

OFFICIAL



Waste Heat from Data Centres

A Strategic Review for Scotland



Contents

Contents	2
Glossary and Acronyms	4
Executive Summary.....	5
Introduction	8
Purpose and Scope	8
Context: Scotland’s Heat Network and Data Centre Strategy and Policies.....	9
Availability of Heat.....	14
Data Centre Typologies (Size).....	14
Key factors impacting heat availability	16
Cooling System types	17
Water Use.....	20
Estimating Available Waste Heat.....	21
Scale of Heat.....	23
Opportunity and Challenges with Heat Reuse	24
Benefits of Heat Reuse.....	27
Identification of suitable Data Centre sites for waste heat reuse	30
Required Infrastructure	32
Data Centres and Heat offtake interface.....	32
Retrofit vs Inclusion in initial design	35
Heat Network Configurations.....	37
Integration with Local Heat and Energy Efficiency Strategies (LHEES)..	42
Governance, Financing, Commercial Models	47
Background.....	47
Delivery and Ownership Models	48
Tariff and Price Considerations	49
Network Readiness	51
Conclusions.....	53
Recommendations	55
Appendix A – Stakeholder Roadmap	59

OFFICIAL

Appendix B – Case Studies	61
A- Existing Examples.....	61
A-1 – Fjernvarme Fyn (Meta).....	61
A-2 – Old Oak and Park Royal Energy Network (OPEN)	62
Appendix C – Data Centre Classifications	65
Appendix D – Heat Network Zoning	66
Appendix E – Supply Chain Opportunities	68



This report was authored by AECOM, with input from HermeticaBlack along with steering group members Scottish Enterprise, Scottish Government, and Zero Waste Scotland.

We would also like to thank the following organisations for their time and insight on the topic. The perspectives shared have directly informed the development of this report:

Apatura, Deep Green, Star Refrigeration, East Lothian Council, SOSE, Office of the Chief Economic Advisor, Turnbull and Scott, City of Edinburgh Council, North Ayrshire Council, Gren, Exuvi8, Fife Council, NESO, Dumfries and Galloway, North Lanarkshire Council, and SPEN.

Front Cover Image description National Manufacturing Institute Scotland - an example of an ambient temperature district heat network, using waste heat recovered from Scottish Water's Laighpark wastewater treatment works.

Glossary and Acronyms

Acronym	Meaning
AI	Artificial Intelligence
ASHP	Air Source Heat Pump
CAPEX/OPEX	Capital/Operational Expenditure
CDU	Coolant Distribution Unit
COP	Coefficient of Performance
DC	Data Centre
DHN/HN	(District) Heat Network
ESCO	Energy Services Company
FNA	First National Assessment
GIS	Geographic Information System
HNZ	Heat Network Zone
HOT	Heads of Terms
JV	Joint Venture
kW/MW/GW	kilowatt, Megawatt, Gigawatt
kWh, MWh, GWh	kilowatt hour, Megawatt hour, Gigawatt hour
LHD	Linear Heat Density
LHEES	Local Heat and Energy Efficiency Strategy
NESO	National Energy Systems Operator
NPF4	National Planning Framework 4
OPDC	Old Oak and Park Royal Development Corporation
OPEN	Old Oak and Park Royal Energy Network
PHEx	Plate Heat Exchanger
PUE	Power Usage Efficiency
SE	Scottish Enterprise
SG	Scottish Government
SHNF	Scottish Heat Network Fund
SOSE	South of Scotland Enterprise
SPEN	SP Energy Networks
SPV	Special Purpose Vehicle
TES/ATES/PTES	(Aquifer/Pit) Thermal Energy Storage
tRESP	Transitional Regional Energy Strategic Plan

Executive Summary

Purpose and Context

Scotland's current operational data centre base remains small (<30 MW IT load in 2025), but the planning and pre-planning pipeline points to the potential for significant growth; potentially multiple GW of IT load. While a proportion of this pipeline is likely to be speculative, it reflects the rapid growth of AI, high-density computing, and wider digital infrastructure demand. Almost all electrical energy consumed by data centres is ultimately converted to low-grade heat, which is typically rejected to atmosphere. If captured and supplied to heat networks and/or major heat users, this waste heat could support heat decarbonisation and reduce overall energy system demand compared with individual property low carbon and heating solutions.

Policies are already in place that could help support data centre heat reuse in Scotland, including **NPF4, Local Heat and Energy Efficiency Strategies, Heat Network Zones under the Scottish heat networks framework, and emerging AI Growth Zone policy**. However, these policies have not been designed or aligned specifically to realise data centre heat reuse opportunities, and they do not yet create a clear route from data centre development to heat network delivery. While the opportunity may be significant, realising it will require data centre growth to be delivered through a coordinated, policy-led, and spatially aligned approach, bringing together data centre planning, heat network zoning, and early engagement.

Scottish Enterprise commissioned AECOM, working with HermeticaBlack, to undertake this strategic review to assess the opportunity for data centre waste heat reuse in Scotland, identify the key technical, spatial, commercial and policy challenges, and propose actions to support delivery.

Opportunity and Benefits

Data centre waste heat represents a potentially significant and continuous source of low-carbon heat. Depending on size and utilisation, a single large or hyperscale data centre could generate hundreds of gigawatt-hours of recoverable heat annually. At the upper end, hyperscale facilities (100–500 MW IT load) could theoretically provide enough heat to serve tens or even hundreds of thousands of homes, provided sufficient heat demand exists within an economically viable connection distance.

This has the benefit of:

- **Lower whole-system electricity demand** compared with building-by-building heat pumps, by reallocating heat pump demand to centralised infrastructure and reducing pressure on local peak electrical capacity;
- **Low-carbon heat for heat offtakers** with the greatest carbon benefit where data centres are supplied by low-carbon electricity and waste heat displaces fossil fuel or less efficient electric heating;

OFFICIAL

- **Potential for lower delivered heat costs**, where waste heat is priced appropriately and integrated into a heat network in a way that supports affordable tariffs for customers;
- **Wider system value**: Heat reuse can reduce cooling demand and associated water consumption within data centres, while supporting local regeneration, fuel poverty objectives, and place-based decarbonisation.

Key Challenges and Risks

There are **material challenges to realising data centre heat reuse at scale**:

- **Location dependency**: Viable reuse hinges on proximity to heat demand;
- **Access to Grid Infrastructure**: Consideration of electricity distribution capacity is needed as Data Centres must be located where the Grid can support their demand;
- **Lifetime**: Heat networks are 20–50+ year assets; data centres often plan for 15–25 years with evolving cooling technologies. DC heat should not be the sole heat source in large networks;
- **Policy and planning gaps**: Although National Planning Framework 4 (NPF4) and Local Heat and Energy Efficiency Strategies (LHEES), provide an enabling framework, none are yet sufficiently explicit or aligned to systematically deliver data centre heat reuse;
- **Commercial complexity**: Data centre operators prioritise uptime, resilience, and control over cooling systems. They typically prefer “heat-when-available” arrangements, clear system boundaries, and limited commercial risk. In parallel, heat networks must deliver affordable and stable tariffs.

Key recommendations

Without early and comprehensive coordination between major stakeholders, there is a substantial risk that many Data Centres in the planning pipeline will become fully operational before their heat reuse opportunity can be realised. **To minimise this risk and encourage the reuse of waste heat from the expected growth in data centres this report recommends that:**

- The compatibility of data centre locations with local heat reuse opportunities should **be explored**, including proximity to high-density heat demand, anchor loads such as hospitals and universities, and planned heat networks. This should consider the scale and timing of both heat supply and heat offtake, to identify where co-location could reduce network development costs and avoid a mismatch between available waste heat and local demand. Planning authorities could support this through guidance that encourages early assessment of heat reuse potential as part of data centre site selection and development planning;
- Designate an organisation or individual with responsibility to coordinate data centre heat reuse opportunities across Local Authorities (via LHEES), data centre and heat network developers, the Scottish Government, and the Department for Science, Innovation, and Technology within AI Growth Zones. This would include responsibility for aligning growth corridors such as the AI Growth Zones (e.g., North Lanarkshire),

OFFICIAL

NESO's Regional Energy Strategic Plan, and LHEES/prospective Heat Network Zones to prioritise areas; ensuring areas prioritised for Data Centre development are positioned to facilitate heat offtake to major areas of heat demand;

- Local Authorities and Scottish Government should consider developing a community benefit/heat infrastructure fund or offset mechanism where near-term heat reuse is not yet possible, ringfenced for future network delivery;
- Planning and Policy recommendations:
 - Scottish Government and planning teams should coordinate where possible to expand planning policy (via National Planning Framework 4/5 or national planning guidance that builds upon NPF 4/5 – the format of which to be determined, but potentially a letter from the chief planner followed by a Planning Advice Note) to introduce clear, enforceable requirements for data centre developments;
 - Local Authorities should use planning conditions/agreements to encourage heat offtake; and consider limiting high-water-use evaporative cooling to drive closed-loop adoption, as closed loop systems offer a material reduction in water consumption;
 - Scottish Government and Local Authorities should collaborate to define “Green Data Centres” within NPF4 – and the required Heat and Power Plans that Data Centre operators need to provide. These definitions would enable developers to meet minimum planning conditions that support heat reuse earlier in the development process;
 - It is recommended that Local Authorities/Heat Network developers should develop Standard offtake principles (pricing, availability, metering), clear system boundaries, and “no-regrets” investment staging;
- Local Authorities to capture data centre pipeline data (location, scale, cooling approach, heat output) in LHEES. This should be done through engagement with Planning teams that are engaged with Data Centre developers, and through Scottish Government Digital Infrastructure coordinators. Use DC heat supply as a zone-strengthening factor in Heat Network Zone refinement and prioritisation;
- Data Centre operators should future proof data centre design with compatible cooling systems and defined recovery interfaces; led by guidance from Scottish Government/Local Authorities where possible.

The Data Centre heat reuse opportunity in Scotland is significant. However, realising this potential will require a shift from opportunistic development to a coordinated, policy-led, and spatially aligned approach.

Introduction

Purpose and Scope

The emergence of AI and increasing demand for cloud and digital services have led to more intensive computational workloads, driving significant growth in both the number and scale of data centres in recent years. As of 2025, Scotland has a relatively limited operational data centre base with total installed capacity estimated at less than 30MW¹. However, the planning pipeline suggests a potential pipeline in excess of 4.5GW², including some individual applications as high as 500MW per Data Centre. Almost all electrical energy entering a data centre is lost as waste heat, and this is currently not being reused by heat networks or individual heat offtakers.

Scottish Enterprise has commissioned AECOM and HermeticaBlack to undertake a strategic review of waste heat offtake from Data Centres in Scotland.

Heat reuse is the process by which waste heat from a data centre's cooling system is captured, typically via a liquid medium, before it is rejected to the atmosphere. This heat can then be transferred, and where required upgraded using a heat pump, for use in a range of applications including space heating, domestic hot water, and wider industrial and agricultural thermal demands.

The review explores the current potential for heat reuse in Scotland, and what the technical and commercial barriers associated with waste heat reuse from existing and proposed Data Centres are. Growing requirements for electricity and for the decarbonisation of heat are going to compete directly with increasing land, electricity, and water demands of Data Centres – and this report intends to explore the possible constructive collaboration between them.

This report focuses on the **reuse of data centre waste heat within district heat networks**, as this is considered the **most likely/most scalable application** in the Scottish context. Heat networks provide a proven mechanism for aggregating sufficient demand to use the vast amount of low-grade heat from a data centre. Other potential uses of data centre waste heat exist but are **outside the focus of this report**, including applications in industrial processes, agriculture, and aquaculture, which tend to be more site-specific and less readily replicable at a strategic level.

The report is intended to provide guidance for a variety of stakeholders – (Scottish Government, Local Authorities, Data Centre developers and operators, heat network developers and operators, and individual heat offtakers) to facilitate the engagement between these stakeholders, and the potential development of waste heat offtake projects across Scotland.

¹ <https://www.gov.uk/government/publications/estimate-of-data-centre-capacity-great-britain-2024/estimate-of-data-centre-capacity-great-britain-2024>

² <https://www.ercs.scot/news/data-centres-in-the-scottish-planning-system-would-double-scotlands-energy-demands-press-release>

OFFICIAL

The report sets out the scale of the opportunity for leveraging Data Centre cooling systems to supply waste heat and assess the extent this opportunity can be realised. It examines the importance of location, both of data centres and potential heat offtakers, and outlines the infrastructure required to enable heat offtake, including the implications for the electricity grid and overall system energy demand. It also highlights key considerations for delivery, including alignment with local authority plans, potential governance and commercial models, and the role of policy in supporting the uptake of data centre heat reuse in the coming years.

Context: Scotland's Heat Network and Data Centre Strategy and Policies

Heat Network Policy and Strategy Overview

Heat Networks are a way to supply multiple heat offtakers from centralised heat sources – rather than individual boilers/heat pumps supplying individual buildings. They provide a solution to connect large, emerging low-carbon heat sources, such as aquifers, rivers, and major waste heat sources, to heat demands – in particular, to customers that would not otherwise be able to access them individually. There are around 1000 heat networks in Scotland though predominantly small-scale communal networks, and in total heat networks only account for 2% of current non-electric building consumption³. Several more significant networks are noted to be in development⁴.

Scotland's heat network policy sits within a wider heat decarbonisation framework, underpinned by the **Heat Networks (Scotland) Act 2021**. The Act sets targets of 2.6 TWh of heat supplied by networks by 2027 and 6 TWh by 2030, requiring a significant expansion in deployment over the coming years. The Act pairs with the recently introduced draft **Buildings (Heating and Energy Performance) and Heat Networks (Scotland) Bill**. This bill proposes stricter minimum energy efficiency standards and mandates the removal of fossil fuel in buildings by 2045. It offers a legal framework to support Heat Network expansion, and introduces the notion of Heat Network Zones, where powers may be used to require certain buildings to connect to a network or adopt other low-carbon heating solutions. Local Authorities are required to prepare **Local Heat and Energy Efficiency Strategies (LHEES)** on a rolling basis (every five years). These strategies set out long-term plans for heat decarbonisation and energy efficiency and provide the evidence base for identifying areas where heat networks are likely to be technically and economically viable.

This evidence is then used to inform the designation of Heat Network Zones, as defined under the Act and detailed in the **Heat Networks (Heat Network Zones and Building Assessment Reports) (Scotland) Regulations 2023**. Heat Network Zones identify areas where heat networks are expected to represent the lowest cost and lowest carbon solution and are intended to support coordinated deployment and provide greater market confidence.

³ <https://www.gov.scot/publications/heat-networks-delivery-models/pages/4/>

⁴ <https://edinburghcentre.org/news/what-is-a-heat-network>

OFFICIAL

While the zoning framework strengthens the strategic case for heat networks, there is currently no statutory requirement for buildings or heat suppliers to connect within a designated zone.

Additionally, the Scottish Government provides capital funding through the **Scotland’s Heat Network Fund (SHNF)**, a £300 million programme offering grant support to public and private sector organisations developing heat networks. The fund supports new low or zero emission networks, the expansion of existing systems, and the decarbonisation of legacy networks, with funding covering up to 50% of eligible capital costs.

Policy support for heat network connections to new build comes from the **New Build Heat Standard**, where connection to a heat network is treated as a compliant low carbon heating route. While the standard does not directly regulate the carbon intensity of the network the LHEES and funding routes collectively drive networks towards lower carbon heat sources, positioning data centre waste heat as a strong opportunity.

