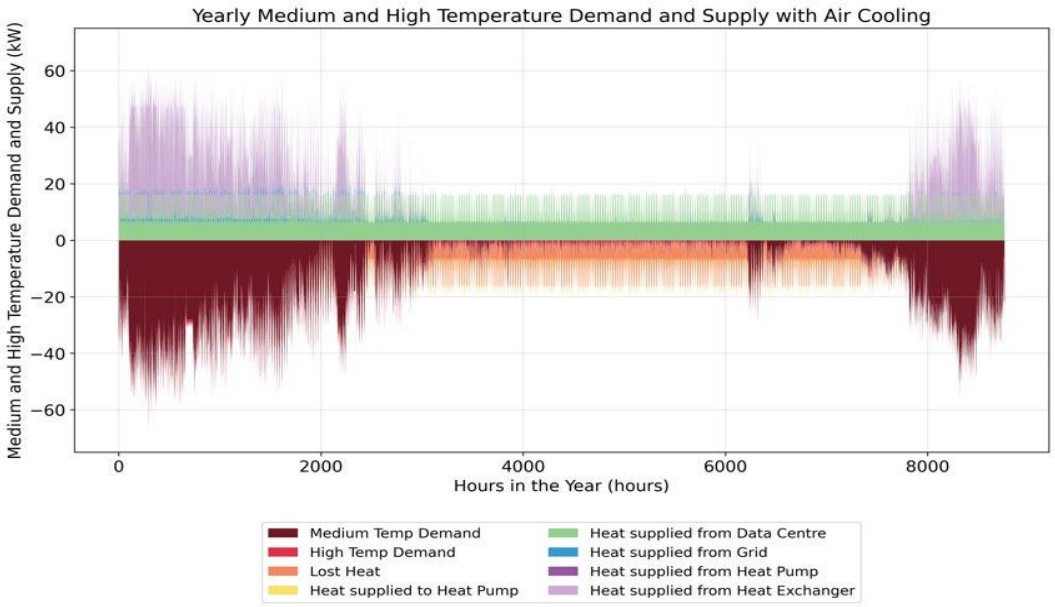

Figure 15 Yearly operation of selected technologies.

Figure 16 illustrates representative weekly profiles of the building's thermal demands and the data centre waste-heat availability for a winter week (top) and a summer week (bottom). Two observations emerge

(i) Operational flexibility plays a limited role in these periods: in winter, the space-heating demand is consistently higher than the waste-heat availability, so all recovered heat can be absorbed without requiring complex control. In summer, by contrast, the space-heating load is negligible, and the domestic hot water (DHW) demand shows only short, low-power peaks, meaning that most of the recovered heat exceeds the building's immediate needs and must be rejected.

(ii) The magnitude and temporal distribution of the DHW demand are such that it can be fully supplied by the data-centre waste heat throughout both weeks. This reinforces the earlier conclusion that DHW is the most suitable end-use for baseload waste-heat integration, while the mismatch between summer availability and demand limits the achievable annual energy-recovery fraction.

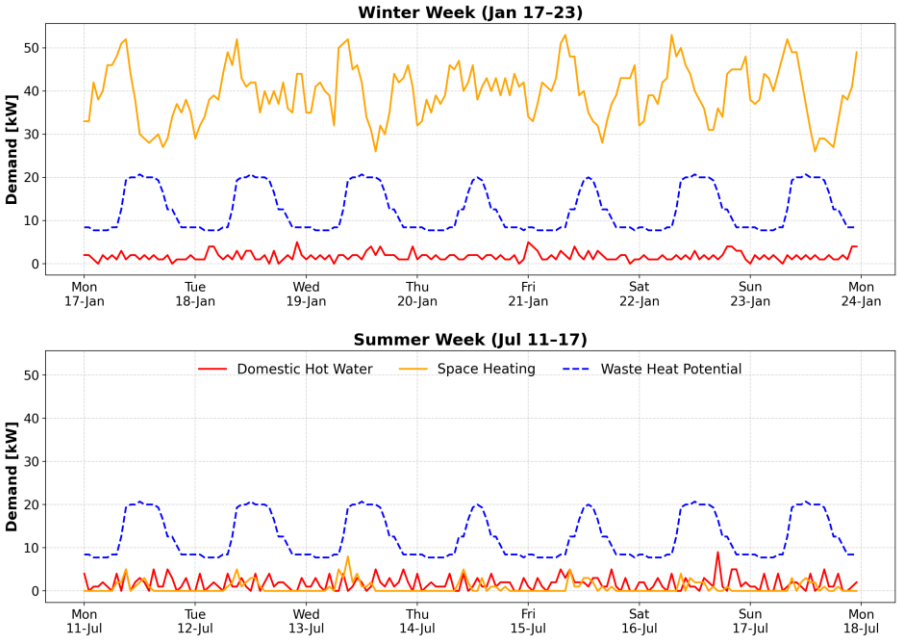

Figure 16 demands and waste heat potential for a winter and summer week.

Figure 17 shows the operation of the medium-temperature (MT) thermal loop during a representative mid-season day (April 10th). This period (mid-seasons) captures conditions where neither heating nor cooling dominates, allowing us to observe how thermal storage interacts with fluctuating waste-heat availability and moderate heating needs.

During this day, the data centre provides a relatively stable amount of medium-temperature heat, shown in blue. The MT demand profile (black line) varies throughout the day, exceeding the direct DC supply in some hours while falling below it in others. The thermal energy storage (TES) system is used to buffer this mismatch: when DC waste heat exceeds the instantaneous MT demand, the surplus is stored (purple segments below zero). Conversely, when MT demand temporarily exceeds direct supply, stored heat is discharged from TES (green segments above zero). This behaviour smooths the MT loop and reduces reliance on auxiliary supply technologies.

Despite this effective short-term balancing, the amount of energy that can be shifted using TES remains limited because the fundamental seasonal mismatch persists: excess waste heat is typically available in summer, while heating needs peak in winter. Nonetheless, even with this limitation, TES is consistently selected by the optimisation as a cost-effective technology, as it enables modest but valuable improvements in waste-heat utilisation and reduces the capacity requirements of other heating technologies.

Overall, the inclusion of MT storage increases the ERF by approximately +0.03, confirming that short-term storage plays a supportive, though not transformative, role. To achieve higher ERF values, the system would need to address seasonal imbalances, for instance

through high-temperature seasonal storage or alternative pathways for summer heat export [49]. However, seasonal thermal energy storage is typically not cost-effective at a single building scale [50].

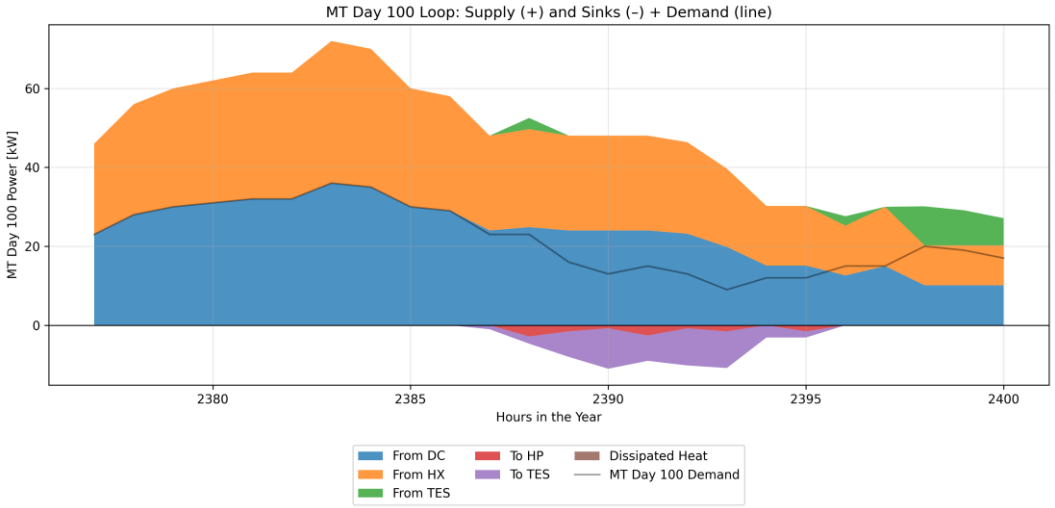

Figure 17 A single day of operation shows the role of short-term flexibility.

3.2. Impact of techno-economic parameters

Figure 18 explores the sensitivity of the techno-economic results to three key assumptions: (i) the outlet temperature of the direct-to-chip (DTC) liquid cooling system, (ii) higher DHN heat prices, and (iii) increased CAPEX per capacity for liquid cooling technologies. The standard parameters (Table 4 and Table 5) presented previously are used as the reference case.