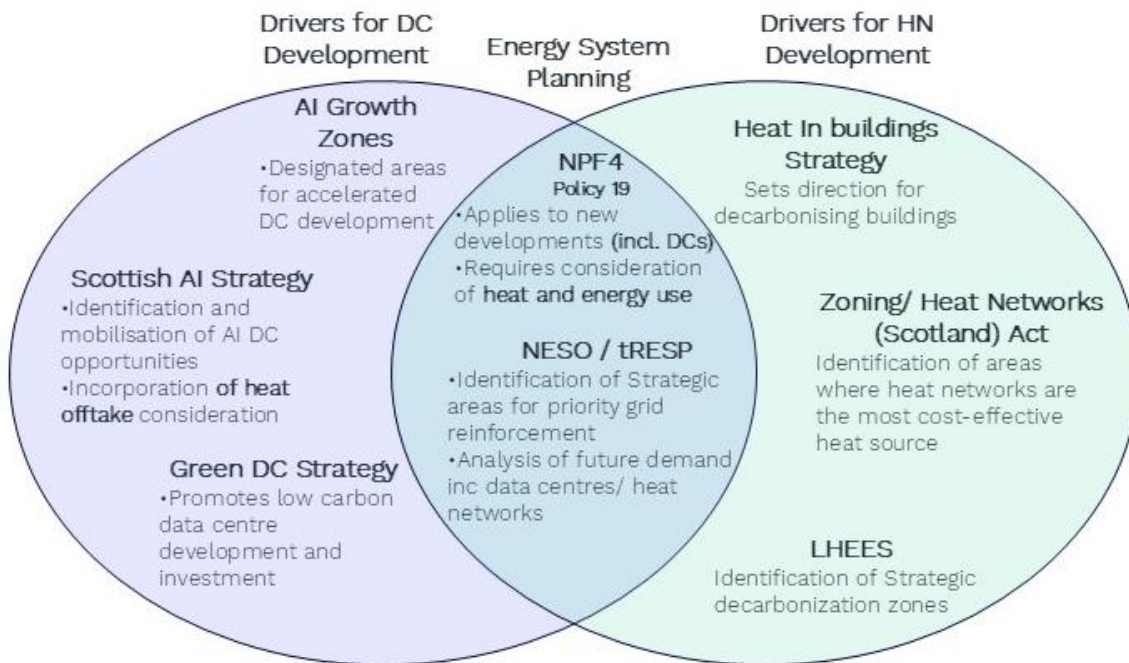


Figure 1. Overview of relevant strategies and policies

Data Centre Policy and Strategy Overview

Scotland is currently an *emerging market* – although Data Centre prevalence is relatively low, its cooler climate, strong renewable energy sources (particularly wind generation which is sometimes curtailed⁵) and availability of land make it an area of interest which increases its attractiveness to data centre developers. To facilitate this, the UK government are assigning selected regions as **AI Growth Zones** – areas to accelerate growth of Data Centres by development of large-scale digital infrastructure by facilitating access to grid capacity, streamlining planning processes, and supporting infrastructure

⁵ <https://www.jbs.cam.ac.uk/wp-content/uploads/2025/03/eprg-wp2503.pdf>

OFFICIAL

delivery. North Lanarkshire has already been designated an AI Growth Zone, with areas such as Fife, Inverness, and Glasgow as contenders for future designation.

Scottish Government have also set out their **AI Strategy for Scotland 2026-2031**, which is essentially a national playbook for turning AI into real economic output and public service improvement, while managing the risks around trust, governance, and public acceptance. In terms of heat there are commitments to:

- Collaborate with partners to identify and mobilise heat offtake opportunities on data centre sites, correlating sites to planned district heat networks by 2027; and;
- Ensure heat reuse, distributed compute and frontier energy solutions are integrated into new developments by 2031.

Scotland's **National Planning Framework 4 (NPF4)** is a long-term spatial strategy and planning policy, setting the framework through which Local Authorities assess and determine development proposals. Within this, **Policy 19 (Heating and Cooling)** introduces a requirement for National and Major developments, including data centres, to produce Heat and Power Plans. These are intended to set out how *“energy recovered from the development will be used to produce electricity and heat”*.

Even relatively modest data centre developments can fall within the National & Major classification. For example, a 10,000 m² facility could range from approximately 5–8 MW for lower density deployments, up to 12–20+ MW for high density /or AI workload sites. However, NPF4 does not currently define what information or detail is suitable for a Heat and Power Plan in relation to heat reuse. Further recommendations on this are set out in [Recommendations](#).

Policy 19 sits within a broader framework of policies that collectively support integrated energy system planning. In particular:

- **NPF4 Policy 1** and **NPF4 Policy 2** establish the overarching requirement for developments to contribute to net zero and climate resilience;
- **NPF4 Policy 11** supports the development of energy infrastructure and low-carbon technologies;
- **NPF4 Policy 18** promotes coordinated infrastructure delivery, aligning spatial planning with energy, heat, and utility systems.

Together, these policies create a framework within which opportunities such as data centre heat reuse can be considered as part of a wider, coordinated energy system.

Grid Capacity Considerations

A key constraint in developing new Data Centres is the lack of available grid capacity. NESO have developed their **transitional Regional Energy Strategic Plan (tRESP⁶)** – to determine areas of priority investment for electrical grid infrastructure upgrades. In Scotland, 32 Grid Supply Point (GSP) areas have been identified as strategically important, representing approximately 16 GW of additional capacity requirements, with 10 areas

⁶ <https://www.neso.energy/what-we-do/strategic-planning/regional-energy-strategic-planning-resp/transitional-regional-energy-strategic-plan-tresp>

OFFICIAL

prioritised for proactive investment by 2034. This includes locations such as Wishaw and Inverarnan, where around 1.6 GW of demand has been identified, driven in part by proposed data centre developments⁷. Future Data Centres and heat networks (see Figure 2) may influence where grid capacity is prioritised and, in turn, where future data centre development is most viable.

This reinforces the need for careful planning and coordination. Data centres should be encouraged where they support wider energy transition objectives and form part of a coordinated energy and infrastructure strategy, including heat reuse, efficient cooling, and grid-aligned siting. This may provide a basis for prioritising strategically aligned data centres over purely commercial developments with limited wider system benefit.

Further assessment of the benefits of coupling DC's and Heat networks for the grid is outlined in [Benefits of Heat Reuse](#).



Figure 2. Existing District and Communal Heat networks⁸ (green, left) vs Shortlisted major planned Data Centre locations⁹ (orange, right) in Scotland

Wider EU Policy

Scotland has an opportunity to build on emerging international practice and strengthen the policy framework for data centre heat reuse. Across the EU, policy and legislation are increasingly being used to improve transparency, encourage energy efficiency, and, in some cases, require data centre operators to consider or enable the reuse of waste heat.

The EU Energy Efficiency Directive now requires data centre operators above relevant thresholds to report key energy and sustainability data to a European database¹⁰, including metrics such as energy consumption, Power Usage Effectiveness, Water Usage Effectiveness, Energy Reuse Factor and Renewable Energy Factor.

Some European countries are going further. Netherlands has implemented provisions through the Dutch Collective Heat Act¹¹. Article 6.2 places obligations on waste heat producers to make heat available at a heat transfer station and limits the basis on which this heat can be charged, focusing on the costs associated with the required equipment and maintenance rather than allowing waste heat itself to become a profit centre. Germany has introduced quantified heat reuse requirements through its Energy Efficiency

⁷ [Strategic Energy Need](#)

⁸ [Edinburgh Climate Change Institute \(ECCI\)](#)

⁹ <https://www.iotinsider.com/industries/ai/scotland-poised-to-become-major-data-centre-hub-amid-surge-in-renewable-powered-investment/>

¹⁰ <https://www.eudca.org/energy-efficiency-directive>

¹¹ <https://zoek.officielebekendmakingen.nl/stb-2026-10.html>

OFFICIAL

Act, with new data centres expected to achieve progressively increasing Energy Reuse Factors, rising from 10% in 2026 to 20% by 2028, subject to feasibility and exemption provisions.

Availability of Heat

Data Centre Typologies (Size)

In this report, data centres have been classified into four typologies. While these are not official industry classifications, they provide a means of differentiating facilities by scale, and characteristics that may influence heat recovery potential.

The relationship between these typologies and traditional data centre business models (e.g. colocation, enterprise, distributed) is set out in Appendix C – Data Centre Classifications.

1. Hyperscale (100MW+)¹²

Large, campus-scale data centres operated by global digital infrastructure providers, supporting cloud services, AI workloads, and other compute-intensive applications. These facilities are typically designed for high utilisation and increasingly high rack densities, may be many multiples of a 100MW's.

2. Large (20–100MW)

Major data centre facilities, often colocation or enterprise-led, serving regional demand, may have a mix of cloud and AI workloads.

3. Medium (5–20MW)

Smaller-scale facilities serving local or sector-specific needs, including enterprise, public sector, or regional colocation deployments.

4. Micro and edge distributed (<5MW)

Highly distributed, small-scale sites located close to end users or data sources, supporting low-latency applications and edge computing use cases.

Table 1 provides an overview of the expected heat reuse potential for each typology.

¹² As of 2025 - all new proposed DC capacity by installed IT load is attributable to hyperscale developments.

OFFICIAL

Table 1. Heat offtake opportunity for each Data Centre typology.

Typology	IT Load (MW)	Available Heat	Offtake Opportunity	Scale of offtaker that could be supported
Hyperscale	>100MW	Very high – large continuous volumes of low-grade heat	Low med: Often located away from existing heat demand, or in areas of planned development	Strategic Heat Mains – Zonal or city scale Heat Networks ¹³
Large	20-100MW	High – of a substantial scale and continuous heat output	Med – more commonly sited closer to urban centres / heat demand.	Strategic Heat Main, Zonal Heat Networks, Med - Large Heat Networks
Medium	5-20MW	Med – lower overall heat volume, but often higher recoverability per MW	Med-high – feasible for co-location with significant demands	Med-Large Heat Networks/ Large + significant individual sites (e.g. Hospitals, universities)
Micro	<5MW	Generally low heat output; higher where liquid cooling is used	High – could be co-located on site with demand	Individual buildings or local loads (e.g. leisure centres, pools, campuses)

Appendix A – provides known and future opportunities for Data Centre heat reuse across these typologies.

The land area of Data Centres depends on the IT load scale above, and the rack density. Micro data centres can occupy less than a few m² of space, with hyperscale data centres occupying areas of 10,000m²+ (the largest exceeding 80,000m²).

¹³ *Strategic Heat Main* – long-distance transmission connecting major sources to multiple networks. >200 GWh/yr.
Zonal Heat Network – networks serving a wider defined zoned area, 50–200 GWh/yr.
Med & Large Heat Network – district-scale systems serving dense clusters of buildings, 10–50 GWh/yr.

OFFICIAL

Key factors impacting heat availability

Whilst broad data centre typologies provide an initial indication of heat reuse potential, the actual quantity and quality of recoverable heat is determined by a range of technical and operational factors. These factors influence both the volume of heat available and its temperature, which in turn affect the feasibility and efficiency of heat reuse within a network.

The key parameters influencing heat availability are:

1. **IT Load** – The total capacity of IT equipment within the site (given in MW or GW). Nearly all electrical energy consumed by servers is ultimately converted into heat, representing the theoretical upper bound of recoverable energy.
2. **Power Usage Efficiency** – Total electrical demand on site is not all recoverable – electricity used for lighting, offices, or pumps is not going to be capturable as heat. The PUE represents the ratio of IT Load: Site load – optimal PUE is 1, with most data centres operating between 1.1 and 1.3.
3. **Cooling System Type** – Cooling system design influences both the efficiency of heat capture and the temperature at which heat can be recovered. Higher-density systems generally require more advanced cooling approaches, which can enable higher-temperature and more efficient heat recovery (see Cooling System types). In addition to these design parameters, several operational and physical constraints further influence the extent to which heat can be recovered in practice.
4. **Utilisation Rate** – The utilisation rate represents the proportion of installed IT capacity that is actively in use over time, and therefore directly determines the quantity of waste heat available for recovery. Utilisation is often quoted as a percentage. The utilisation of a data centre varies on several factors:
 - A. **Cloud/on premises** – Cloud DCs tend to have higher average utilisation by dynamically allocating resources across customers compared to fixed allocation of onsite centres.
 - B. **Workload type** – Data Centres may have different primary workloads for example: AI training, video rendering, or computation simulation. These activities vary in their usage pattern – some are more consistent, and some fluctuate significantly, altering average overall utilisation.
 - C. **Design** - Facilities prioritising redundancy and availability may have lower utilisation rates due to reserved capacity.
5. **Heat recovery factor** – This factor represents the proportion of heat that is physically recoverable (heat demand extracted as a proportion of electricity consumed as computing load). It depends on the ability of the installed equipment to recover heat before it is lost – air-to-air systems may see recovery factors of 10% or less, and a high focus on recovery during design and development can lead to factors as high as 90%.
6. **Spatial and infrastructure constraints** – The scale of heat that can be recovered from a data centre is often limited by the space available for heat offtake

OFFICIAL

infrastructure. This becomes particularly significant for larger facilities (typically 20 MW+), where the size of heat exchangers, pipework and pumping systems increases materially.

If space is not allocated during the design and construction stage, there may never be space for the equipment, directly limiting the volume of recoverable heat. Early engagement between stakeholders is therefore essential to ensure sufficient space is safeguarded and integrated into the overall design.

In the absence of adequate provision, heat recovery may be significantly reduced or not feasible at all.

Implications of factors for heat network development

Collectively, these factors mean that the theoretical heat recovery proportion of a data centre is rarely realised in practice. The quantity and temperature of recoverable heat are highly site-specific and can vary significantly depending on design choices and operational profiles. As a result, robust assessment of heat reuse potential requires early engagement with data centre developers and access to design and operational criteria. At an early-stage conservative assumptions or the use of sensitivity ranges for utilisation and heat recovery are recommended to avoid overestimating available heat.

Cooling System types

Overview

The cooling system ensures IT equipment does not overheat, by collecting heat emitted from the processing units and removing it. It is via the cooling system that waste heat reuse can be realised. The most prevalent external systems are **Air Cooled** or **Water Cooled**.

OFFICIAL

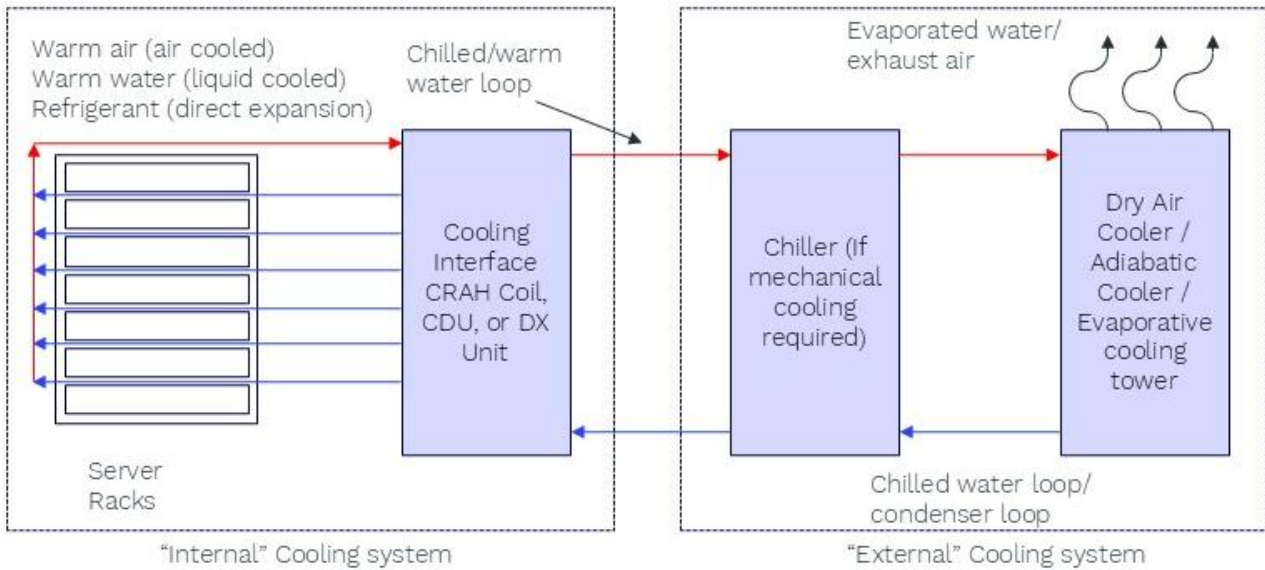


Figure 3. Typical demarcation between internal and external cooling system components for both air and water-based systems

Table 2 outlines the cooling system types and their suitability for heat offtake. Currently, most Data Centres use Air cooling as it is sufficient for lower rack densities prevalent in older designs (Figure 4). However, the emerging higher rack densities used for AI workloads requires significantly greater heat transfer than air is capable of, requiring developers to introduce liquid-cooling systems. Modern liquid cooling systems enable higher heat transfer and can support lower PUE values (typically ~1.1–1.2), particularly for high-density workloads. However, similar efficiencies can also be achieved in optimised air-cooled facilities, with higher PUE values (c.1.4–1.6), generally associated with older designs. Most modern facilities use some combination of both, combining liquid cooling at chip or rack level with air systems to manage residual heat within the data hall.




Figure 4. Relationship between ITE (Information Technology Equipment) rack density and cooling system

From a heat recovery perspective, **closed-loop/liquid-led cooling systems are typically better suited than air-based or direct evaporative cooling systems, as they capture heat in a controlled circuit that can be more readily recovered and reused.** Where heat recovery forms part of the project strategy, data centre developers should therefore consider closed loop liquid or water-based cooling systems wherever feasible.

Table 2. Heat recovery opportunity from major Cooling Systems

Type	Recoverable temperature	Description	Recovery opportunity
Air Cooling	15°C-25°C	Outside air is fed into the data hall via louvres or air handling units (AHUs). In adiabatic systems wet pads can reduce the outdoor air temperature. Can be <i>direct</i> or <i>indirect</i> – the difference being whether outside is directly used, or if it is passed over a heat exchanger to prevent external contamination.	Low to Medium – requires installation of water coils in AHU or louvres to collect heat. Transfer of heat from air to water is inefficient and will require significant additional infrastructure. May increase fan power of cooling system.
Liquid-Air Cooling	20°C-30°C	Chilled water (via air- or evaporative-cooled chillers/cooling towers) cools a circulating air stream, which cools the servers.	Medium to High – presence of a water loop provides direct interface for a heat exchanger and transfer from water to water is more efficient, will still require a boost from heat pumps. – though these would come with greater efficiency than many other sources.
Liquid Cooling	30°C-50°C	Liquid (coolant) makes direct contact with the IT equipment via immersion or cold plates. A Coolant Distribution Unit (CDU) transfers the heat to a chilled water loop to be ejected via a chiller. The most efficient (2-phase) systems use a coolant that is vaporised to capture more heat.	Very High – same as above, usage of water directly allows for easy interface, but immersion also offers higher temperatures due to direct contact with IT equipment.

Use of Free Cooling

Beyond the usage of mechanical equipment to cool, data centres often use ‘free cooling’ - using external air to either meet all or a proportion of the cooling demand. The ability to do this is dependent on the design of the data centre, the rack density and external air temperature & humidity.

OFFICIAL

Free cooling is not inherently incompatible with heat recovery, but its effectiveness in supporting reuse depends on how it is integrated into the overall cooling strategy. Where systems are designed with appropriate heat recovery interfaces, such as air to water heat exchangers, the rejected heat can still be captured and utilised. However, direct air based free cooling (the equivalent of opening a window) would be difficult to recover heat from. There is also a seasonal interaction to consider. Free cooling is most effective during colder periods, which often coincide with peak heat demand. Without a mechanism to recover heat, this can result in useful thermal energy being rejected at the time it is most valuable.

Water Use

Water consumption associated with data centre cooling is an increasing area of concern, particularly where evaporative (open loop) systems are used. These systems can require significant volumes of water to reject heat and are more commonly deployed in regions such as the US, where climatic conditions and design approaches favour their use. In the UK there is a greater tendency towards air-cooled or closed-loop systems, which substantially reduce ongoing water demand. A comparison of indicative water use for a notional 100 MW data centre operating at 75% utilisation is set out below.

Table 3. Comparison of water usage between open- and closed- loop systems (100MW example data centre)

Parameter	Evaporative Cooling	Closed-Loop Cooling
Cooling approach	Open-loop, water lost via evaporation	Recirculating system with minimal losses
Water use for cooling	0.5–1.0 L/kWh	Negligible (closed system)
Annual electricity use	~657 GWh/year (assumed for both cases)	
Annual water use	300-600 million litres/year	~100,000 litres/year
Initial water requirement	Not material (continuous use dominates)	~120,000 litres (one-off fill)
25-year total water use	7.5-15 billion litres	~2.7 million litres
Relative water use	Baseline	<1% of a single year of evaporative use

Closed-loop cooling systems therefore offer a material reduction in water consumption compared to evaporative systems, albeit typically at higher capital cost and with additional technical complexity. Conversely, evaporative systems may be simpler in some respects but require continuous water input and ongoing water treatment to maintain system performance.

OFFICIAL

Emerging alternatives, such as the use of seawater for heat rejection, are also being explored. These approaches have the potential to eliminate freshwater demand but introduce additional considerations, including corrosion, biofouling, and potential impacts on the marine environment.

From a policy perspective, increasing focus on water consumption are likely to drive wider adoption of closed loop cooling. This shift would have a positive secondary effect on heat recovery potential, as closed-loop systems retain heat within a controlled liquid circuit, enabling more straightforward integration with heat exchangers and heat network infrastructure. Local authorities concerned with the environmental impact of data centres should therefore consider introducing planning requirements that limit water consumption or restrict the use of evaporative cooling systems. Such measures would both reduce water demand and encourage heat reuse.

Estimating Available Waste Heat

Table 3 and Figure 5 sets out a simplified methodology for estimating the quantity of recoverable heat from a data centre. It is intended to provide an order-of-magnitude assessment based on a small number of key parameters, rather than a detailed design calculation. The approach reflects typical industry assumptions for PUE, utilisation, and cooling system performance, and can be used to inform early-stage feasibility.

Notes / assumptions:

- The Heat Recovery Factor is highly dependent on cooling system configuration and from where in the cooling system heat can be recovered from;
- The heat pump input electricity contributes additional useful heat. Consequently, the heat that is suppliable to the end user is calculated to be slightly higher than the heat extracted from the waste heat source (depending on the COP of the heat pump). The relationship between these terms is defined by the heat pump coefficient of performance (COP):
- $COP = \frac{Heat_{supplied}}{Electricity_{input}}$. Rearranging:
- $Heat_{supplied} = Electricity_{input} \times COP$
- Since the supplied heat is the sum of the waste heat and the electrical input:
- $Heat_{supplied} = Heat_{waste} + Electricity_{input}$
- Combining these relationships gives:
- $Heat_{supplied} = Heat_{waste} \times \frac{COP}{COP-1}$
- Annual heat supplied assumes a steady average load across the year. More detailed assessments should apply hourly or monthly load profiles where available;
- With an average annual heat demand of 10-15MWh, the waste heat in the example in Figure 5 could supply up to 870 Scottish homes. Due to peaking requirements (10kW per home, occurring at similar times), a diversified heat network would realistically be able to supply approximately 400 homes.

OFFICIAL

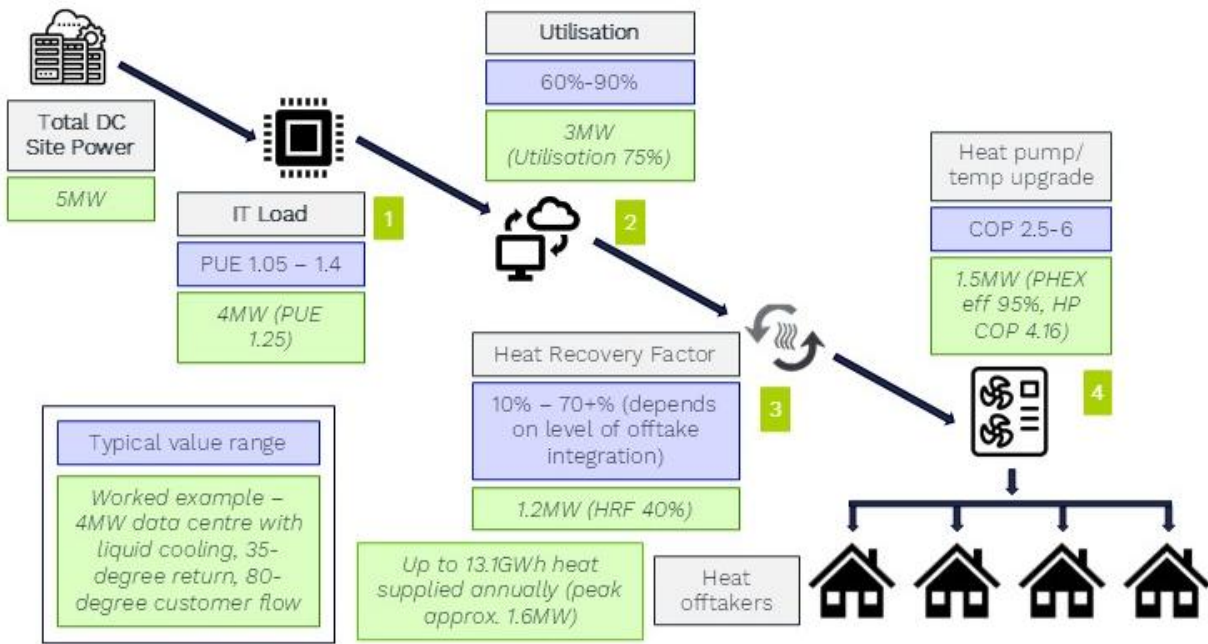


Figure 5. Process for estimating waste heat potential (with 5MW worked example)

Seasonality of Heat

There is a typically a mismatch between heat availability and heat demand. Data centres operate continuously, meaning heat output is relatively stable throughout the year. In contrast, building heat demand varies both hourly and seasonally, with significantly lower demand in summer and peaks in winter.

At a network level, aggregation of multiple customers and thermal storage can smooth short term demand fluctuations, but seasonal variation remains. As a result, not all available heat can be utilised year-round (without significant intraseasonal storage – see Thermal Storage) and additional heat sources are typically required to meet peak winter demand.

For example, a data centre with an IT load of 5MW, providing a relatively constant heat output of around ~1.3MW (as per above example) could meet a considerable proportion of demand for a network with a peak load of 3 MW. However, supplementary heat would still likely be required during peak periods. See Figure 6.

The scale of potential heat recovery is therefore dependent not only on data centre size, but also on the characteristics and diversity of connected heat demand.

OFFICIAL

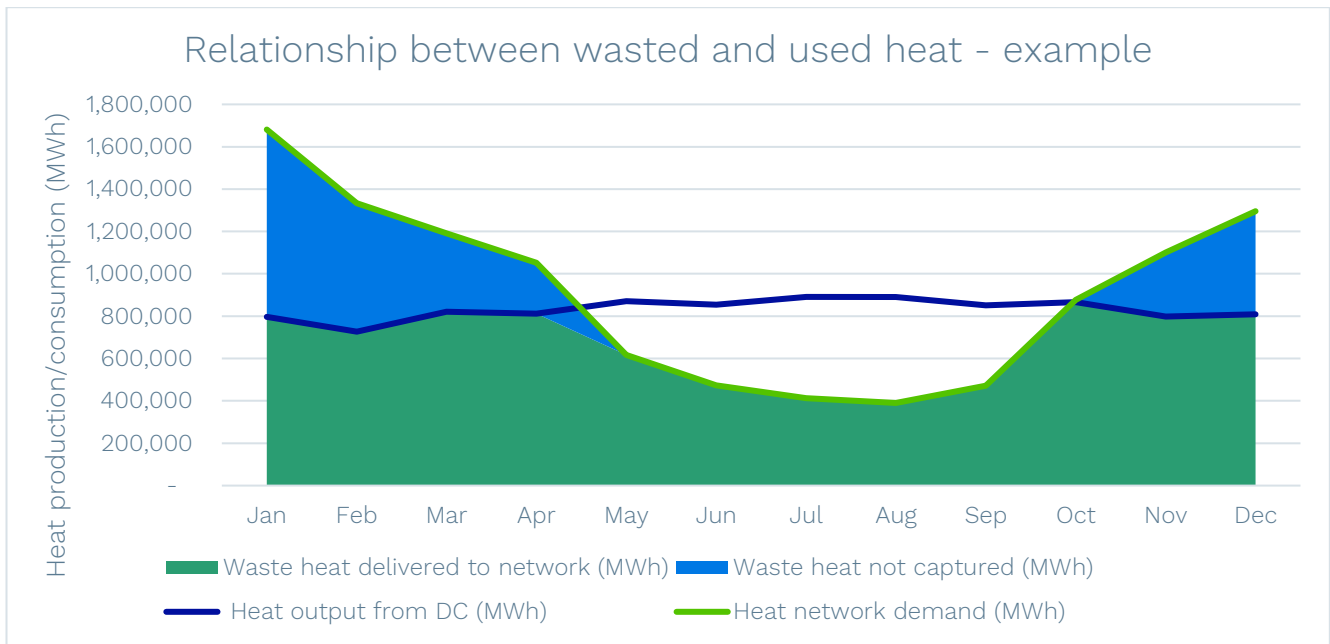


Figure 6. Example heat network - relationship between wasted DC heat and used heat for the network. Note in the summer the Data Centre emits more heat than the network requires - this heat remains unutilised. Further detail on this is provided in Thermal Storage.

Scale of Heat

While a “homes per MW” provides a useful benchmark, actual heat utilisation is strongly influenced by the profile, diversity, and seasonality of connected demand, and will therefore vary on a site-by-site basis.

Nevertheless, Table 4 indicates the number of homes that could be heated by different typology of data centre.

Table 4. Range summary for heat offtake opportunity

Typology	IT Load (MW)	Potential Annual Heat (GWh)	Homes Equivalent
Micro	<5	2–21	200–2,000
Medium	5–20	5–52	500–5,000
Large	20–100	25–260	2,700–24,000
Hyperscale	>100	370–750	49,000–69,000
Upper end Hyperscale	500	1,500–2,000	180,000–240,000

Opportunity and Challenges with Heat Reuse

Proximity to Heat Demands

Heat network viability is primarily governed by the relationship between available heat supply and the density of nearby heat demand. This is commonly assessed using linear heat density (LHD), defined as the annual heat demand served per metre of network (MWh/m/year). As a rule of thumb, schemes typically require a minimum LHD of around 4 MWh/m/year for initial viability, with higher densities supporting stronger commercial opportunities.

In practice, this means that data centre heat reuse is only viable where there is sufficient and proximate demand to use the recovered energy. The cost of distribution infrastructure, including pipework, civils and pumping, increases with distance, and without sufficient demand along the route, longer connections quickly become uneconomic. At typical costs of £3,000-£5,000 per metre, an additional kilometre of pipework can add £3-5 million in capital cost, reinforcing the importance of adequate heat demand density.