Regarding the cooling-system temperature level, assuming a reduced DTC outlet temperature of 35 °C does not lead to a significant reduction in TAC. This result is primarily explained by the structure of the heat demand in the NEST building, where the dominant share is space heating delivered through medium-temperature floor heating systems operating below 35 °C. As a consequence, MT waste heat remains largely usable even when high-temperature recovery (65 °C) is not available. The main penalty associated with the lower DTC temperature arises in meeting DHW demand, which still requires heat imports from the DHN. However, because DHW represents a relatively small share of the total thermal demand, the additional imports have only a marginal impact on overall costs and emissions. The predicted DoC increases due to the reduced heat losses in supplying MT heat, i.e. no need for the HE_WW with 95% efficiency.

For the case with T_DTC = 35 °C, the ERF and the TESU could not be reliably computed due to a numerical artefact. Specifically, the optimiser exploited the medium-temperature storage as a heat sink to dissipate excess recovered heat, resulting in non-physical operation. While this limits the interpretability of ERF and TESU in this specific scenario, the consistency of other key indicators suggests that the overall trends remain robust. **Importantly, these results indicate that, for new buildings, high-temperature waste heat (≈65°C) is not strictly required to achieve cost-effective integration in systems where the dominant demand lies at medium temperature levels (≈35 °C).**

Increasing DHN heat prices, reflecting the Empa tariff structure instead of the ewz reference (Section **Hiba! A hivatkozási forrás nem található.**), leads to moderate changes in system behaviour. Higher import costs incentivise greater utilisation of locally recovered waste heat, resulting in a slight increase in thermal storage installation and utilisation. This, in turn, contributes to a higher ERF, as a larger fraction of generated waste heat can be effectively reused. However, the higher DHN prices inevitably translate into increased TAC, despite improved heat recovery performance.

The increased DHN cost case also shows higher imported heat-related expenditures and emissions compared to the baseline, demonstrating that price signals alone are insufficient to fully decouple the system from external heat sources, especially under seasonal demand–supply mismatches. Nevertheless, the results confirm that price assumptions influence optimal design choices and should be carefully adapted to site-specific conditions in future analyses.

Finally, doubling the CAPEX per capacity of the LC system provides an important robustness test for the economic viability of advanced cooling solutions. Even under this conservative assumption, liquid cooling remains part of the optimal system configuration. **While annualised CAPEX increases, the total system cost remains competitive, and no switch back to purely air-cooled or non-recovery configurations is observed.** This is a key result, as it demonstrates that the benefits of improved waste heat recovery, reduced chiller sizing, and lower reliance on imported heat outweigh higher upfront investment costs within the optimisation framework.

Across all sensitivity cases, photovoltaic capacity remains unchanged, constrained by the available installation area. This indicates that the observed variations in system costs and performance are driven primarily by interactions within the heating, cooling, and storage subsystems rather than by changes on the electricity generation side.

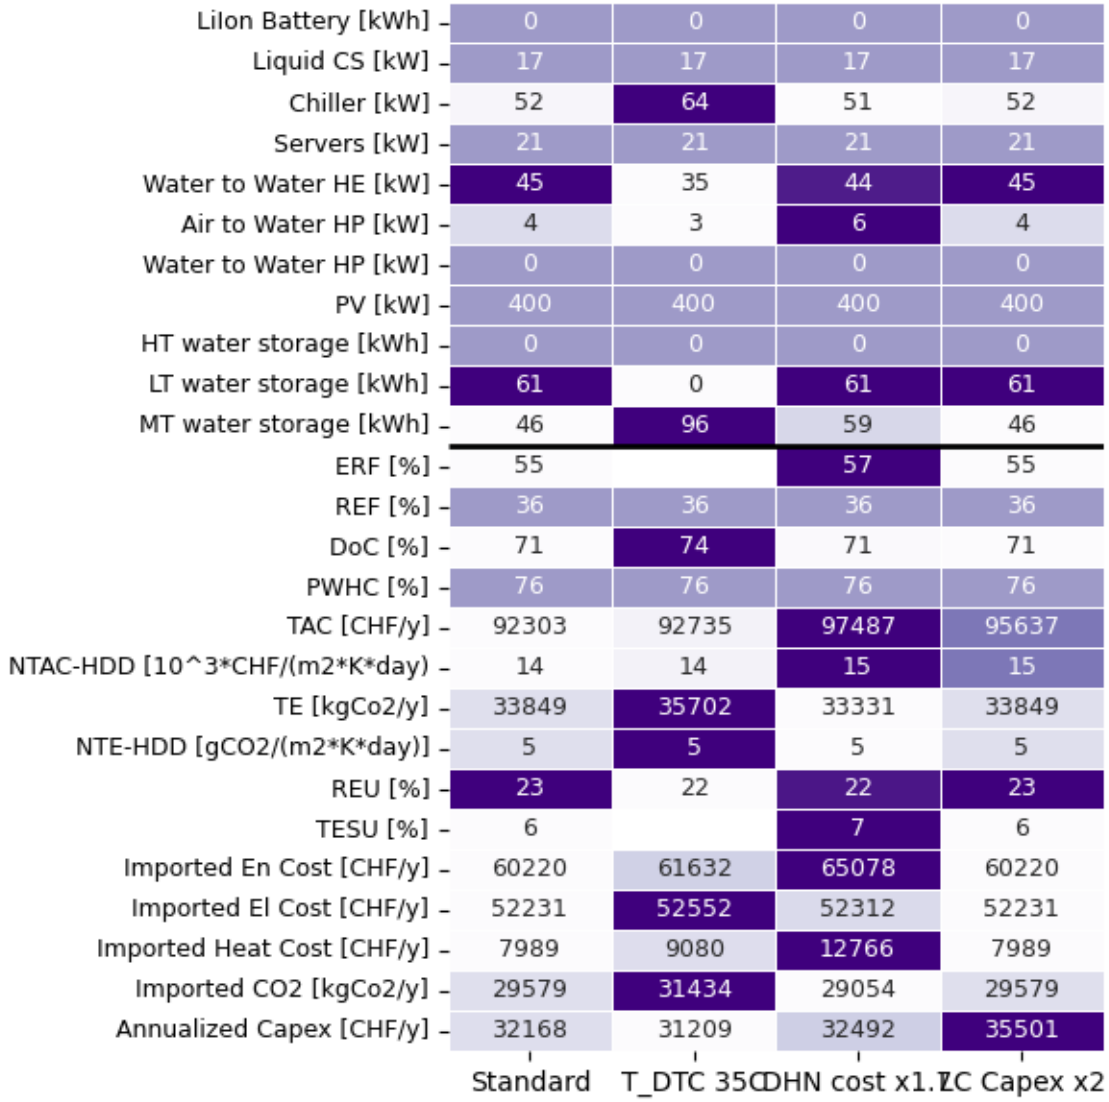

Figure 18 Impact of direct to chip (DTC) cooling temperature, DHN import cost and LC capex per capacity on the optimal sizing and KPIs.

Figure 19 depicts the comparison of the NTAC-HDD for different scenarios. The DTC liquid cooling system remains the most cost-effective solution also under conservative price and performance assumptions.

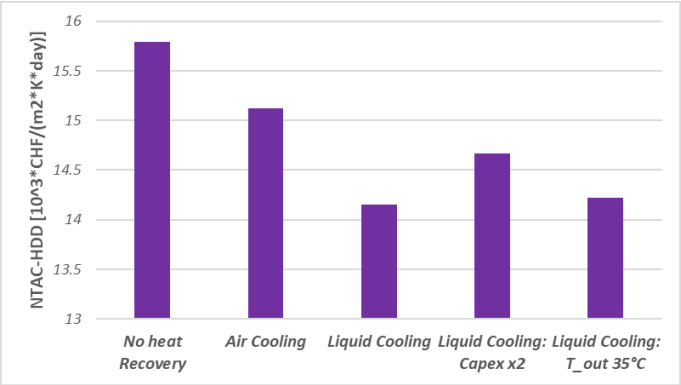

Figure 19 Comparison of the normalized total annualized cost per heating degree day.

3.3. Impact of tertiary building type

The results discussed earlier highlighted that buildings with stable heating needs throughout the year benefit the most from data-centre waste heat recovery, because a steadier demand profile allows for higher and more consistent utilisation of the available waste heat. To explore this aspect further, demand profiles from a broader set of tertiary buildings were analysed.

The demand data were taken from the SCCER JASM hourly demand dataset [48]. This dataset provides synthetic, building-type specific hourly energy demands for space heating, DHW, electricity, and cooling for Swiss building archetypes constructed in different periods and under different retrofit conditions. While these profiles are numerically generated and may not capture all the nuances of real operation, they are well suited for trend identification, seasonality assessment, and relative comparison across building types.