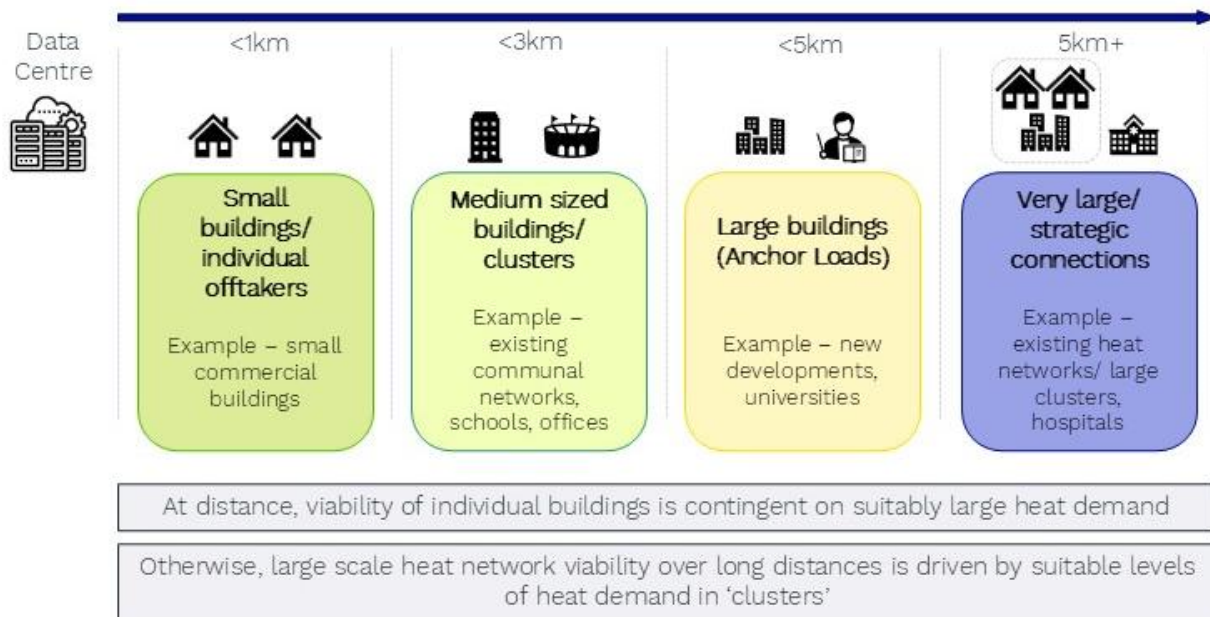


Figure 7. Impact of distance from Data centre on heat reuse viability.

However, this relationship is not solely driven by proximity. Larger heat sources can justify longer transmission distances where sufficient volumes of heat can be delivered, and demand can be aggregated along the route. Strategic Heat Mains, large-diameter transmission networks designed to transport significant volumes of heat over several

OFFICIAL

kilometres, can therefore enable cost-effective long-distance heat transfer in the right conditions.

Even in these cases, viability is determined by the level and distribution of connected demand, rather than the size of the data centre alone.

This relationship is contextualised in Figure 7. As distance to Data Centres increases, the economically viable opportunities move from smaller buildings and individual offtakers to larger offtakers – to ensure the level of heat demand is sufficiently high. **Note “Larger offtakers” includes larger collections of close-by offtakers - clusters of demand.**

The Example - A-2 – Old Oak and Park Royal Energy Network (OPEN) scheme shows that data centre heat reuse is most viable where new development is located in close proximity to the heat source. In this case, planned residential and mixed-use development within ~0.5km enables a linear heat density of ~8 MWh/m/year making network delivery commercially viable. The example highlights that proximity and demand density, rather than heat availability alone, are the critical enablers of successful heat reuse.

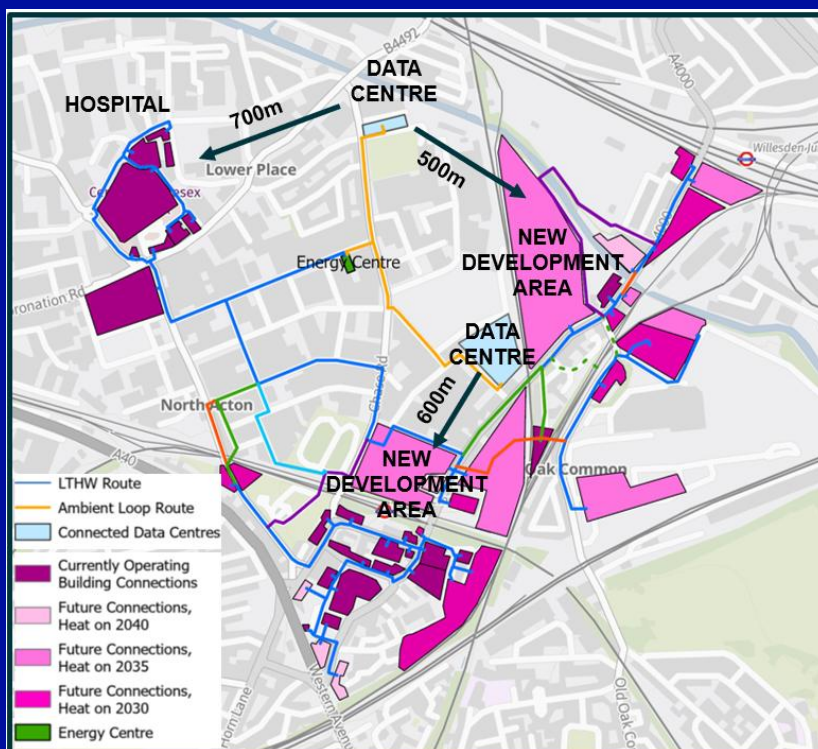


Figure 8. OPDC Case Study with Expansion opportunities

Role of Anchor Loads in Heat Network Viability

Anchor loads are large, consistent heat demands (typically >500 MWh/year as defined in The First National Assessment of Potential Heat Network Zones (FNA)). Anchor Loads are very useful connections for heat networks:

OFFICIAL

- Successful procurement would require confidence in the commitment of a considerable proportion of end user demand – Anchor Load commitment provides this;
- Anchor Loads form significant revenue streams that ensure economic viability;
- Larger revenue means networks can reach further out to Anchor Loads, paving the way for smaller “on route” demands, further increasing viability.

Early and proactive engagement with prospective Anchor Loads is essential to the successful development of a Heat Network. However, their own decarbonisation targets may raise conflicts with the Data Centre heat reuse opportunity:

- Independent decarbonisation e.g. on-site heat pumps can undermine the economic viability of any future heat network by removing the major offtaker;
- “Cherry picking” – A Data Centre supporting nearby major loads may reduce the viability of an independently developed heat network. Engagement between the Operator and heat network developer is critical to prevent network opportunity sterilisation.

Interconnection with nearby, or purposefully co-located, micro and medium data centres can provide access to a consistent low-grade heat source, supporting decarbonisation and enabling compute to be deployed where heat can be effectively reused. Engagement with NHS estates teams indicates clear interest in this approach. Hospitals are particularly well suited, as large, single-site heat users with a strategic focus on utilising nearby low carbon heat sources where available. Alignment between data centre development and NHS facilities could therefore support deployment of these smaller, distributed “edge” data centres, but this engagement should be held in tandem with wider area stakeholders (Local Authority, possible DHN developer) to ensure this does not sterilise wider network opportunities.

An aligned approach in selecting a Data Centre location between the developer and Councils is imperative to mitigate these risks.

Trenching and Infrastructure Constraints

In dense urban environments such as Glasgow or Edinburgh, where underground space is already occupied by a high density of utilities (e.g. water, gas, electricity, telecoms, and drainage), the physical space required for large heat network pipes and associated trenching can become a significant constraint. Street works requirements under the Specification for the Reinstatement of Openings in Roads (Scotland), and the presence of existing services can further limit available corridor space for new infrastructure.

This challenge is exacerbated when considering low temperature or ambient pipework required to transport heat between a DC and a heat pump (see Required Infrastructure). These sections of pipe operate at lower temperature differentials (ΔT), requiring significantly larger pipe diameters to transfer the same amount of heat. This can result in wider trench requirements and increased spatial constraints. For these reasons, distances between the data centre and heat pumps should be minimised wherever possible. Note

OFFICIAL

that pipe is required both sides of the heat pump – but after temperature uplift, the diameter of the pipe decreases significantly, so costs and impacts decrease.

Infrastructure Lifespan

Heat networks are typically designed as long term infrastructure assets, with operational lifetimes of 40–50 years or more¹⁴. In contrast, data centres are developed with shorter lifespans of around 15–25 years, with major refurbishment or replacement occurring within this period¹⁵.

This misalignment introduces uncertainty when integrating data centre heat into long term heat network planning. In addition, evolving cooling technologies and operational changes may alter the temperature, availability, or recoverability of waste heat over time, making long term supply of heat difficult to guarantee.

As a result, heat supply agreements are typically limited to shorter durations (e.g. ~20 years), and reliance on a single data centre as a primary heat source may introduce long term risk. Data centre heat can form a primary component of a heat network; however, heat networks should be supported by a diversified heat supply mix to manage long term risk.

Benefits of Heat Reuse

The benefits of heat reuse are realised at a whole system level, benefitting the data centre, electricity network, and heat offtakers. It enables the recovery and redistribution of otherwise rejected energy, improving overall system efficiency and reducing reliance on standalone heat generation. Key benefits include:

Reduced overall electricity demand

Compared to the counterfactual solution (e.g. standalone ASHPs), heat reuse can reduce both overall system electricity demand and peak electricity requirements, as higher source temperatures improves the heat pump efficiency.

This is illustrated in Figure 9, which expands on the example in Figure 5 and considers the combined impact of data centre led heat networks serving low carbon residential demand. The example shows that the same heat demand is met with materially lower electrical input, reducing both peak load and overall grid impact.

¹⁴ CIBSE Journal (2017). *Burying an investment, not an expensive problem*. Available at:

<https://www.cibsejournal.com/technical/burying-an-investment-not-an-expensive-problem/>

¹⁵ Evaluation Ltd. *Managing Critical Mechanical and Electrical System Replacement in Data Centres – A Business Challenge*.

Available at: <https://www.evaluation.co/news/managing-critical-mechanical-and-electrical-system-replacement-in-data-centres-a-business-challenge>

OFFICIAL

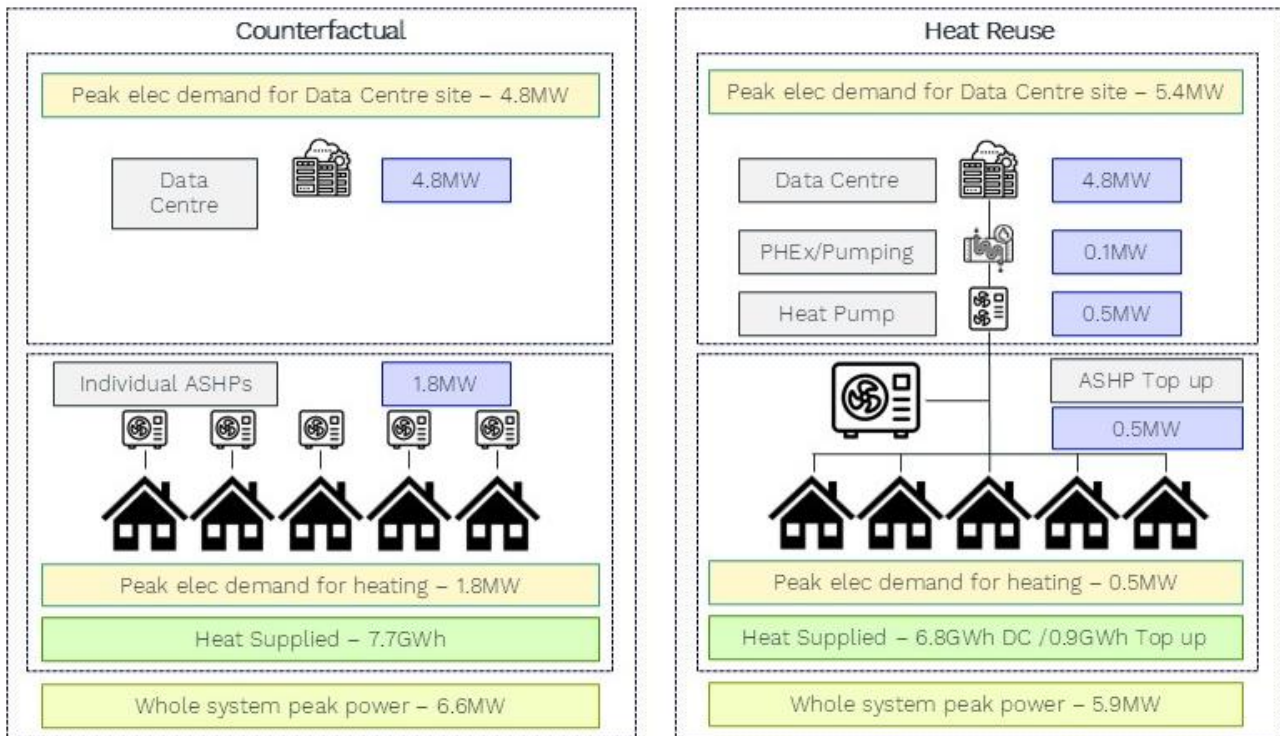


Figure 9. Impact of heat reuse on system peak electricity demand. Example heat network requiring 7.7GWh of heat annually

The impact of this connection is also summarised in Table 5.

Table 5. Peak and annual electricity savings due to Data Centre Heat Reuse

	Counterfactual	Heat Reuse	Saving
Heat source	Individual ASHPs	DC waste + centralised ASHP	-
System peak electricity demand (MW)	6.6MW	5.9MW	10%
Data Centre annual electricity consumption (GWh)	29.95GWh	29.49GWh	2%
Annual electricity consumption for heat (GWh)	3.64GWh	2.33GWh	36%
Annual heat supplied (GWh)	7.7GWh	7.7GWh	-

Reduced cooling demand and water use

Typically, a Data Centre would retain their own cooling plant and use it only when heat reuse does not sufficiently cool the system. Operationally, heat reuse would result in

OFFICIAL

reduced cooling demand and associated water use, as heat rejection is partially transferred to the network rather than dissipated via air or evaporative cooling systems.

Lower cost heat

The cost of heat supplied from a heat network is influenced by factors including upfront infrastructure CAPEX, financing arrangements and the volume of heat sold. Access to low-cost heat sources, such as data centre waste heat, can improve the overall economics of heat networks. Higher source temperatures than that of equivalent environmental sources (rivers, air etc.) would allow for higher COPs, and therefore lower heat pump running costs.

The price of heat depends on commercial arrangements between the data centre operator and the heat network developer. However, lower generation costs can support more affordable and stable heat pricing, helping to improve the viability and scalability of heat networks (see [Tariff and Price Considerations](#)).

Alignment with European Policy

In recognition of the high energy use associated with data centres, several international policies have been introduced to improve efficiency and encourage the reuse of waste heat.

The European Union Energy Efficiency Directive (EU EED) includes reporting obligations for larger data centres (≥ 1 MW), incorporating the ISO 50600-4-6 standard and the Energy Reuse Factor (ERF), which measures the proportion of energy reused outside the facility. The Directive requires operators to assess and use waste heat where it is technically and economically feasible.

The Energy Reuse Factor (ERF) measures the proportion of a data centre's total annual energy consumption that is reused outside the facility. It is defined as:

$$ERF = E_{Reuse} \div EDC$$

where E_{Reuse} is the energy reused externally, and EDC is the total data centre energy consumption, including IT equipment and supporting systems such as cooling and power infrastructure.

Germany's Energy Efficiency Act (2023) goes further by introducing minimum reuse targets, requiring data centres to reuse at least 10% of energy, rising to 20% by 2028, subject to certain conditions.

While these requirements are not directly applicable in Scotland, they reflect a clear direction of travel towards the utilisation of data centre waste heat. Integrating heat reuse within developments would therefore align with emerging international best practice and help future-proof schemes against potential regulatory tightening.

OFFICIAL

Identification of suitable Data Centre sites for waste heat reuse

This checklist sets out factors to consider when identifying potential data centre locations in the context of waste heat reuse. It is intended to support early opportunity screening only, as a sense check of whether a location may have the characteristics needed to facilitate heat offtake.

If most answers are “Yes” the site is likely to be a strong candidate for further assessment; if several are “Uncertain” targeted feasibility work is required; and if multiple are “No”, particularly in relation to proximity to demand, the likelihood of viable heat reuse may be low.

Table 6. Checklist for identifying data centre locations in the context of heat reuse

Factor	Screening Question	What to look for (Guidance)
Electricity Network Capacity	Is there sufficient grid capacity (or committed reinforcement) to support the data centre and associated heat recovery plant (PHEX or Heat Pump)?	Areas with existing capacity or planned upgrades. Consider whether additional load from heat pumps can be accommodated. Consider if Data Centre would be a barrier to further heat decarbonisation (exacerbates grid constraints).
Alignment with existing Heat Networks	Is the site located near to an existing or in-development heat network?	Consult with Local Authorities on emerging and developed heat network plans.
Alignment with LHEES / Heat Network Zones	Is the site located within or near an identified heat network opportunity area?	Review local LHEES and Delivery Plan. Prioritise areas identified for heat networks or with recognised anchor loads.
Proximity to Heat Demand	Is there significant heat demand within a viable connection distance?	Typically, <1–2 km straight line distance to major loads or large clusters. Consider both distance and route complexity. Long connections (>2–3 km) and complex crossings may undermine viability unless connecting to significant demand.
Heat Density	Is there sufficient linear heat density to support a network?	Target ~2-4 MWh/m/year or higher. Look for dense urban areas or clustered developments.