Figure 20 compares normalized annual heating demand profiles for two representative cases: hospitals and offices. Although both buildings require space heating predominantly in winter, hospitals exhibit a flatter and more continuous demand, whereas office buildings show stronger weekday patterns and larger fluctuations between peak and off-peak periods. This contrast illustrates how building type directly influences the potential to integrate and utilise waste heat throughout the year.

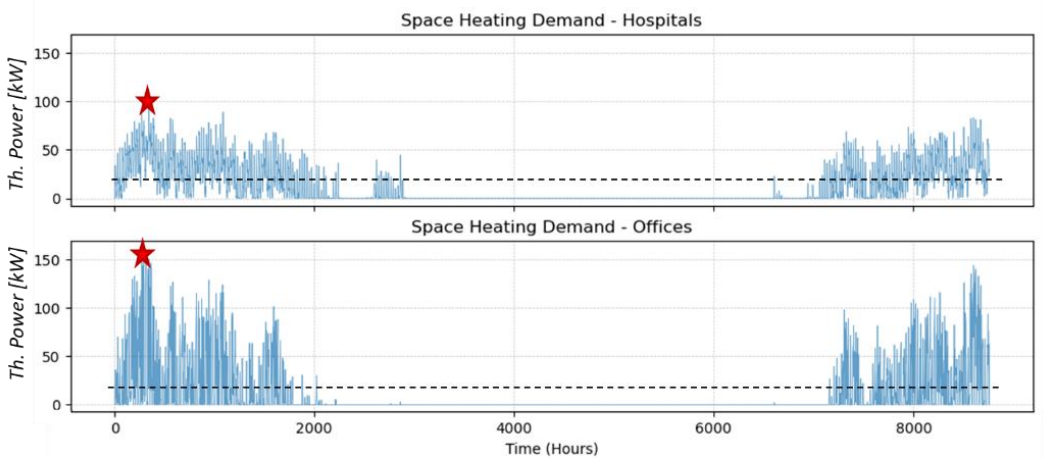

Figure 20 Examples of space heating demands for hospitals and offices from [48].
The red star refers to the peak power, while the black dashed line represents the average power.

To quantify this behaviour more systematically, we computed a peak-to-average ratio (PAR) defined as:

$$PAR = \frac{Q_{peak}}{Q_{average}}$$

Lower PAR values indicate steadier demand and, consequently, a higher likelihood of absorbing waste heat continuously.

The full dataset was analysed for multiple building categories (schools, offices, shops, restaurants, hospitals, etc.), also considering the spread due to the construction period and renovation state. Figure 21 summarizes the PAR distributions. Error bars represent the variability introduced by different envelope qualities and thermal characteristics across the archetypes.

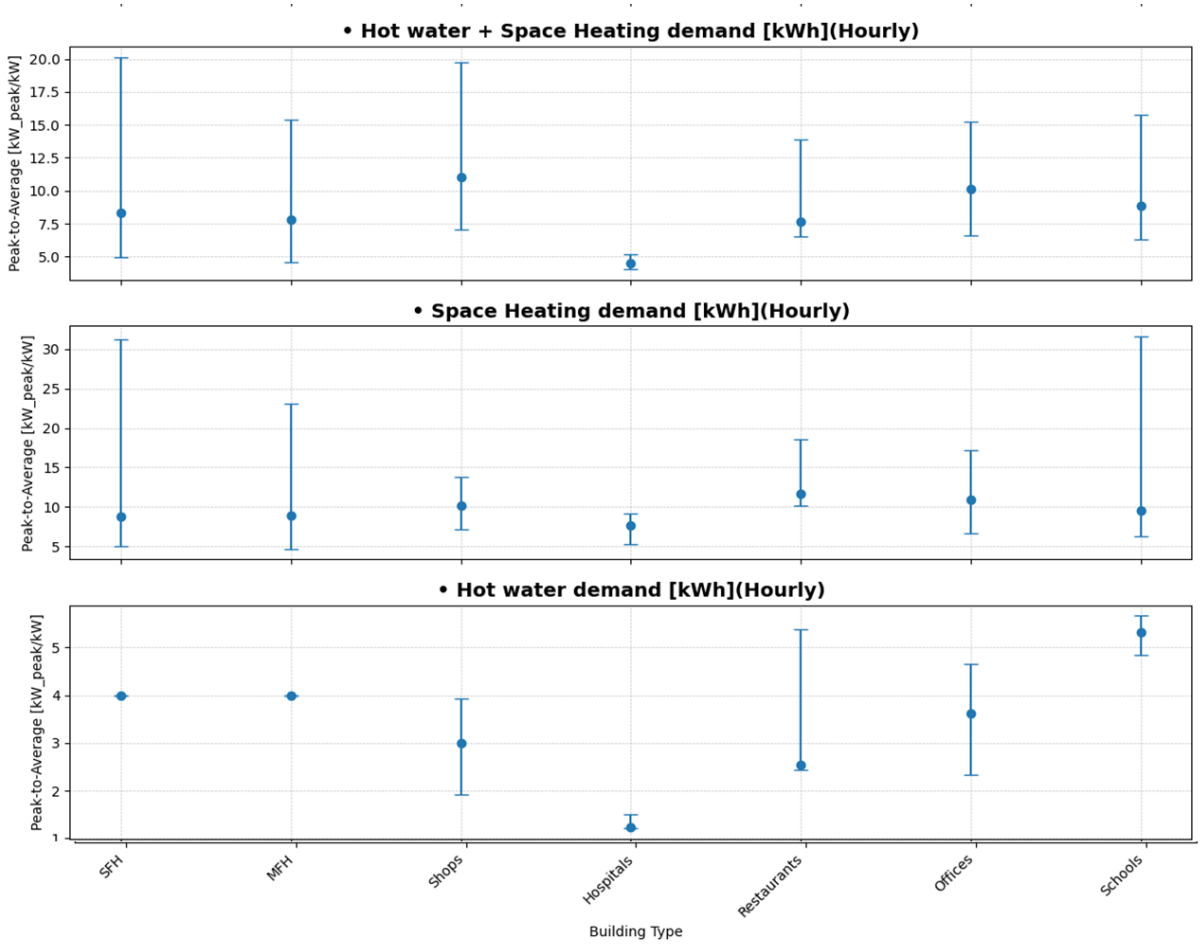

Figure 21 Peak to average distribution for different building types.
SFH: single-family house; MFH: multi-family house.

Two observations emerge clearly:

1. Domestic hot water demand exhibits consistently low PAR values across all building types, meaning that DHW is comparatively stable and a reliable sink for waste heat. However, DHW demand is generally small in magnitude, which limits its contribution to heat recovery potential when considered alone.
2. Space heating demand shows significant differences between building types. Hospitals stand out as having the lowest PAR, confirming the qualitative observation from Figure 21. They maintain a relatively steady thermal load and are therefore strong candidates for high waste-heat utilisation. Schools and office buildings, on the other hand, show much higher PAR values and strong intermittency, which would inherently limit the achievable ERF.

Overall, the analysis indicates that waste heat recovery from data centres is most attractive in buildings with continuous thermal demand, particularly hospitals or similar 24/7 operated structures. Buildings with intermittent operation would require larger

storage volumes or district-level integration to reach comparable heat recovery performance.

3.4. Impact of the data center size

Figure 22 summarizes how the optimal system design and performance indicators change as the installed DC capacity is scaled down (×0.5) or up (×2) with respect to the standard case. While the standard configuration has been discussed in detail previously, this section focuses on the observed trends resulting from variations in DC size.

From a system-design perspective, the installed chiller capacity decreases as the DC size increases. This counterintuitive result is explained by the optimisation trade-off between chiller capacity and LT thermal storage. For larger DC sizes, the combination of chiller and LT storage becomes increasingly less cost-effective compared to directly valorising larger amounts of waste heat, leading the optimizer to reduce reliance on active cooling. This illustrates how waste-heat integration strategies directly influence the sizing of conventional cooling components in greenfield optimisation.

PV capacity exhibits a moderate dependence on DC size. As the DC capacity increases, the additional electrical demand associated with servers and auxiliary cooling systems incentivises higher PV installations in order to partially offset grid electricity imports. However, the PV capacity is ultimately bounded by the available surface and thus saturates for larger DC sizes. As a result, the marginal contribution of PV to additional DC scaling decreases.

TES capacities scale with DC size, particularly at low- and medium-temperature levels, reflecting the increased availability of recoverable waste heat. Nonetheless, the overall TESU remains low and remarkably stable across all DC sizes. This confirms that, despite increased storage installation, storage plays a secondary operational role and is mainly effective during mid-season periods, rather than enabling substantial temporal shifting at the annual scale.

From an economic standpoint, the TAC increases with DC size. This increase is primarily driven by higher capital expenditures for cooling systems, heat exchangers, and auxiliary components, as well as increased electricity consumption for IT equipment and cooling. While heat imports decrease with larger DC sizes, the economic benefit from reduced heat purchases is insufficient to fully compensate for the added electrical and investment costs.

Interestingly, total annual CO_2 emissions reach a minimum at the standard DC size (20 kW). This non-monotonic behaviour results from a trade-off between reduced

emissions from displaced heat imports and increased emissions associated with higher electricity consumption. For smaller DC sizes, waste-heat recovery is limited, resulting in higher reliance on district heating. For larger DC sizes, additional grid electricity imports dominate, partially offsetting heat-related emission savings.

The DoC increases with DC size, indicating that a larger share of the building heat demand can be covered by locally recovered waste heat. However, this increase is not linear, reflecting temporal mismatches between heat supply and demand, particularly during summer months and off-peak periods. Conversely, the ERF decreases with DC size. While larger DCs generate more recoverable heat in absolute terms, a growing share remains unused due to limited local demand and export constraints, underscoring the importance of demand–supply alignment rather than absolute heat availability.

Overall, optimal waste-heat utilisation emerges from a balance between DC capacity, building heat demand, and the ability to absorb surplus heat.

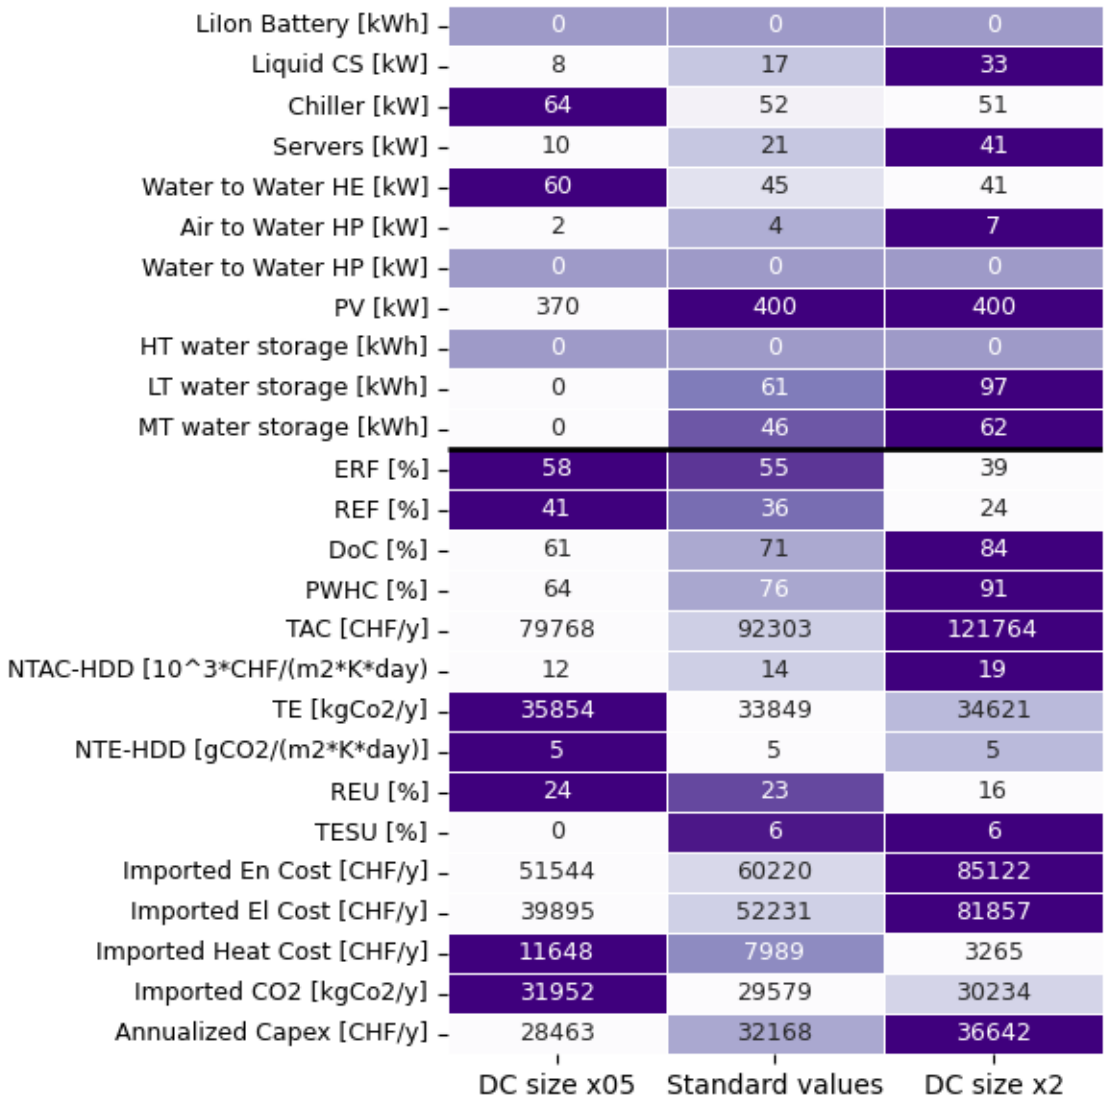

Figure 22 Impact of DC size on optimal sizing and KPIs.

3.5. Impact of location

To evaluate the impact of location on the optimal energy hubs and KPIs, the AAU pilot is also considered. The boundary conditions for the AAU pilot are depicted in Figure 23. The total annual energy demand is larger than the NEST pilot, so the considered IT load was scaled up to match the IT/heating demand peak ratios.

The ambient temperature is lower over the year, with the HDD metric raising from 2452 to 3042 K*day. The heating demand is also more prominent at temperatures above 35 °C (97% of demand requires a supply temperature > 35°C). Regarding energy costs, a DHN cost of 0.099 €/kWh was considered, with a CO2 intensity of 0.1123 kg CO2/kWh. No feed-in tariff for the electricity was considered.