OFFICIAL

Anchor Loads and new developments.	Are there large, consistent heat users nearby?	Hospitals, universities, leisure centres, large residential or mixed-use developments. These improve load factor and project viability.
Cooling System Compatibility	Is the proposed cooling system compatible with heat recovery?	Preference for liquid or water-based systems with recoverable heat at useful temperatures. Air-only systems without recovery are less favourable.
Temperature Suitability	Is the recoverable heat temperature suitable for reuse (or viable for uplift)?	Typically, ~20–30°C requiring heat pumps. Newer liquid-cooled systems may be able to provide higher temperatures (~30–60°C+), improving efficiency and reducing uplift requirements.
Planning / Policy Support	Is there local authority support or policy drivers for heat networks or low carbon heat?	Planning policy, net zero targets, Local Development Plans and other Local Authority Climate/Economic strategies, or potential constraints on water use / cooling systems.
Site Constraints / Routing	Are there any major physical or consenting barriers to connecting to demand?	Rivers, rail, highways, utilities congestion, land ownership, and designation (e.g. national park; SSSIs; conservation areas; etc).

Required Infrastructure

Data Centres and Heat offtake interface

Interfacing Infrastructure

A schematic between a Data Centre and heat network is shown in Figure 10. The required infrastructure is:

- Cooling system - Waste heat needs to be transferred from the return of the cooling loop (after heat has been collected from the servers) to a liquid system;
- Water coil – if an air-cooling system, a liquid coil is needed in the Air Handling Unit (AHU);
- Plate Heat Exchanger (PHEX) –Hydraulic separation between DC-side cooling loop and heat network-side primary pipes;
- Heat Pump – For upgrading the temperature of the waste heat (and for hydraulic separation) – can be used with or instead of a PHEX;
- Pumps – To push the fluids around each loop;
- Controls/metering – To monitor temperature, control flow rates, and valves;
- Isolation valves – Ensure ability to disconnect/stop fluid for maintenance.

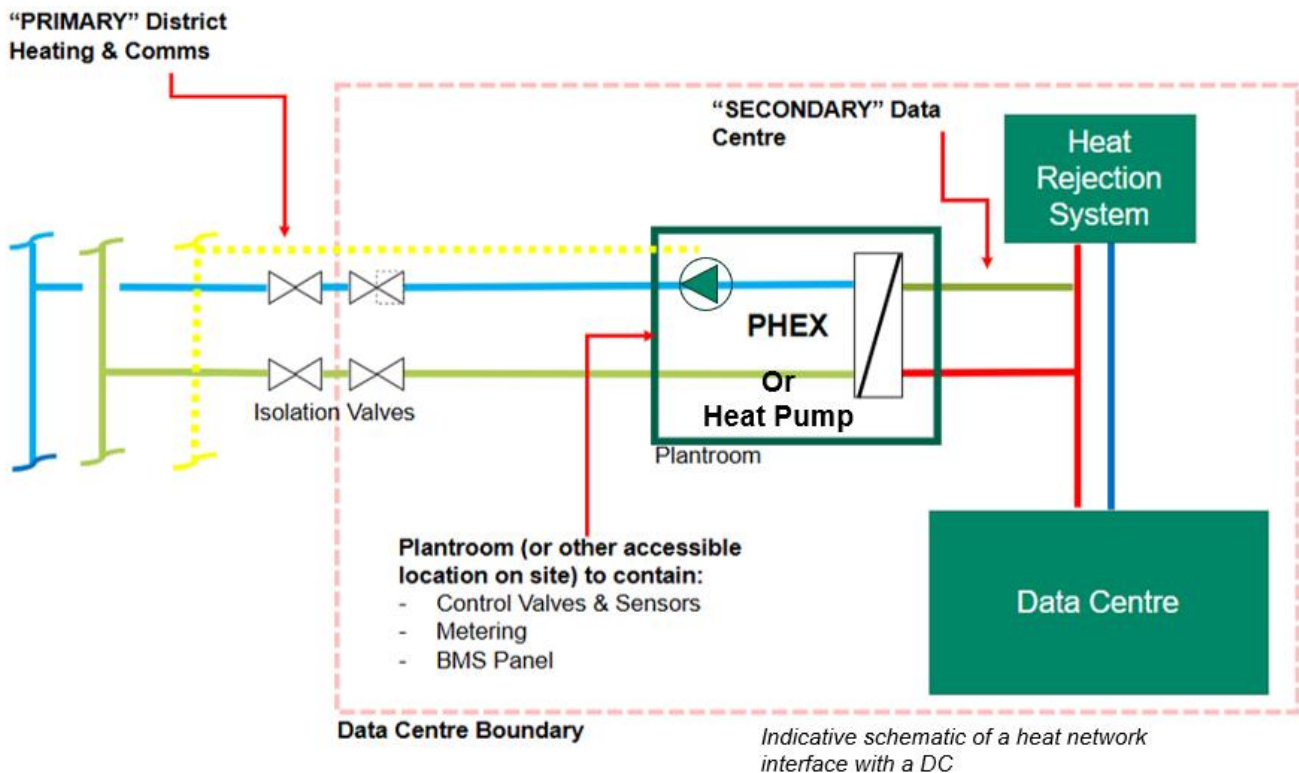


Figure 10. Interface between Data Centre and Heat Network

OFFICIAL

Technical options for connection to Cooling System

There are two options for interfacing between the DC cooling system and a heat network– **Heat Transfer Substations**, and **Chiller/Heat Pump Coupling**.

Heat Transfer Substation - Plate Heat Exchangers enforce a hydraulic separation between the Data Centre and the heat offtaker – this also often forms a commercial ownership boundary. The return from the DC cooling system is passed across the PHEX, where waste heat is transferred to the network-side, for uplift and/or controlling and distribution.

To minimise thermal losses and reduce the extent of primary pipework, the substation should be located as close as practicable to the data centre cooling system. Where the district heating operator owns or maintains the equipment, suitable access arrangements will also be required.

The size of the PHEX installation is driven by heat duty and temperature differential (ΔT). Individual units are typically 1–3 m in height and 0.4–1.2 m in width, with larger schemes (10MW+) using multiple units in parallel to provide flexibility.

Substation space is therefore largely determined by the number and arrangement of PHEX units. Redundancy is typically incorporated (e.g. N+1) to ensure continuity of operation. A high-level approximation of total space take (footprint) is detailed in Figure 11.




Figure 11. Indicative dimensions for Plate Heat Exchanger Substations (m²)

Chiller/Heat Pump Coupling - Instead of upgrading the heat at an offsite energy centre, a Water Source heat pump can directly connect across the DC cooling system and a heat network. This can facilitate the recovery of heat from the cooling system directly, as the heat pump can be used as the primary cooling system, upgrading heat instantly and

OFFICIAL

delivering to end users- removing the need for the PHEx. In the case of Figure 12, the “Water Source” would be the cooling loop in the data centre.




Figure 12. Heat Pump across a Data Centre and Heat Network

Compared to a PHEx-only solution, heat pump installations require more space due to the inclusion of compressors, heat exchangers, and associated plant. For larger schemes, multiple units are typically installed in parallel, resulting in a plantroom-scale installation. Additional space may also be required for buffer storage and pipework interfaces.

Heat pumps located on the Data Centre site would introduce an additional electrical load to the Data Centre development, typically requiring around 20–35% input electricity relative to the heat delivered, depending on performance. Where this equipment forms part of the Data Centre-side infrastructure, this additional demand would need to be accounted for in grid connection capacity. By contrast, where heat pumps are located off-site, the associated power supply whilst still required, would typically sit with the heat network operator rather than the Data Centre developer.

Lastly while heat pumps can contribute to cooling, the data centre must retain independent heat rejection capability. As such, parallel systems (e.g. chillers or dry coolers) are still required to ensure resilience when heat offtake is unavailable. Some purpose-built solutions are starting to become available that integrate cooling and heat recovery functions into a single system.

Comparison of interface options

Engagement with Data Centre developers in Scotland and case studies suggests that Heat Transfer Stations are the preferred interface arrangement for developers. This is due to the relative ease of establishing a clear commercial boundary and the reduced technical complexity of the interface. Futureproofing for a plate heat exchanger connection requires fewer modifications to a conventional cooling system than direct integration with heat reuse infrastructure.

OFFICIAL

Table 7. Comparison between PHEX Connection and Integrated HP Connection

Aspect	Plate Heat Exchanger Heat Transfer Substation (PHEX)	Chiller/Heat Pump Coupling
Description	Hydraulic separation between data centre and heat network via heat exchanger	Heat pump directly integrated with data centre cooling system and heat network
Integration / Retrofit	Low complexity, minimal impact on existing cooling system	More complex, fundamental change in cooling strategy
Efficiency	Minor temperature and efficiency losses across PHEX	Higher efficiency with fewer interfaces and losses
Cooling Operation	Independent of heat demand (cooling always available)	Dependent on heat demand; requires backup cooling (e.g. chillers)
Commercial / Ownership Boundary	Clear separation between data centre and network	Greater onerous on DC to maintain and use heat pump. Heat provided may come at a cost
Infrastructure Requirements	Requires space for heat transfer station on-site	Greater space take compared to PHX but substitutes for chiller
Cost Implications (CapEx, OpEx)	Some limited additional equipment and interface costs. Low OpEx	Greater up-front CapEx and long term OpEx, however can substitute for chiller

Retrofit vs Inclusion in initial design

The ability to integrate heat reuse infrastructure is a key determinant of whether data centre waste heat can be used. This is strongly influenced by whether provision is made at the design stage or whether integration is attempted retrospectively.

The heat reuse infrastructure is most effectively delivered when designed into the data centre from the outset, or where provision is made to allow for future connection. While policy (e.g. NPF4 and Heat Networks (Scotland) Act 2021) encourages the use of waste heat, the market remains immature and most facilities are not yet future-proofed.

Retrofit solutions are generally feasible but come with trade-offs in cost, performance, and operational risk. They are typically more expensive than integrating heat recovery from the outset and often fall short on the comparative heat-recovery potential. Crucially,

OFFICIAL

retrofitting requires interfacing with live cooling infrastructure, and unless systems were designed with this in mind, integration can necessitate partial or full shutdowns. This combination of higher cost, constrained performance, and disruption risk drives both complexity and reluctance to connect. Table 8 provides an overview of the key considerations, challenges, and benefits when considering retrofit and future proofing.

Table 8. Impact of retrofit vs futureproofing

	Retrofit	Future proofing
Equipment and Installation Cost	Potential additional costs for bespoke infrastructure or longer pipe sections if space not suitably allocated	Relative lower cost – design considers necessary components and sizes in advance
Time to install	Medium – additional measures for access and security will be needed	Low – programmed in as part of overall construction
Disruption to Operations	High – Depending on the location of the cooling system, the Data Centre may need to be completely shut down temporarily	Low – installation of equipment can be built into programme for developing site
Spatial allocation	Medium – substation and pipework must use what space is available and accessible	Good – considerations for access and disruption should have already been accounted for
Impact on DC energy usage	Good – will reduce work done by chillers, but depends on what points of cooling system can be accessed	Best – location of interface can be selected to optimise this
Electrical Supply	Additional capacity may need to be sought for the substation as not considered in initial Plans	Capacity can be considered during design of DC

Designing for heat offtake upfront is typically lower cost and less disruptive than retrofit. However, developers are often reluctant to invest in infrastructure without certainty of

OFFICIAL

use. To address this, local authorities, through LHEES and relevant energy planning should be encouraged to provide visibility of network plans and zonal designations.

The approach to futureproofing depends on the cooling system. As a minimum, **provision should be made for future connection points** on chilled water circuits to enable integration of plate heat exchangers. Space should also be safeguarded for a substation. A more initiative-taking approach would include installing key components (e.g. PHEX, pumps, controls), subject to developer appetite and project certainty.

Heat Network Configurations

District heat networks are significant infrastructure that transport heat from a central source to multiple buildings. Heat networks are often classified into “generations”, reflecting their supply temperatures, system architecture, and approach to heat delivery. The most common in current deployment are:

- **3rd & 4th Generation** – Centralised heat production with a distribution network circulating hot water at a flow temperature of typically ~50–60°C (4th generation) and <100°C (3rd generation);
- **5th Generation** – Ambient temperature networks (~10–40°C) with decentralised local heat pumps, enabling both heating and cooling.

Existing buildings utilising gas boilers use temperature regimes typically consistent with 3rd generation systems. Fabric upgrades and other efficiency measures could be introduced to offtakers to make them compatible with lower flow temperatures, which would have a positive impact on the efficiency of the heating plant. However, these upgrades are typically expensive and disruptive and require all offtakers to undertake them to make such a network feasible – introducing potential issues in commitment, affordability, and programme.

The requirements for connecting data centres to heat networks differ between high-temperature (3rd/4th generation) and low-temperature (5th generation) systems, primarily due to whether temperature uplift is required before the heat reaches the customer.

OFFICIAL

3rd / 4th Generation Networks

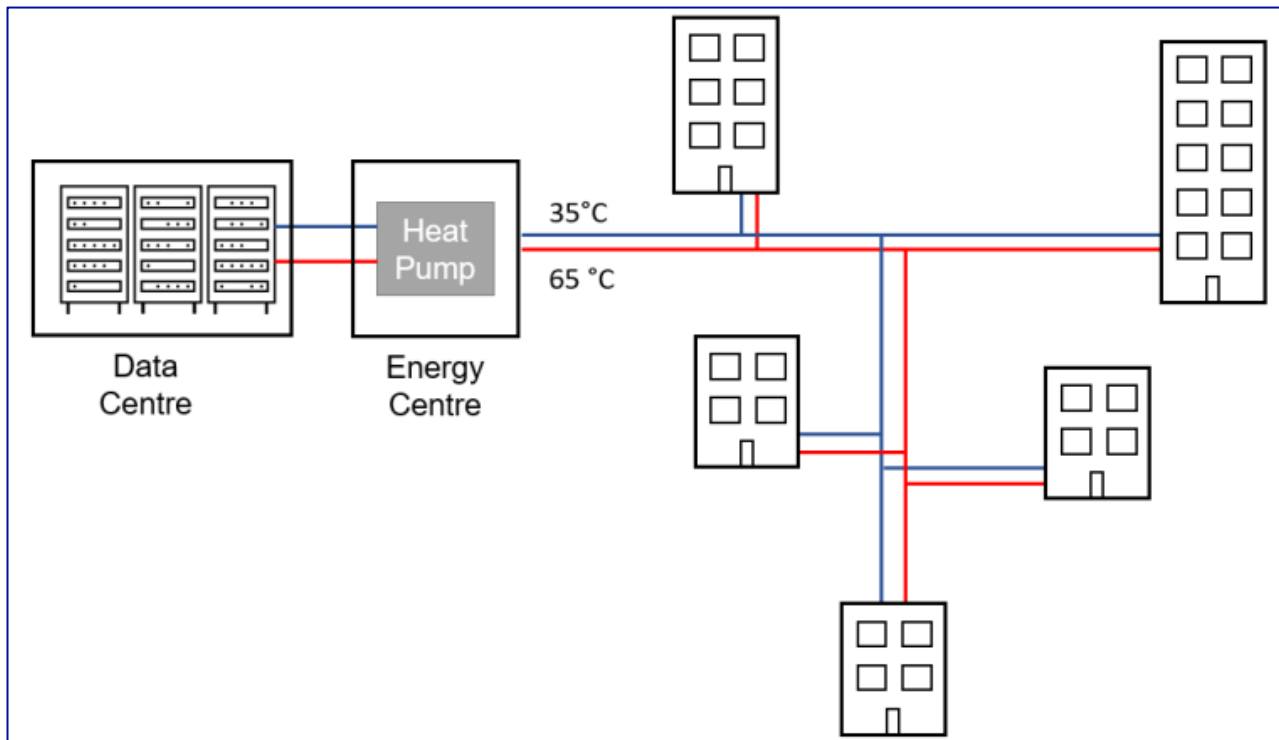


Figure 13. 3rd & 4th Generation Heat Network (indicative F/R temperature)

Data centre waste heat (typically 20–40°C) must be upgraded to network temperatures before it can be used. This is typically achieved via water source heat pumps, located within a centralised or third-party energy centre rather than at the data centre itself.