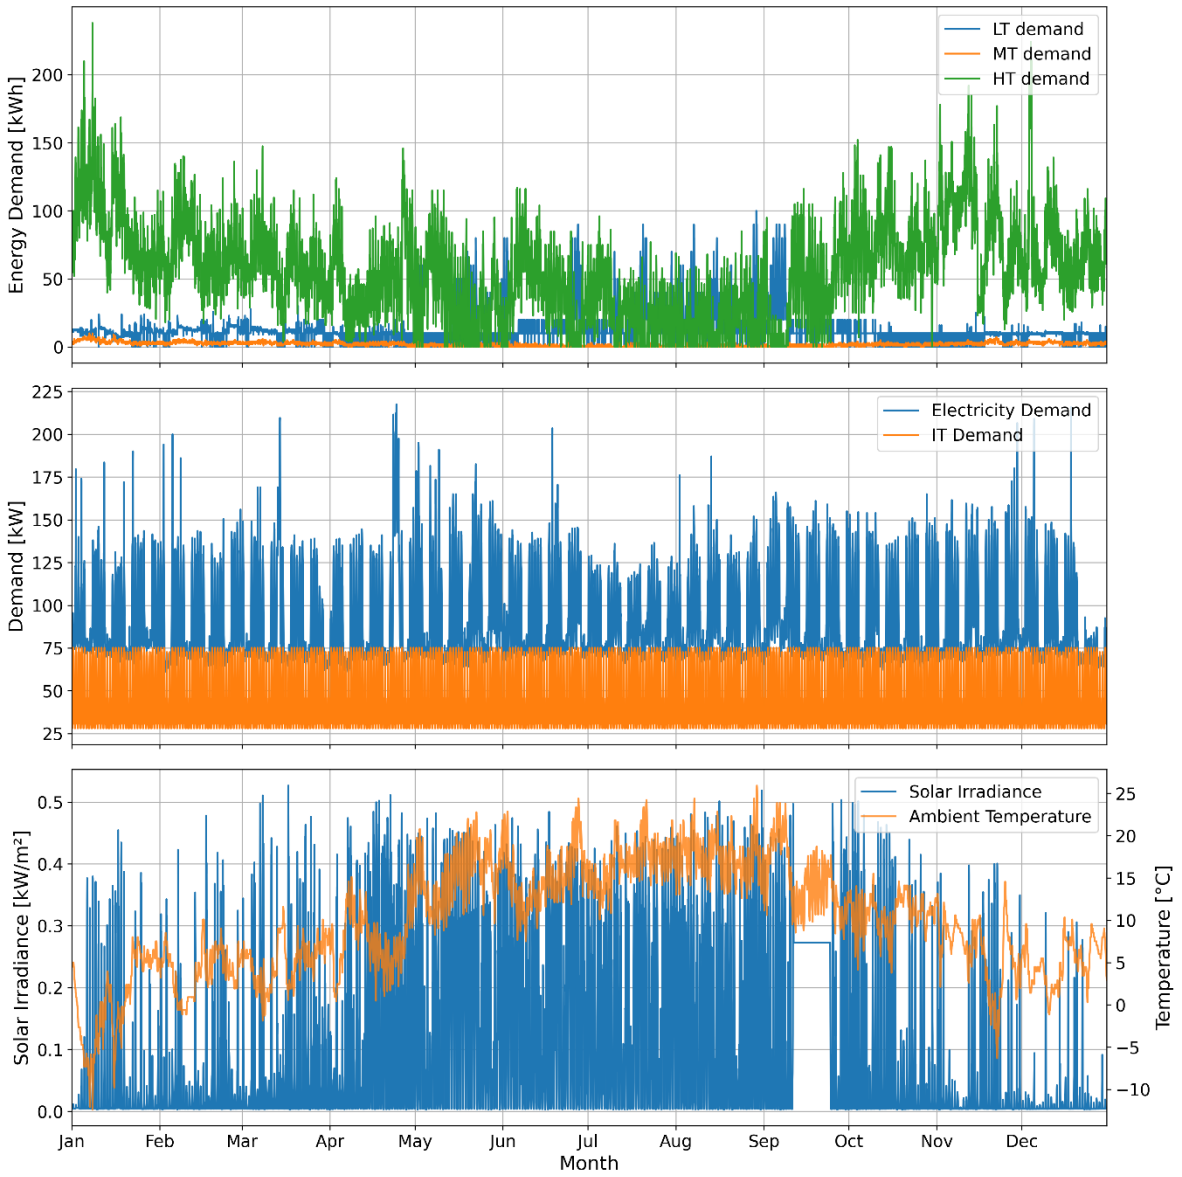

Figure 23 Input data for the optimal design case based on the AAU pilot.

Figure 24 compares the optimal system configuration and resulting KPIs for the energy demands and costs for the NEST and the AAU pilots, highlighting how climatic conditions, demand structure, and system boundaries influence the techno-economic performance of data-center waste-heat recovery.

A first clear difference concerns the role of temperature levels in cooling and heat rejection. In the AAU case, a larger LT storage (200 kWh) is installed, while no MT storage is selected. This reflects a strategy to limit the chiller capacity and operational costs. In contrast, the NEST case relies more strongly on MT storage, as medium-temperature heat is directly compatible with the dominant space-heating demand.

The heating demand profile at AAU is relatively flat over the year, with limited intra-day variability compared to the short-term fluctuations of the IT workload. As a result, short-term thermal storage provides limited additional value: most of the recovered heat can either be consumed almost immediately or must be dissipated due to sustained demand levels. This contrasts with the NEST pilot, where mid-season variability and stronger temporal mismatch between supply and demand create a clearer economic role for MT storage. Further, most of the heating demand takes place at the same temperature level of the DHN, for which a constant tarif is considered.

The different demand structures are also reflected in the sizing of conversion technologies. The AAU case exhibits substantially larger installed capacities for liquid cooling, chillers, servers, and heat pumps. This is driven not only by higher absolute heating demand but also by the requirement to recover heat at higher temperature levels. From an economic perspective, the absolute total annualised cost (TAC) is significantly higher for the AAU pilot, reflecting the larger system size, higher electricity consumption, and increased investment in cooling and heat-upgrade technologies. However, when normalised with respect to heating degree days (NTAC-HDD), the AAU system performs better than the NEST case. This indicates that part of the higher absolute cost is compensated by the colder climate, which leads to a more consistent utilisation of heating infrastructure over the year.

Overall, the comparison demonstrates that waste-heat recovery remains technically feasible and often economically attractive across different pilot sites, but its performance is strongly shaped by temperature requirements and demand variability. Colder climates (AUU pilot or PSNC pilot) with higher temperature heating demands may achieve higher ERF values, but at the cost of increased electrification and emissions. Warmer climates, e.g., TOFAS pilot, require lower heat inputs, but can rely on higher RES availability.

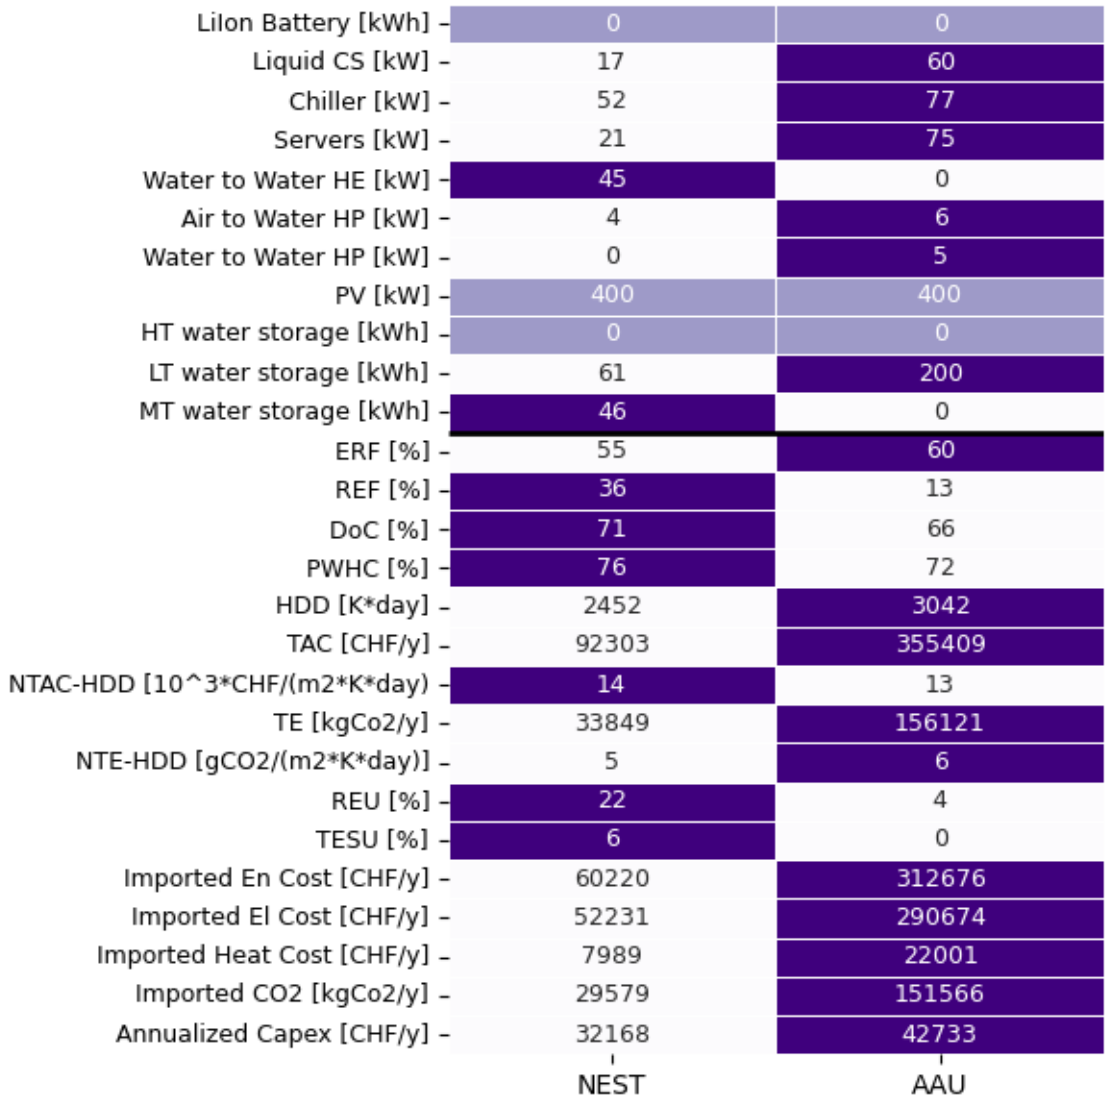

Figure 24 Comparison of optimal sizing and KPIs for the NEST and AAU pilots demands and energy costs.

3.6. On the KPIs calculation

This section discusses how modelling assumptions, system boundaries, and mathematical definitions influence the estimation of key performance indicators (KPIs). In particular, the Energy Recovery Fraction (ERF) is highly sensitive to how recovered heat is accounted for.

According to the European directive [5], the ERF is defined as the amount of energy (heat) that can be reused:

$$ERF = \frac{Q_{DCtoBuilding}}{P_{DC}}$$

For a building with the multi-temperature architecture described in Section 2.3, this definition requires careful interpretation. In the liquid-cooling configuration, heat recovered at a high temperature level is always injected into the HT network and subsequently exported to the district heating network (DHN) if a surplus is present. If DHN exchanges are not explicitly modelled, all heat entering the HT network may be (incorrectly) classified as "recovered," even when it is not actually used within the building. This leads to an overestimation of the recovery potential.

Even when DHN interactions are included, another source of bias arises from ignoring heat losses along the internal heat-transfer chain. Measuring recovered heat directly at the DC outlet, before heat exchangers, storage, and distribution, typically yields values higher than the usable heat that eventually meets thermal demands.

A third source of overestimation comes from temporal aggregation. Computing ERF using yearly cumulative totals implicitly allows surplus summer heat to compensate winter deficits, masking the intrinsic temporal mismatch between supply and demand.

To better quantify such elements, we calculate here three different versions of the ERF as follows:

- Considering all recovered heat as locally consumed:

$$ERF_{max} = \frac{Q_{liq} + Q_{air}}{P_{DC}} = 0.93$$

- Considering heat exports due to heat surpluss but not considering heat losses within the building:

$$ERF_{no\ losses} = \frac{Q_{liq} + Q_{air\_to\_HP} - Q_{air_{to_{LT}}} - Q_{export}}{P_{DC}} = 0.73$$

- Considering heat losses within the building:

$$ERF = \frac{Q_{demand} - Q_{DHN,in} + Q\_\exp - P_{HP}}{P_{DC}} = 0.54$$

The last formulation incorporates losses implicitly by relying on the building-level thermal balance:

$$Q_{demand} = Q_{DHN} + Q_{DC} - Q\_\exp + P_{HP}$$

Under the assumptions adopted in this study, heat delivered by heat pumps is ultimately attributed to the DC (through recovered low-temperature heat), while DHN imports are treated as the residual external supply required to close the heat balance. Heat exported to the DHN appears with a negative contribution, thus penalising hours in which the

building cannot absorb the recovered heat. This formulation ensures that losses and inefficiencies are reflected in the ERF value.

For the Degree of Coverage (DoC), we similarly observed that bounding the index to unity is essential to avoid spurious values:

$$DoC = \min\left\{\frac{Q_{demand} - Q_{DHN,in} + Q_{exp} - P_{HP}}{Q_{dem}}, 1.0\right\}$$

Moreover, DoC must be computed over short time intervals (hourly in this study). Using annual or monthly aggregates would mask the temporal mismatch between heat supply and demand, leading to a significant overestimation of the true degree of coverage.

4. Discussion

The results of this study confirm that waste heat recovery and utilization from edge data centers are cost-effective at the building scale. Among the configurations analyzed, liquid cooling, despite its higher installation cost, consistently yields the best techno-economic performance, mainly because it delivers high-quality heat at ≈65 °C with minimal auxiliary consumption. The higher temperature lift offered by liquid cooling translates into lower dependence on heat pumps and reduced overall system complexity.

Air-based cooling, on the other hand, supplies waste heat at significantly lower temperatures, requiring a dedicated air-to-water heat pump to upgrade the temperature to 35 °C for modern space heating applications. While this introduces additional investment costs, the optimization results indicate that these costs remain acceptable because the heat pump enables reductions in total annualized cost once the recovered heat is integrated into the building energy hub. In other words, the heat pump cost is outweighed by the avoided heat procurement from the DHN and by the reduced need for cooling capacity.

The role of thermal storage strongly depends on the building thermal demand profile and the seasonal timing of heat recovery. In a system with characteristics similar to the NEST pilot, short-term storage, particularly medium-temperature tanks, supports flexibility only during mid-season periods, when short mismatches between supply and demand occur regularly. During the summer, when the building thermal demand is minimal, thermal storage rapidly reaches its capacity limits and excess heat must be exported. Conversely, in winter, the waste heat recovery potential is insufficient to fully cover the heating demand, so storage plays only a marginal role. Increasing the ERF beyond these limits would require long-term or seasonal storage, but the cost-effectiveness and efficiency of such options remain challenging at the single-building scale.

Regarding temperature levels, the analysis shows that DHW can be fully supplied by the liquid-cooled data center in a NEST-like building. At the same time, most of the building's heating demand occurs at the medium-temperature level (around 35 °C) for space heating applications. This implies that the high-temperature heat typically associated with liquid cooling is not strictly necessary to unlock the economic value of waste-heat recovery for edge data centers integrated in buildings. If a lower outlet temperature (but above 35 °C) simplifies the cooling technology or reduces auxiliary consumption, liquid-cooling systems could be operated at lower temperatures without compromising their usefulness for meeting the dominant medium-temperature load. This highlights the importance of

matching WHR strategies not only to the total amount of available heat but also to the temperature requirement of the building's demands.

Finally, the exploration across different building types indicates that the building thermal demand profile is a critical determinant of WHR potential. Buildings with flat, steady heating loads, such as hospitals, exhibit the highest compatibility with continuous waste-heat streams from data centers. Their low peak-to-average demand ratios ensure that a large share of recovered heat can be effectively absorbed and self-consumed without requiring oversized storage or backup systems. Conversely, buildings with highly variable or intermittent heating profiles, such as schools or restaurants, are less suitable unless integrated into larger multi-building networks or district heating, where, in the former, the variability can be smoothed through aggregation.

4.1. Limitations and challenges

Uncertainties in techno-economic data.

CAPEX, OPEX and emission factors for cooling technologies and edge data-center equipment vary significantly across countries, vendors, deployment scales, and years. Even for mature energy technologies, e.g., heat pumps, TES, etc., cost and performance data can differ depending on location, installation conditions, and operating environment. The analysis conducted in this work captures part of this uncertainty through sensitivity analyses on selected parameters, but it does not fully explore the underlying variability. More advanced approaches for decision-making under uncertainty, such as stochastic KPI, further improve the reliability of the conclusions [51].

Simplified system representation without detailed temperature and flow dynamics.

The modelling framework was intentionally designed to scan a wide design space and minimise the number of heuristic assumptions, following common practice in early-stage energy-system design studies [52] [53]. As a result, the system is represented using simplified steady-state relationships that do not capture thermal inertia, pipe heat losses, thermal flow dynamics, supply–return temperature variability, or part-load behaviour of conversion technologies. While this approach is appropriate for conceptual design and technology screening, more detailed physics-based models are needed for final system design and control development. These aspects will be further addressed in Task 7.3 of the HEATWISE project.

Perfect forecast and optimality rather than robustness.

The optimisation assumes perfect foresight of all time-series inputs, including demands, energy prices, PV production, and waste-heat availability. This assumption overestimates the achievable performance and can bias the solution toward highly coordinated operational strategies that would be difficult to replicate in real-time conditions. The results, therefore, represent an upper bound on technically feasible performance. Introducing forecast errors and evaluating robust or receding-horizon control strategies would provide a more realistic assessment of the flexibility potential.

Impact of IT workload

In this study, the IT workload of the data center is treated as an exogenous input and is assumed to be fixed over time. As a result, potential interactions between computational demand and the energy system are not explicitly explored. Including a revenue stream associated with performing IT tasks locally, e.g. [54], together with the possibility to shift or modulate IT workloads in time, could unlock additional flexibility and improve both economic and energy performance. For instance, workload shifting could be used to better align waste-heat availability with heating demand (although section 3.1 highlights limited benefits from short-term flexibility for NEST-like buildings) or to minimize electricity costs under time-varying tariffs. However, such an analysis is currently limited by the availability of representative and transparent IT workload profiles for edge data centers, particularly at high temporal resolution.