From a data centre perspective, connection requires:

- A heat transfer interface (plate heat exchanger) or heat pump (see Interfacing Infrastructure);
- Hydraulic connection to the cooling system;
- Provision for heat export at stable flow and temperature conditions.

Implications:

- Additional system-level infrastructure (e.g. heat pumps, pumps, controls);
- Increased overall system cost and electrical demand (though not necessarily borne by the data centre).

OFFICIAL

5th Generation Networks

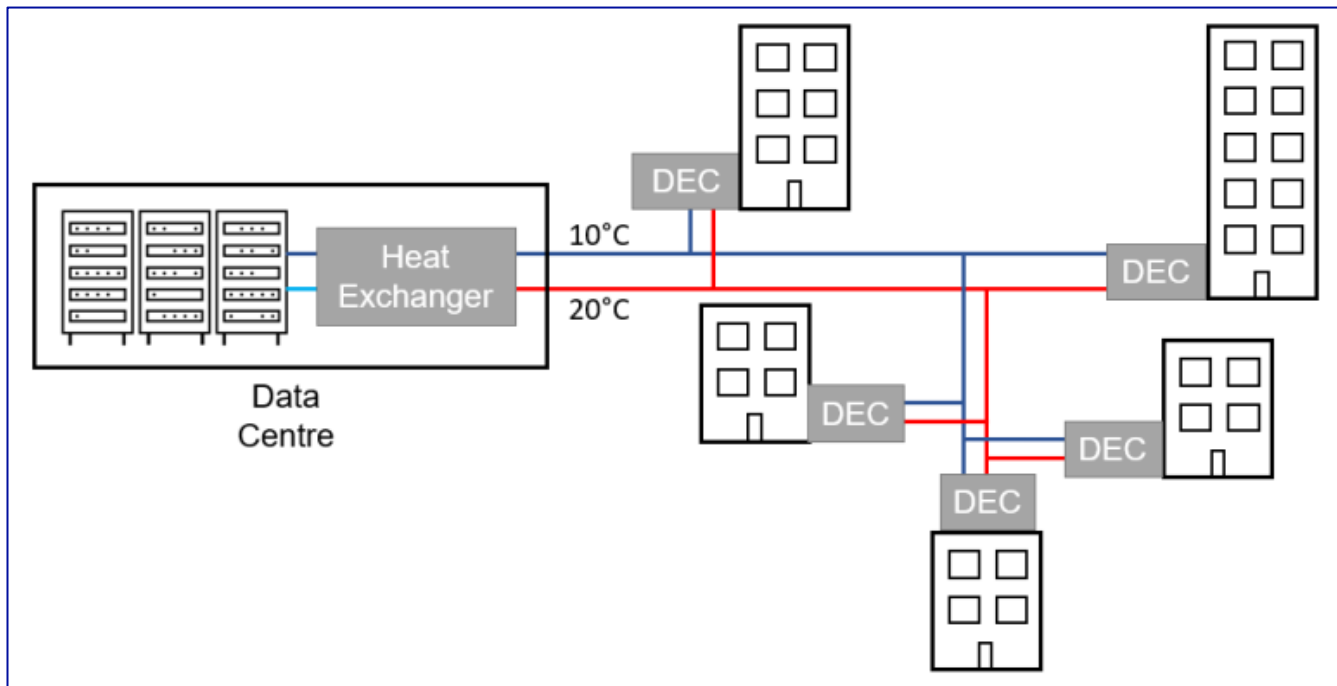


Figure 14. 5th Generation Heat Network (DEC – Decentralised Energy Centre)

These systems operate at temperatures aligned with data centre waste heat, meaning no central temperature uplift is required.

From a data centre perspective, connection typically requires:

- A heat transfer interface (plate heat exchanger);
- Hydraulic connection to the cooling system;
- Heat pumps are instead located at the point of use (buildings), rather than within a central energy centre.

Implications:

- Larger pipework to transfer heat at lower temperatures and differentials. These pipes are more expensive per unit length than their smaller counterparts;
- Building level heat pumps and centralised or decentralised pumping. Not all existing buildings may be able to take additional equipment (architectural constraints, listed building etc).

OFFICIAL

Key Differences

The primary distinction is that:

- 3rd/4th generation networks require temperature uplift somewhere in between the Data Centre and the offtakers – preferable closer to the data centre (often offsite from the datacentre);
- 5th generation networks do not require temperature uplift prior to distribution but may necessitate significant works to existing building systems to integrate localised heat pumps.

From the data centre perspective, the physical connection requirements are similar, but the wider system infrastructure and complexity differ significantly.

Thermal Storage

Thermal storage is a core component of heat networks and is most commonly deployed as short-duration buffering, typically over timescales of hours to days. This is usually achieved through above-ground hot water tanks. Such systems enable surplus heat to be captured and stored for later distribution when heat supply is insufficient, as well as to optimise the operation of heat pumps, for example by shifting heat generation to periods of lower electricity prices.

In the context of data centre heat reuse, short-duration thermal storage can also be supplemented with longer-duration solutions to better align heat supply and demand. Unlike many conventional heat sources, data centres generate waste heat continuously throughout the year and are far less subject to seasonal variation (such as river or air source).

Heat demand within connected heat networks, however, is highly seasonal, with significantly higher loads in winter than in summer. Where viable inter-seasonal thermal storage could be incorporated to support improved utilisation of recovered heat on an annual basis.

As the waste heat is low-grade (typically 20–30 °C), it can either be stored at this temperature or boosted to district heating temperatures (70 °C+ if not 5th gen) then stored. The scale of heat available, and (for the “pre-boost” storage option) the large storage volume required, is enabled by large-scale, long-duration storage solutions, including:

- Large pit thermal storage;
- Minewater (with thermal storage) systems;
- Aquifer Thermal Energy Storage (ATES).

These options offer:

- Sufficient scale to store large volumes of low-grade heat, and substantial opportunity for high-grade heat;

OFFICIAL

- Lower spatial impact compared to surface tanks, noting that pit thermal energy storage, as a surface-based solution, can require substantial land take;
- The ability to store heat over months, rather than days.

Note that storing high grade heat would be preferable due to its greater energy density but would be subject to a greater level of temperature degradation in the long-term without adequate insulation/ground warming, making it more difficult to store seasonally. Scotland is well suited to this approach, with large areas of land that could be developed for Pit Thermal storage – shown to be very effective at scale in places such as Denmark.

There is further potential in disused mine workings (e.g. Midland Valley) and aquifers, which can function as natural thermal stores with low heat loss. Councils, Universities, and private institutions (University of Edinburgh, Mining Remediation Authority) in Scotland are exploring the ability of the disused mines below Scotland, to store and supply water for heating – noting that this concept is not yet fully proven.

While “small scale” Thermal Storage is already routinely considered in the development of heat networks, certain geographic conditions such as the presence of former mines, geothermal features or available land may, in limited or future cases, offer additional flexibility for the management of surplus data centre waste heat.

Integration with Local Heat and Energy Efficiency Strategies (LHEES)

Relevance of Data Centre Waste Heat to LHEES

Each local authority in Scotland has developed a Local Heat and Energy Efficiency Strategy (LHEES), setting out long-term plans for improving building energy efficiency and decarbonising heat. LHEES form part of Scotland's statutory heat planning framework under the Heat Networks (Scotland) Act 2021 and are subject to review on a five-year cycle. While all local authorities have now produced their LHEES, **the methodology is currently under review and may be updated for the next cycle, potentially leading to adjustments to the process described in the following sections.**

LHEES aim to improve building energy efficiency, reduce fuel poverty, and decarbonise heat. Within this framework, heat networks are a key delivery mechanism, enabling the use of low-carbon and waste heat sources that are not accessible at the individual building level. Data centre waste heat can support heat networks, and more broadly, LHEES considerations by:

- Reducing emissions through the displacement of fossil fuel heating;
- Improving heat supply stability due to consistent baseload generation (environmental sources, for example, may provide reduced heat in the winter);
- Supporting affordability, particularly in areas affected by fuel poverty.

Overview of LHEES Development Process

This section summarises how the assessment of data centres as potential waste heat sources can be incorporated within local heat planning in Scotland. As shown in Figure 15 the recommended LHEES methodology follows an eight-stage process that moves from strategic review and data collection through to strategy development and delivery planning. It should be noted, however, that this methodology was not mandatory; while some local authorities followed it closely, others adapted elements of the approach.

OFFICIAL

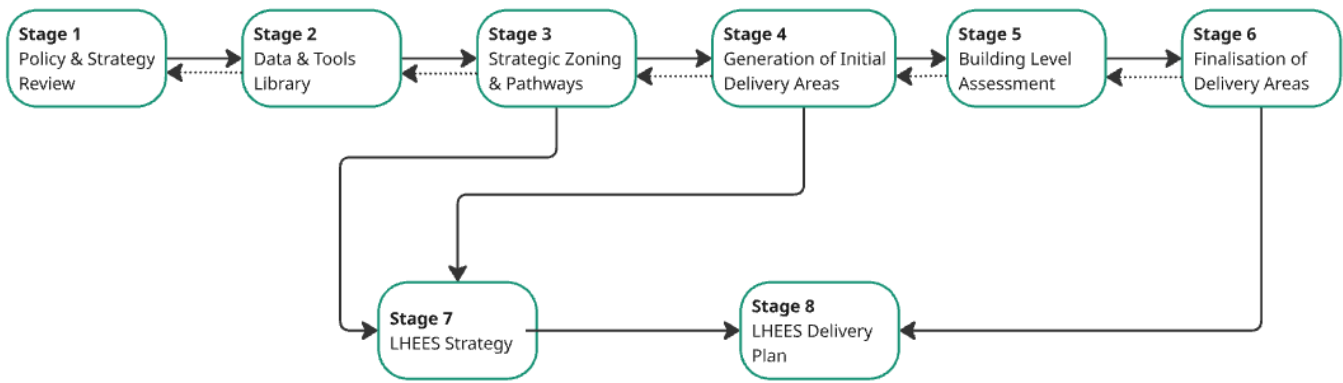


Figure 15. Summary of LHEES Stages

The LHEES process is outlined below:

- **Stage 1:** Reviews the relevant policy and strategic context;
- **Stage 2:** Compiles the datasets, tools, and evidence required for the analysis;
- **Stage 3:** Undertakes strategic zoning and identifies potential pathways;
- **Stage 4:** Generates initial delivery areas, including consideration of constraints, resources, and new developments;
- **Stage 5:** Undertakes building-level assessment;
- **Stage 6:** Finalises delivery areas;
- **Stage 7:** Prepares the LHEES Strategy;
- **Stage 8:** Develops the LHEES Delivery Plan.

Data Centre Waste Heat Opportunities within the LHEES Methodology

In the LHEES Methodology: Version 4.0,

- Strategic consideration of energy resources occurs during Stage 4, where constraints, resources and new developments are incorporated into the analysis;
- Section 2.7.5 “Constraints, resources and new developments” explains that heat network analysis enables the strategic consideration of factors such as development sites identified within Local Development Plans and potential sources of renewable or waste heat, allowing these resources to be integrated into the identification and refinement of potential heat network delivery areas.¹⁶

While the LHEES methodology may be updated or refined in future iterations, the overall approach is likely to continue recognising the importance of opportunities to use waste heat. In this context, data centres represent a potentially valuable source of recoverable heat and should be considered alongside areas of heat demand when exploring heat network opportunities.

¹⁶ Scottish Government (2022). *Local Heat and Energy Efficiency Strategies (LHEES) Methodology Version 4.0*. Zero Waste Scotland.

OFFICIAL

As part of the LHEES process, engagement and collaboration with key stakeholders are essential to develop a comprehensive understanding of opportunities and constraints across the local authority area. Notably:

- This engagement should occur early in the LHEES process and continue throughout its development, ensuring that relevant information is incorporated into the evidence base;
- Obtaining spatial location data from data centre developers not available in national datasets can determine whether a data centre is located in proximity to areas of significant heat demand and therefore represents a potential waste heat resource for district heating networks;
- Where potential heat network zones are identified in proximity to significant waste heat sources, engagement with operators can help establish the approximate quantity and grade of heat available, as well as any drivers or constraints associated with supplying a heat network.

Influence of a Major Heat Source such as a Data Centre on Heat Zone Designation

Within the LHEES analytical process, the presence of a major heat source such as a data centre does not determine the initial identification of Potential Heat Network Zones. These zones are identified primarily using demand-based indicators such as linear heat density, anchor loads and clusters of heat demand¹⁶.

- The detailed assessment stage then considers whether the identified Potential Zones may be particularly suitable for heat network development;
- This assessment considers a range of factors, including the potential use of renewable and waste heat sources, existing infrastructure, major heat demand buildings, and wider development constraints¹⁷;
- The outputs generated provide supporting evidence for stakeholder discussion and decisions on whether Potential Zones should be progressed to designation, which sits outside the formal remit of LHEES.

In this context, the presence of a data centre may strengthen the case for designation by demonstrating that a zone could access a significant and consistent waste heat source, particularly where the facility is located near clusters of heat demand or anchor loads.

Integration of DC Waste Heat within the LHEES and Heat Network Zone designation Process

Figure 16 illustrates how data centre waste heat considerations can be incorporated across the Local Heat and Energy Efficiency Strategies (LHEES) and Heat Network Zone (HNZ) designation process. The diagram presents a simplified overview of the core stages, with green boxes highlighting where specific actions relating to data centres may be

¹⁷ Scottish Government (2024). *Heat Networks – Detailed Practitioner Approach LHEES Stage 4 – Generation of Initial Delivery Areas*, Section 3.4, and Step 4.

OFFICIAL

applied. A more detailed summary of the Heat Network Zoning process is provided in Appendix D – Heat Network Zoning .