5. Key takeaways

The analysis demonstrates that WHR from edge data centres can be both technically viable and economically attractive when integrated into tertiary buildings. The main insights from this study are:

- **Liquid cooling offers the highest overall performance:** Two-phase direct-to-chip cooling provides high-temperature heat that can directly supply domestic hot water and space heating with minimal upgrading. This results in the largest reduction in total annualised cost and the highest energy-recovery factor (ERF). Liquid cooling systems remain cost-effective even when a doubled specific CAPEX per unit of capacity is assumed.
- **Air-cooled systems remain attractive when combined with heat pumps:** Although air cooling alone provides only low-grade heat, upgrading it via an air-to-water heat pump is cost-effective in many cases. The additional investment is compensated by lower operational costs, particularly in buildings with stable medium-temperature heating demand.
- **Flexibility needs depend strongly on demand patterns and system size:** For the NEST-type building, short-term flexibility, primarily via water-based thermal storage, provides benefits during mid-season periods but plays a limited role during summer and winter, when availability–demand mismatches are structural. Overall, ERF improvements from TES is around 0.03 only.
- **Seasonal mismatch limits achievable recovery factors:** In summer, waste-heat availability exceeds the building heating demand, reducing the ERF despite high recovery efficiency at the data-centre level. Without long-term (seasonal) thermal storage or district-level integration, this seasonal constraint remains unavoidable.
- **Domestic hot-water demand is a reliable target for WHR:** DHW loads are stable throughout the year and can be fully met by the heat recovered from the data centre. However, for a NEST-like building, most of the building's heat consumption is at medium temperature (≈35 °C), meaning that high-temperature recovery is not strictly necessary for economic viability. Higher recovery temperatures are instead required for a AAU-like building.
- **Hospitals and similar buildings are strong candidates:** Analysis of hourly heating-demand profiles shows that hospitals exhibit low peak-to-average ratios, i.e., stable demand, which maximises the utilisation of recovered heat and improves cost-effectiveness. Other promising types include care homes and certain office buildings with flat annual loads.
- **The design of the energy hub is influenced by WHR integration:** Allowing heat recovery reduces required chiller capacity and shifts the optimal mix of storage

and conversion technologies. While this is partially an artefact of perfect foresight and optimisation, it highlights strong design synergies between WHR and building energy systems.

Overall, the results indicate that edge data-centre waste heat can meaningfully contribute to decarbonising building heating systems, provided that the cooling technology, building demand profile, and integration strategy are appropriately matched.

5.1. Outlook

The insights from this techno-economic assessment provide several directions for future work within HEATWISE, particularly in Tasks T7.3, T7.4, and across WP9.

The modelling results underline a number of practical aspects that will guide the ongoing development and validation activities:

i. *Identification of key operational challenges*: The study highlights that robust year-round heat-management strategies are essential for all HEATWISE pilots. For the NEST pilot, system operation is relatively straightforward in winter, when heating demand is high, and in summer, when excess heat can be exported to the district heating network. However, mid-season periods emerge as particularly challenging, as heating demand fluctuates and the system becomes more sensitive to control strategies and smart use of short-term flexibility solutions. More generally, the results show that seasonal imbalances between waste-heat availability and building demand significantly limit achievable ERF, regardless of the cooling technology employed. This underlines that WHR planning cannot focus solely on maximising winter heat recovery, but must also address summer surplus management through appropriate operational strategies. Task 7.3 will therefore focus on detailed dynamic modelling and control development to ensure stable and efficient operation across all seasons, with particular attention to transitional periods.

ii. *Implications for space-heating demand and building control:* Since the highest-quality waste heat is produced by liquid cooling, and most building demand is at medium temperature, coordinating building heating control with DC operation will be critical. Therefore, Task 7.3 needs to include also heating and storage elements for the MT grid of the NEST pilot.

iii. *Evidence for the economic viability of WHR:* The optimisation results demonstrate that waste-heat recovery from edge data centres is not only technically feasible

but also cost-effective under realistic conditions. This supports the business-case work in WP9, helping compare WHR pathways across diverse business cases.

iv. *Importance of clear and consistent KPI definitions and calculation*: The study showed that KPI values, such as ERF or DoC, vary significantly depending on how system boundaries and losses are accounted for. System-level KPIs need to be applied consistently across the different WPs.

6. References

[1] E. Cremona and P. Czyzak, “Grids for data centres: ambitious grid planning can win Europe’s AI race,” EMBER, June 2025. [Online]. Available: https://ember-energy.org/app/uploads/2025/06/Grids-for-data-centres-in-Europe.pdf

[2] Y. Hao *et al.*, “Data centers waste heat recovery technologies: Review and evaluation,” *Applied Energy*, vol. 384, p. 125489, Apr. 2025, doi: 10.1016/j.apenergy.2025.125489.

[3] X. Yuan *et al.*, “Data center waste heat for district heating networks: A review,” *Renewable and Sustainable Energy Reviews*, vol. 219, p. 115863, Sept. 2025, doi: 10.1016/j.rser.2025.115863.

[4] X. Yuan, Y. Liang, X. Hu, Y. Xu, Y. Chen, and R. Kosonen, “Waste heat recoveries in data centers: A review,” *Renewable and Sustainable Energy Reviews*, vol. 188, p. 113777, Dec. 2023, doi: 10.1016/j.rser.2023.113777.

[5] “Energy Efficiency Directive.” Accessed: Nov. 21, 2025. [Online]. Available: https://energy.ec.europa.eu/topics/energy-efficiency/energy-efficiency-targets-directive-and-rules/energy-efficiency-directive_en

[6] E. Wheatcroft, H. Wynn, K. Lygnerud, and G. Bonvicini, “The role of low temperature waste heat recovery in achieving 2050 goals: a policy positioning paper,” Dec. 16, 2019, *arXiv*: arXiv:1912.06558. doi: 10.48550/arXiv.1912.06558.

[7] Cerre, “Data Centres & the Grid – Greening ICT in Europe,” 2021. [Online]. Available: https://cerre.eu/wp-content/uploads/2021/10/211013_CERRE_Report_Data-Centres-Greening-ICT_FINAL.pdf

[8] European Commission, “European industrial technology roadmap for the next generation cloud-edge offering,” 2021. [Online]. Available: https://ec.europa.eu/newsroom/repository/document/2021-18/European_CloudEdge_Technology_Investment_Roadmap_for_publication_pMdz85DSw6nqPppq8hE9S9RbB8_76223.pdf?utm_source=chatgpt.com

[9] Group Technopolis, “Europe’s Race for Computing: Progress and Challenges in Edge Computing Infrastructure Deployment,” Technopolis-group. Accessed: Nov. 13, 2025. [Online]. Available: https://technopolis-group.com/europes-race-for-computing-progress-and-challenges-in-edge-computing-infrastructure-deployment/

[10] “Europe Edge Computing Market Size & Trends Report, 2033.” Accessed: Nov. 13, 2025. [Online]. Available: https://www.imarcgroup.com/europe-edge-computing-market

[11] Melgaard, Simon Pommerencke, Jensen, Rasmus Lund, Humbert, Gabriele, and Yilmaz, Cagatay, “D3.3 of the HEATWISE project - KPIs for energy performance and waste heat recovery efficiency from IT infrastructure in buildings,” 2024.

[12] Y. Zhang *et al.*, "Cooling technologies for data centres and telecommunication base stations – A comprehensive review," *Journal of Cleaner Production*, vol. 334, p. 130280, Feb. 2022, doi: 10.1016/j.jclepro.2021.130280.

[13] Q. Zhang *et al.*, "A survey on data center cooling systems: Technology, power consumption modeling and control strategy optimization," *Journal of Systems Architecture*, vol. 119, p. 102253, Oct. 2021, doi: 10.1016/j.sysarc.2021.102253.

[14] E. Oró, P. Taddeo, and J. Salom, "Waste heat recovery from urban air cooled data centres to increase energy efficiency of district heating networks," *Sustainable Cities and Society*, vol. 45, pp. 522–542, Feb. 2019, doi: 10.1016/j.scs.2018.12.012.

[15] X. Yuan, Y. Liang, X. Hu, Y. Xu, Y. Chen, and R. Kosonen, "Waste heat recoveries in data centers: A review," *Renewable and Sustainable Energy Reviews*, vol. 188, p. 113777, Dec. 2023, doi: 10.1016/j.rser.2023.113777.

[16] M. Azarifar, M. Arik, and J.-Y. Chang, "Liquid cooling of data centers: A necessity facing challenges," *Applied Thermal Engineering*, vol. 247, p. 123112, June 2024, doi: 10.1016/j.applthermaleng.2024.123112.

[17] H. Alissa *et al.*, "Using life cycle assessment to drive innovation for sustainable cool clouds," *Nature*, vol. 641, no. 8062, pp. 331–338, May 2025, doi: 10.1038/s41586-025-08832-3.

[18] B. Ott, P. M. Wenzel, and P. Radgen, "Analysis of Cooling Technologies in the Data Center Sector on the Basis of Patent Applications," *Thermal Management*, 2024, Accessed: Nov. 17, 2025. [Online]. Available: https://www.mdpi.com/1996-1073/17/15/3615

[19] H. your say, "King of Hill in data center cooling." Accessed: Nov. 17, 2025. [Online]. Available: https://www.datacenterdynamics.com/en/opinions/king-of-hill-in-data-center-cooling/

[20] O. Van Geet and D. Sickinger, "Best Practices Guide for Energy-Efficient Data Center Design," NREL/TP--5R00-89843, DOE/GO--102024-6283, 2417618, July 2024. doi: 10.2172/2417618.