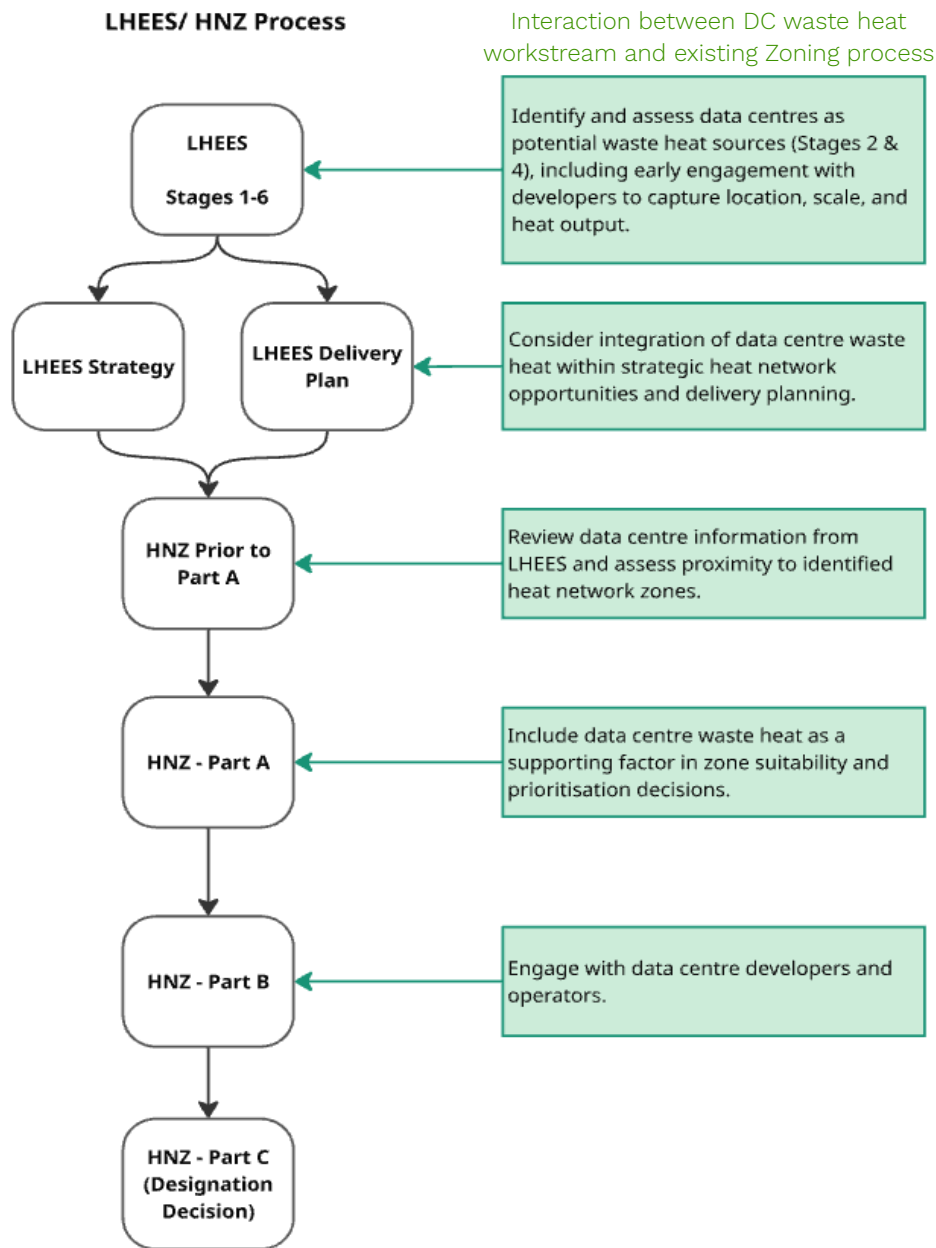


Figure 16. Integration of Data Centre Waste Heat within the LHEES and Heat Network Zoning (HNZ) Process.

Recommendations for Future LHEES Development

LHEES are refreshed every 5 years - whilst data centre developments can come forward at any point within that cycle. There is a risk that significant new waste heat opportunities may emerge after the relevant LHEES evidence base has been prepared and may not be appropriately considered until the next formal update. A more flexible approach is needed to ensure that major new heat sources, including (but not limited to) data centres, can be considered in a timely way. As a recommendation:

OFFICIAL

- New data centres emerging between LHEES review cycles could trigger a proportionate opportunity review rather than a full LHEES refresh;
- Maintain a live register of strategic heat sources;
- Allow relevant heat network opportunity areas to be reviewed when major new developments come forward.

This would enable data centre waste heat to be considered alongside existing or emerging zones and affect whether zones should be refined or expanded.

The review should consider whether the scale and location of recoverable heat could support opportunities beyond the host authority, requiring cross-boundary collaboration with neighbouring authorities and relevant stakeholders.

Governance, Financing, Commercial Models

Background

Alongside technical viability being ascertained, it is important that Commercial and Governance structures are considered to ensure contractual viability. This section outlines the key considerations and drivers for both heat offtakers and Data Centre operators in the provision of waste heat.

Operational Drivers

The commercial position between data centre operators and heat network developers is shaped by the operational priorities of each party. Data centres operate within highly risk-sensitive environments where uptime, resilience and security are critical. Any third-party interface with cooling systems therefore introduces potential operational risk, and as a result data centre operators prefer commercial structures that minimise intervention within their operational boundary.

System Boundaries

This typically leads to a clear demarcation between the data centre cooling system and the external heat network infrastructure. Waste heat is commonly exported via a defined interface point at or near the site boundary, following hydraulic separation from the internal cooling system. This approach allows the data centre operator to maintain full control of its critical infrastructure while allowing heat network developers to design, own and operate the infrastructure required to capture and distribute heat externally.

Heat Supply Uncertainty

The commercial structure of heat supply arrangements is often influenced by the uncertain nature of waste heat availability. Data centres may be reluctant to guarantee volumes or temperatures of heat export, particularly as computing loads fluctuate and cooling technologies evolve. For this reason, early projects have often adopted a “heat when available” model, under which the heat network can use waste heat when it is available but does not rely on it as a guaranteed supply. This allows both parties to avoid exposure to operational or financial penalties if conditions change.

Affordable Heat Pricing

From the perspective of heat network developers and local authorities, the commercial objective is to ensure that the cost of waste heat supports and enables the overall commercial and financial viability of the heat network. Local authorities interviewed during this study emphasised that the price of waste heat must allow the network to deliver heat to consumers at competitive and socially acceptable tariffs. This is important in Scotland, where many heat network projects are expected to contribute to addressing fuel poverty and wider social outcomes.

Delivery and Ownership Models

Delivery Models Overview

Several potential ownership and delivery models exist for heat reuse projects involving data centres. The most appropriate model will depend on local policy objectives, the risk appetite of stakeholders, and the scale of the opportunity.

For further detail on available delivery and ownership models, refer to the Heat Networks Delivery Models publication, by Scottish Futures Trust (SFT):

<https://www.gov.scot/publications/heat-networks-delivery-models/pages/1/>.

ESCO Model

ESCO-led delivery models are widely considered the most straightforward and common approach. Under this model, a heat network developer or Energy Service Company (ESCO) is responsible for financing, designing, building, and operating the infrastructure required to capture and distribute waste heat. The data centre provides the interface point and may supply heat under an agreed commercial arrangement. This structure aligns well with the operational priorities of data centre operators, as it avoids requiring them to take on energy infrastructure risks outside their core business.

Public-Private Joint Venture

Public-private joint venture models are also emerging as a potential approach, particularly where local authorities wish to maintain strategic oversight of heat network infrastructure. Several Scottish local authorities indicated an interest in participating directly in delivery vehicles for heat networks. Such vehicles can allow local authorities to influence network expansion, integrate multiple heat sources, and ensure that projects deliver local social benefits.

Joint venture models can also help align infrastructure investment with wider place-based development strategies. For example, large housing developments or regeneration projects may benefit from being integrated into the governance structure of a heat network scheme.

Special Purpose Vehicle (SPV) Delivery Model

A third potential approach is the creation of a special purpose vehicle (SPV) to develop and operate the heat reuse infrastructure. In this model, a project company is established specifically for the heat reuse project, allowing liabilities and financing structures to be ring-fenced from the core businesses of the participating organisations. This model can enable blended financing from multiple sources, including private investors, public funding programmes, and infrastructure funds.

Comparison

Across these models, the balance between network-led and data centre-led delivery is an important consideration. The evidence from the stakeholder engagement suggests that network-led delivery is more likely to be viable in most cases. Data centres have limited appetite to assume responsibility for heat network infrastructure, as doing so would expose them to risks outside their core business model. In contrast, some heat

OFFICIAL

network developers are accustomed to managing heat supply risks and customer relationships.

However, there may be circumstances where data centre developers take a more active role. This is more likely where the cooling system has been specifically designed to enable heat export, or where the data centre forms part of a wider integrated energy development.

Tariff and Price Considerations

The pricing of waste heat varies between projects, but several consistent principles are emerging, shaped by the need to align the commercial positions of both heat network developers and data centre operators.

At a fundamental level, tariffs must reflect the incremental costs incurred by the data centre in enabling heat export. While waste heat is a by-product of core operations, it is not inherently “free” to supply. Delivering usable heat to a network typically requires additional capital investment in heat recovery systems, heat exchangers, pumping and control equipment, as well as ongoing operational and maintenance input. In some cases, it may also influence the design of the cooling strategy itself. Cost recovery is therefore a key requirement for data centre operators.

This is often where the perception of “free heat” arises. The term is frequently used because the underlying thermal energy itself is effectively a by-product that would otherwise be rejected to atmosphere. In addition, data centres may derive some operational or strategic benefit from heat export, for example through reduced cooling loads, improved energy efficiency metrics, or enhanced ESG performance. However, while these factors may reduce the cost of producing recoverable heat, they do not remove the costs associated with delivering usable heat to customers.

This distinction is critical. Even where the heat source itself is considered “free” at the point of generation, substantial infrastructure is still required to capture, upgrade, distribute, and manage that heat within a functioning network. In practice, customer heat tariffs must recover costs associated with:

- Heat recovery infrastructure within the data centre;
- Connection assets and energy centre equipment;
- Heat network infrastructure;
- Pumping energy and electricity consumption;
- Thermal storage;
- Operation, maintenance, and asset replacement;
- Network management arrangements.

As a result, there is effectively no scenario in which heat can be delivered to end users at zero cost, even if the raw heat source itself carries little or no commodity value. The phrase “free heat” should therefore be understood primarily as referring to the absence of a conventional fuel cost, rather than the absence of overall system cost.

OFFICIAL

In most real-world schemes, the largest component of heat pricing is not the heat source itself, but the capital and operational expenditure associated with the heat network and associated plant. This is particularly true where low-grade waste heat requires upgrading via large-scale heat pumps before it can be used within district heating systems. Electricity consumption associated with these heat pumps can become a major determinant of operating cost and tariff structure.

Existing international examples support this position. Projects where waste heat is described as “free” typically still involve significant public or private investment in network infrastructure and long-term commercial arrangements to recover those costs. In many cases, the attractiveness of waste heat lies not in eliminating customer heat charges, but in providing a lower-carbon and potentially lower-cost alternative to conventional gas-fired heat generation over the long term.

The position emerging in the market is therefore that waste heat should be regarded as a potentially low marginal cost heat source, rather than genuinely free for consumers. The economic value arises because there is no need to purchase primary fuel for heat generation in the conventional sense, however, substantial costs can remain in capturing, upgrading, transporting, and reliably supplying that heat to customers. Robust commercial structures must therefore ensure that these costs, alongside associated risks, and investment requirements, are appropriately recovered across the network lifecycle.

From a data centre developer’s perspective, the primary drivers remain operational resilience, uptime, and flexibility. Any heat export arrangement must not introduce risk to core operations or constrain future expansion. As such, participation in heat networks is viewed as an opportunity where it aligns with these priorities, rather than a primary objective. This reinforces the need for commercial structures that appropriately value the heat while recognising the data centre’s risk profile and investment requirements.

Alongside cost recovery, pricing structures must also support the delivery of affordable heat to end users. In the Scottish context in particular, where heat networks are often closely linked to public sector decarbonisation strategies and fuel poverty objectives, affordability is a central consideration.

Robust metering and monitoring arrangements are also essential to underpin any commercial agreement. Clear and transparent frameworks for measuring heat export, verifying performance, and reporting are critical to building trust between parties and ensuring that payments accurately reflect delivered value.

Two additional factors can materially improve commercial performance, and help bridge the gap between these objectives:

- Thermal storage, which allows heat to be captured when available and used when demand arises, improving utilisation of recovered heat and reducing reliance on backup systems;
- Temperature-based pricing, which recognises that higher temperature heat provides heat delivers greater system utility to the network (reducing or eliminating the need for further upgrading) and should therefore attract a higher tariff.

Overall, successful pricing arrangements tend to reflect a balance: acknowledging that waste heat has intrinsic value to the heat network, while also recognising the costs, benefits, and operational constraints faced by data centre operators.

Network Readiness

Ensuring that data centres can support future heat reuse requires a combination of supportive planning policy, technical design, and commercial structuring.

Planning and Policy Mechanisms

A key intervention is the requirement for a heat reuse plan as part of the planning process. This should set out expected heat capacity, temperature ranges, and potential local demand, ensuring that opportunities are considered early in the design phase.

Planning agreements can also secure provision of heat export infrastructure or future connection points. This helps avoid situations where developments are technically incompatible with heat networks once they emerge. A related gap identified through stakeholder engagement is the absence of a clear definition of a “green data centre.” Introducing a formal definition that includes waste heat recovery would help embed heat reuse as a standard expectation within planning frameworks.

Commercial Agreements and Heads of Terms

Commercial arrangements define the technical and operational interface between systems. Early development of Heads of Terms (HoT's) is particularly valuable in establishing key principles before detailed design and negotiation.

Typical HoT's cover:

- The parties involved and interface location;
- Responsibilities for connection infrastructure;
- Expected heat supply characteristics;
- Metering, reporting, and performance monitoring.

They also address operational matters such as maintenance, access, and outage procedures, alongside liability, insurance, and termination provisions.

Capacity Constraints

A recurring issue is limited capacity within local authorities to engage with complex commercial arrangements. For example, one stakeholder reported being unable to progress Heads of Terms discussions due to resource and capability constraints. This highlights the need for targeted capacity building support as the market develops.

Community Benefits

The potential for community benefit is a key driver for public sector involvement.

The most direct benefit is the potential provision of lower-cost, lower-carbon heat. By utilising waste heat, networks can reduce reliance on fossil fuels while offering more stable long-term pricing. This is particularly relevant in Scotland, where projects are often expected to contribute to reducing fuel poverty.

Public–private partnership models can support these outcomes by enabling strategic oversight of pricing and ensuring that revenues are reinvested locally. They also allow integration with wider place-based objectives.

OFFICIAL

Large public sector estates, such as those operated by the NHS, can play an important role as anchor loads, providing stable demand and improving network viability. However, stakeholders emphasised the need to manage long-term risks, particularly given uncertainty around future data centre technologies and heat outputs.

Beyond direct delivery models, there is growing interest in mechanisms to secure community benefit where heat reuse is delayed or not immediately viable. One option is the creation of a community benefit or heat infrastructure fund, supported by developer contributions.

A complementary approach is a form of offset mechanism, where developers make financial contributions in lieu of immediate heat reuse. These funds could be ring-fenced for future heat network development or wider decarbonisation activity.

Such mechanisms help address the current lack of consistent policy levers to require or incentivise heat reuse. Importantly, they should complement, rather than replace, direct heat utilisation, which remains the primary objective.

Conclusions

This report has examined the opportunity to reuse waste heat from data centres in Scotland, considered the scale of the potential resource, the technical and spatial factors that influence viability, and the policy, commercial and delivery conditions required to realise this opportunity. Drawing on stakeholder engagement, case studies, and analysis, this section sets out the key conclusions arising from the study.

Market maturity: The data centre (DC) heat reuse market in Scotland is at an early stage. Both the heat network sector and data centre pipeline are still developing, with no operational examples of large-scale heat reuse from data centres, and delivery models are still emerging. Small to mid-scale examples, such as the University of Edinburgh¹⁸, show that data centre heat reuse can work where heat demand is close by.

Growth trajectory: Data centres are expected to become a significant component of Scotland's infrastructure, driven by increasing digital demand and the emergence of AI Growth Zones. These, alongside areas of planned grid reinforcement, are likely to shape where future heat reuse opportunities arise. *(Key examples are the North Lanarkshire AI Growth Zone, and the "Central Belt" area of dense population/heat demand).*

Economic opportunity vs environmental pressure: Data centres offer significant investment potential (with the UK/ Scottish government's AI strategy for Scotland estimating an additional annual generation of £23 billion GDP by 2035), but this is increasingly balanced by scrutiny of their environmental impact (electricity demand, grid capacity, water use). There is a growing challenge to the scale and pace of DC deployment in Scotland. Heat reuse is a clear opportunity to help balance these competing factors.

Role in heat network development: Data centres can function as a stable low-grade heat source (~20–40°C), which can be upgraded via heat pumps for network use. At scale, the forecast pipeline 3-5GW IT load could support heat supply to hundreds of thousands of homes, depending on utilisation and network integration. Co-location of data centres and heat networks serving new or existing developments can reduce grid stress compared to alternative low carbon heating. This is demonstrated by potential opportunity case studies such as Ravenscraig.