[21] "ECOQube." Accessed: Nov. 21, 2025. [Online]. Available: https://ecoqube.org/

[22] P. Richner, Ph. Heer, R. Largo, E. Marchesi, and M. Zimmermann, "NEST – A platform for the acceleration of innovation in buildings," *Inf. constr.*, vol. 69, no. 548, p. e222, Dec. 2017, doi: 10.3989/id.55380.

[23] "Waterless Liquid Cooling | Reduce Energy Costs | ZutaCore." Accessed: Nov. 17, 2025. [Online]. Available: https://zutacore.com/

[24] "Energy sector | Federal Statistical Office - FSO." Accessed: Nov. 17, 2025. [Online]. Available: https://www.bfs.admin.ch/bfs/en/home/statistics/construction-housing/buildings/energy-sector.html

[25] ehpa, "European heat pump market report," 2025. [Online]. Available: https://www.ehpa.org/wp-content/uploads/2025/07/EHPA-Market-Report-2025-executive-summary.pdf

[26] FWS, "Jahresbericht," 2023. [Online]. Available: https://www.fws.ch/wp-content/uploads/2024/06/web_fws_jahresbericht-23_de.pdf

[27] I. Keskin and G. Soykan, "Optimal cost management of the CCHP based data center with district heating and district cooling integration in the presence of different energy tariffs," *Energy Conversion and Management*, vol. 254, p. 115211, Feb. 2022, doi: 10.1016/j.enconman.2022.115211.

[28] L. A. Bollinger and V. Dorer, "The Ehub Modeling Tool: A flexible software package for district energy system optimization," *Energy Procedia*, vol. 122, pp. 541–546, Sept. 2017, doi: 10.1016/j.egypro.2017.07.402.

[29] M. Mohammadi, Y. Noorollahi, B. Mohammadi-ivatloo, and H. Yousefi, "Energy hub: From a model to a concept – A review," *Renewable and Sustainable Energy Reviews*, vol. 80, pp. 1512–1527, Dec. 2017, doi: 10.1016/j.rser.2017.07.030.

[30] E. Marzi *et al.*, "Power-to-Gas for energy system flexibility under uncertainty in demand, production and price," *Energy*, vol. 284, p. 129212, Dec. 2023, doi: 10.1016/j.energy.2023.129212.

[31] H.-K. Ringkjøb, P. M. Haugan, and I. M. Solbrekke, "A review of modelling tools for energy and electricity systems with large shares of variable renewables," *Renewable and Sustainable Energy Reviews*, vol. 96, pp. 440–459, Nov. 2018, doi: 10.1016/j.rser.2018.08.002.

[32] M. Macías and J. Guitart, "SLA negotiation and enforcement policies for revenue maximization and client classification in cloud providers," *Future Generation Computer Systems*, vol. 41, pp. 19–31, Dec. 2014, doi: 10.1016/j.future.2014.03.004.

[33] "GLW Tariflblaetter glattpower business," Glattwerk AG. Accessed: Nov. 20, 2025. [Online]. Available: https://www.glattwerk.ch/assets/docs/GLW_Tariflblaetter_glattpower_business_2025.pdf

[34] "District heating connection: Check availability for your location," ewz. Accessed: Nov. 20, 2025. [Online]. Available: https://www.ewz.ch/en/business-customers/heating-climate/energy-solutions/connection-district-heating-district-cooling.html

[35] "Interactive App | Electricity Maps." Accessed: Dec. 09, 2025. [Online]. Available: https://app.electricitymaps.com/map/live/fifteen_minutes

[36] "Downloads," Glattwerk AG. Accessed: Dec. 09, 2025. [Online]. Available: https://www.glattwerk.ch/downloads

[37] B. für E. BFE, "Künftige Rolle von Gas und Gasinfrastruktur in der Energieversorgung der Schweiz." Accessed: Dec. 09, 2025. [Online]. Available:

https://www.bfe.admin.ch/bfe/de/home/versorgung/gasversorgung/gasversorgungsgesetz.html

[38] P. Marocco, D. Ferrero, E. Martelli, M. Santarelli, and A. Lanzini, "An MILP approach for the optimal design of renewable battery-hydrogen energy systems for off-grid insular communities," *Energy Conversion and Management*, vol. 245, p. 114564, Oct. 2021, doi: 10.1016/j.enconman.2021.114564.

[39] "How Long Do Air Conditioners Last?" Accessed: Dec. 10, 2025. [Online]. Available: https://www.lennox.com/residential/lennox-life/consumer/how-long-do-air-conditioners-last

[40] "Capital Cost Analysis of Immersive Liquid-Cooled vs. Air-Cooled Large Data Centers | Schneider Electric." Accessed: Dec. 10, 2025. [Online]. Available: https://www.se.com/ww/en/download/document/SPD_WP282_EN/

[41] B. Monzavi, "Taking a look at the lifespan of Data Centre components - CIRKLA." Accessed: Dec. 10, 2025. [Online]. Available: https://blog.cirkla.tech/2022/05/13/taking-a-look-at-the-lifespan-of-a-data-centre/

[42] W. Torell, "Liquid vs. Air Cooling. Which is the Capex winner?," Schneider Electric Blog. Accessed: Dec. 10, 2025. [Online]. Available: https://blog.se.com/datacenter/architecture/2020/02/24/liquid-vs-air-cooling-which-is-the-capex-winner/

[43] J. Howell, "Breaking Down Data Center Cost: Building vs. Outsourcing," ENCOR Advisors. Accessed: Dec. 10, 2025. [Online]. Available: https://encoradvisors.com/data-center-cost/

[44] S. Zimmermann, I. Meijer, M. K. Tiwari, S. Paredes, B. Michel, and D. Poulikakos, "Aquasar: A hot water cooled data center with direct energy reuse," *Energy*, Jan. 2012, doi: 10.1016/j.energy.2012.04.037.

[45] "CROSSDat – SWEET-CROSS." Accessed: Dec. 10, 2025. [Online]. Available: https://sweet-cross.ch/data/

[46] J. Petersson and Ç. Yılmaz, "D6.1 - HEATWISE assessment methodology and data model," 2025.

[47] G. Humbert, J. Hannson, and B. P. Koirala, "KPIs for an efficient energy management system in buildings with integration of waste heat recovery from data centres - HEATWISE," 2025. Accessed: Dec. 10, 2025. [Online]. Available: https://heatwise.eu/result/kpis-for-an-efficient-energy-management-system-in-buildings-with-integration-of-waste-heat-recovery-from-data-centres/

[48] "JASM - Data platform." Accessed: Nov. 20, 2025. [Online]. Available: https://data.sccer-jasm.ch/demand-hourly-profile-retrofits-cesar/2020-09-17/

[49] M. Fiorentini, P. Heer, and L. Baldini, "Design optimization of a district heating and cooling system with a borehole seasonal thermal energy storage," *Energy*, vol. 262, p. 125464, Jan. 2023, doi: 10.1016/j.energy.2022.125464.

[50] D. Romanov, I. Chakraborty, and S. Holler, “Comparative analysis of scenarios of data center waste heat utilization for district heating networks of different generations,” *Energy Conversion and Management*, vol. 334, p. 119856, June 2025, doi: 10.1016/j.enconman.2025.119856.

[51] G. Mavromatidis, K. Orehounig, and J. Carmeliet, “Design of distributed energy systems under uncertainty: A two-stage stochastic programming approach,” *Applied Energy*, vol. 222, pp. 932–950, July 2018, doi: 10.1016/j.apenergy.2018.04.019.

[52] M. Wirtz, M. Hahn, T. Schreiber, and D. Müller, “Design optimization of multi-energy systems using mixed-integer linear programming: Which model complexity and level of detail is sufficient?,” *Energy Conversion and Management*, vol. 240, p. 114249, July 2021, doi: 10.1016/j.enconman.2021.114249.

[53] M. Sulzer, M. Wetter, R. Mutschler, and A. Sangiovanni-Vincentelli, “Platform-based design for energy systems,” *Applied Energy*, vol. 352, p. 121955, Dec. 2023, doi: 10.1016/j.apenergy.2023.121955.

[54] “Energy Arbitrage: Analyzing Bitcoin Mining’s Historical Revenue per kWh,” Hashrate Index. Accessed: Dec. 09, 2025. [Online]. Available: https://hashrateindex.com/blog/energy-arbitrage-analyzing-bitcoin-minings-historical-revenue-per-kwh/

7. List of figures

8. List of tables