Location and economic viability: Data centre waste heat can support the economic viability of a heat network. As a higher-grade heat source than many environmental sources, it can enable improved heat pump efficiency, lower electricity consumption, and potentially lower heat generation costs than alternative low-carbon counterfactual sources (such as ASHP). However, these benefits depend on connection distance, demand density, heat availability, capital cost, electricity prices, utilisation, and the commercial terms of the heat supply agreement. Heat networks are particularly sensitive to pipework length and infrastructure cost, with viability improving where linear heat density is high. Larger data centres may support longer-distance transmission due to the scale of recoverable heat available, while smaller or edge facilities are more likely to require close

¹⁸ Heat from servers to keep university students warm - <https://www.bbc.co.uk/news/articles/c72pr1z854ro>

OFFICIAL

proximity to anchor loads. Data centre heat reuse should therefore be assessed on a case-by-case basis.

Commercial and delivery models: Data centre heat reuse is most likely to be delivered through network-led or third-party delivery models, where heat network developers or ESCOs finance, own and operate the heat recovery and distribution infrastructure. Data centre operators typically have limited appetite to assume heat network risk, preferring clear system boundaries and offering “heat-when-available” arrangements that protect their operational resilience. As a result, successful projects rely on clear demarcation of ownership, standardised interfaces, and proportionate commercial terms.

Need for standardisation and capacity building: The absence of standard commercial frameworks, Heads of Terms, and pricing principles remains a barrier to wider deployment. In parallel, capacity constraints within local authorities limit the ability to progress complex commercial negotiations at pace. Targeted guidance, template agreements, and supported policy mechanisms would help reduce transaction costs, build confidence among data centre operators, and accelerate delivery of viable heat reuse projects.

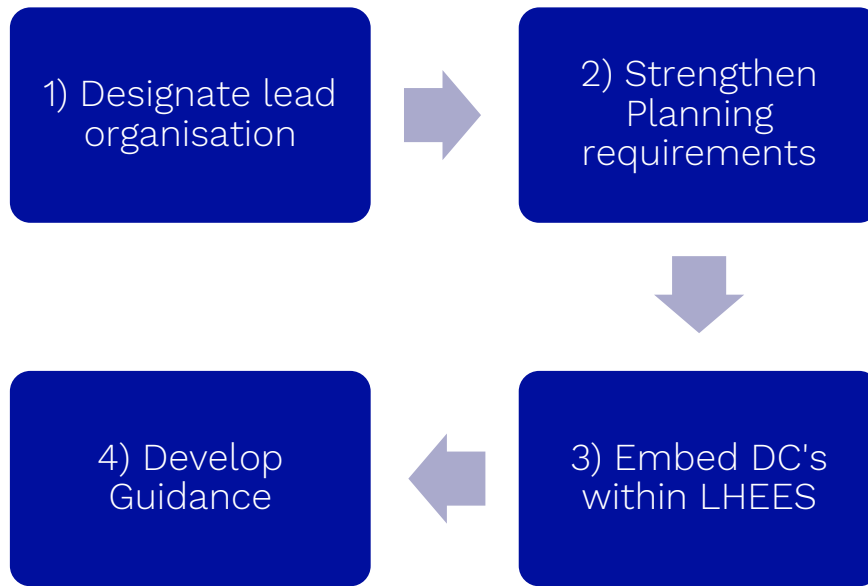
Policy context and opportunity for early integration: Realising data centre heat reuse at scale will benefit from early-stage planning and coordinated policy support. While National Planning Framework 4 supports “green data centres” as national developments and promotes heat networks through Policy 19, there is currently limited agreement of what constitutes a “green data centre”, and limited guidance on Heat and Power Plans or heat recovery expectations.

In addition, there can be misalignment between LHEES update cycles and the typically shorter development timelines of data centres. Given the strong location dependency of heat reuse, this highlights the importance of identifying and securing opportunities at the earliest stages of site selection and design.

Overall conclusion: While the opportunity for data centre heat reuse in Scotland is significant, realising this potential will require a shift from opportunistic development to a coordinated, policy-led, and spatially aligned approach.

OFFICIAL

Recommendations



1) Designate Lead Organisation for coordination

The designation of a lead organisation or named individual with a clear mandate to coordinate data centre heat reuse opportunities across Local Authorities data centre and heat network developers, the Scottish Government, and Department for Science, Innovation, and Technology within AI Growth Zones would enable the facilitation of heat reuse opportunities.

Sitting between public sector heat planning and private sector delivery, this role would provide technical oversight of heat and power strategies, align LHEES objectives with developer plans, and actively identify and support potential heat offtake opportunities at both planning and pre planning stages. A specific organisation/individual is not currently proposed. However, without coordination, it is likely that emerging Data Centres will continue to be developed without reasonable provision for heat reuse, making future heat reuse opportunities more difficult to realise.

This organisation would further benefit from the support of a community benefit/heat infrastructure fund or offset mechanism to be allocated where near-term heat reuse is not yet possible, ringfenced for future network delivery. This would encourage heat reuse and support the facilitation of the development of these projects.

2) Define and strengthen planning requirements for data centres

Expand planning policy (via National Planning Framework 4/5 or national planning guidance that builds upon NPF 4/5 – the format of which to be determined, but potentially a letter from the chief planner followed by a Planning Advice Note) to introduce clear, enforceable requirements for data centre developments. Whilst NPF4 supports “green data centres” as national developments, the absence of further guidance creates ambiguity. It is acknowledged that factors that may influence if a data centre is viewed as “Green” extend beyond heat reuse – for example may include source of electricity, waste, whole-life carbon, water usage etc. Actively assessing the opportunity for Heat reuse

OFFICIAL

should be one element of this (non-exhaustive) list of ‘tests’ a developer should meet in seeking planning consent.

Requirements should apply, as a minimum, to data centres classified as national and major developments, ensuring consistent treatment through the planning system. This should include mandatory Heat and Power Plans to be developed and explicit expectations for heat recovery where viable.

It is recommended that the contents of a Heat and Power plan are better established to ensure the priorities of Data Centre developers better align with Planning and Energy Policy:

- Provide a technical design outlining how provisions for heat offtake have been accounted for in future;
- Develop estimates for the level of heat supply & Engage with Local Authorities on how the DC heat offtake can interface with LHEES.

Local Authorities should consider implementing restrictions on cooling system types – seeking to reduce the usage of water-intensive evaporative cooling systems, in favour of closed loop systems. This would facilitate the usage of systems more compatible with heat reuse alongside aligning with wider environmental policy on reducing wastage of water.

Other referenceable guidance to note includes the OPDC industrial supplementary planning document¹⁹ and European guides such as the Dutch Collective Heat Act, which mandates the public ownership of heat networks, protects consumers against excessive heat prices and encourages owners of waste heat to make this heat available to networks where feasible.

¹⁹ <https://www.london.gov.uk/who-we-are/city-halls-partners/old-oak-and-park-royal-development-corporation-opdc/planning-policy/opdc-planning-policy/industrial-supplementary-planning-document-spd>

OFFICIAL

Heat and Power Plan – Suggested Content for Data Centres

As part of the planning application, developers should be required to submit a Heat and Power Plan that clearly sets out the potential for waste heat recovery and how this has been considered within the design. As a minimum, this should include:

1. Heat reuse potential and technical parameters

- Assessment of recoverable heat, including indicative capacity (MW), annual export (MWh), and temperature ranges
- Key operational parameters, including peak IT load, data centre typology (e.g. enterprise, colocation, hyperscale), expected cooling systems, and indicative PUE
- Estimated cooling demand and profile of the facility
- Development programme, including phasing and expected utilisation over time

2. Evidence of opportunity identification

- Demonstration of engagement with the Local Authority, heat network developers, and/or potential anchor loads during site selection and design
- Summary of options considered for heat reuse, including network connection and alternative local uses

3. Design integration and futureproofing

- Description of how heat recovery has been incorporated into the cooling system design, and how the design maximises heat reuse potential
- Provision of infrastructure required to enable heat export, including space allocation for heat exchange and/or heat pump plant
- Where no immediate offtaker is identified, evidence of futureproofing, including safeguarded space and routes for pipework and connection points

3) Embed data centres within LHEES and zoning processes

Amend Local Heat and Energy Efficiency Strategies (LHEES) guidance to embed the identification and assessment of data centres as potential waste heat sources within the LHEES process, particularly during data collection and identification of constraints, resources, and new developments. Local Authorities should proactively engage with data centre developers and operators early in the LHEES process to capture indicative spatial, technical, and operational data, including location, scale, cooling systems, and potential heat output. This should be supported through existing stakeholder engagement processes and pre-application discussions, recognising that detailed information may not be available at this stage.

Where available, information provided through developer-led Heat and Power Plans should subsequently be used to refine and validate the LHEES evidence base at planning stage, ensuring alignment between project-level assessments and wider heat network planning. Whilst LHEES remains a demand-led process, the presence of a large and

OFFICIAL

consistent heat source such as a data centre should be incorporated as a supporting factor in zone refinement and prioritisation and carried through into Heat Network Zone review & designation where appropriate.

As LHEES are refreshed periodically, new data centres may emerge after the evidence base has been prepared. A more flexible approach for future LHEES guidance is therefore recommended, subject to further consideration of the appropriate mechanism and implications for the LHEES process.

Major new heat sources would trigger a proportionate opportunity review rather than a full LHEES refresh. This could be supported by a live register of strategic heat sources, allowing relevant heat network opportunity areas to be reviewed, refined or expanded where appropriate, including across local authority boundaries where the scale and location of recoverable heat justify cross-boundary collaboration.

4) Develop technical and commercial guidance

Local Authorities or Central Government should develop guidance and standards to ensure data centre heat reuse opportunities are deliverable in practice. This should include technical standards for Data Centre operators to ensure developments are “connection-ready” at planning stage, such as minimum heat export parameters, provision of safeguarded space for heat exchange and heat pump plant, and defined expectations for futureproofing, including protected pipe routes and connection points.

In parallel, develop standardised commercial principles for operators and heat network developers to support consistent delivery, including standardised Heads of Terms for data centre operators covering key aspects of heat offtake arrangements (e.g. pricing and availability).

OFFICIAL

Appendix A – Stakeholder Roadmap

The following table summarises the key actions that local authorities, data centre developers/operators, and heat network developers/operators should consider across the development lifecycle. It is intended to show how responsibilities evolve from early strategic planning and site identification through to detailed design, delivery, and operation. The actions are not intended to be prescriptive but to function as prompts during the project lifecycle.

Stage	Actions for Local Authorities	Actions for Data Centre Developers/Operators	Actions for Heat Network Developers/Operators
Coordination and Strategic Planning	<p>Set clear expectations for green / low-carbon data centres</p> <p><i>Define local expectations for energy, water use, and heat reuse through LDP / SPG; clarify how NPF4 Policy 19 is applied locally.</i></p>	<p>Early pre-application engagement</p> <p><i>Engage planners and energy officers early; recognise heat reuse as a material planning consideration.</i></p>	<p>Share strategic heat network context</p> <p><i>Provide visibility of existing or planned heat networks, anchor loads, and expansion plans.</i></p>
	<p>Signal strategic opportunity areas</p> <p><i>Highlight priority locations via LHEES, Heat Network Zones, regeneration sites, and growth corridors.</i></p>		
	<p>Establish coordination route</p> <p><i>Identify single points of contact across planning and heat/network teams.</i></p>		
Site Identification	<p>Embed DCs in LHEES evidence gathering</p> <p><i>Treat data centres as potential waste-heat sources within LHEES opportunity identification.</i></p>	<p>Assess site suitability</p> <p><i>Review grid capacity, planning alignment, network proximity, and spatial constraints.</i></p>	<p>High-level heat network screening</p> <p><i>Assess distance to demand, linear heat density, and network configuration.</i></p>

OFFICIAL

Site Identification		<p>Define outline technical parameters</p> <p><i>Set out IT load range, cooling approach, indicative heat quantity, and temperature.</i></p>	<p>Confirm role of DC heat</p>
		<p>Prepare outline Heat & Power Plan</p> <p><i>Document potential heat reuse, options considered, and alignment with local plans.</i></p>	<p><i>Determine whether DC heat is primary, enabling, or supplementary.</i></p>
Site Development	<p>Support coordinated feasibility</p> <p><i>Use planning conditions, policy support, and funding guidance to enable feasibility work.</i></p>	<p>Design for heat recovery ("heat-ready")</p> <p><i>Safeguard space for PHEX / heat pumps; define connection points; protect pipe routes.</i></p>	<p>Develop detailed feasibility & business case</p> <p><i>Define technical configuration, energy centres, storage, and phasing.</i></p>
		<p>Confirm operational constraints</p> <p><i>Clarify availability assumptions, resilience requirements, and outage scenarios.</i></p>	<p>Define system boundaries</p> <p><i>Establish ownership, responsibilities, and interface points.</i></p> <p>Lead commercial structuring</p> <p><i>Develop Heads of Terms, availability model, and tariff principles.</i></p>
Construction and Operation	<p>Deliver planning commitments</p> <p><i>Monitor compliance and support phased expansion of networks and connections.</i></p>	<p>Construct heat interface infrastructure</p> <p><i>Build agreed substations, PHEX, and connections.</i></p>	<p>Deliver network and energy centre</p> <p><i>Construct heat network pipework, energy centre, and storage.</i></p>
		<p>Maintain operational resilience</p> <p><i>Ensure redundancy and uninterrupted cooling operation.</i></p>	<p>Operate and optimise system</p> <p><i>Run, maintain, monitor, and optimise network performance.</i></p> <p>Enable future growth</p> <p><i>Integrate additional heat sources and loads over time.</i></p>

Appendix B – Case Studies

A- Existing Examples

A-1 – Fjernvarme Fyn (Meta)

Project Context

Fjernvarme Fyn operates one of the largest district heating networks globally, supplying over 65,000 connections through approximately 120 km of transmission network (80–90°C) and 2,200 km of distribution pipework (70–75°C). As part of its heat supply portfolio, the network integrates waste heat from the Meta data centre in Odense, which has an IT capacity more than 500 MW. The data centre represents a significant, but not sole, source of heat to the network.

Technical Approach

The data centre uses an air-cooling system, with waste heat recovered from exhaust air via retrofitted water coils. This low-grade heat (~27°C) is upgraded using large-scale heat pumps to network supply temperatures (~70°C). The system has expanded over time, increasing from an initial capacity of 24 MW to approximately 45 MW of installed heat recovery capacity.

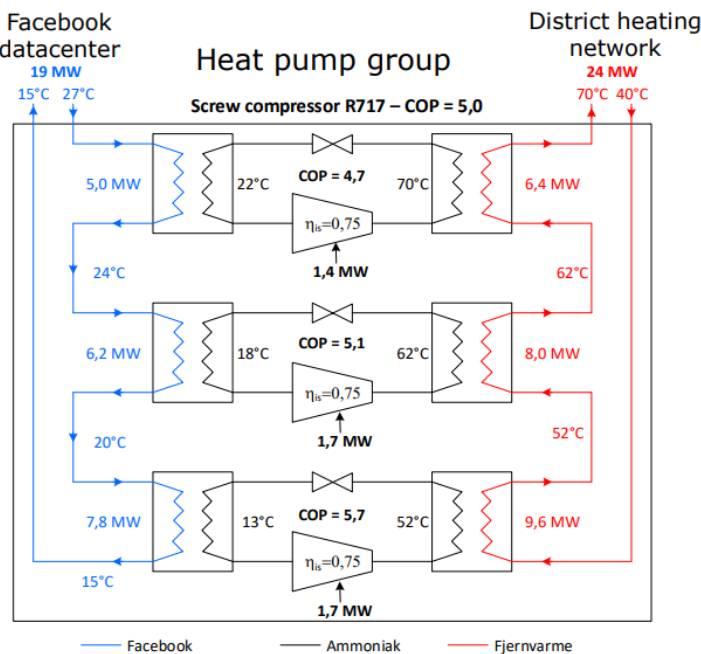


Figure 17. Schematic of Heat Pumps between Data Centre and DHN in Odense Commercial and Delivery Model

The heat network operator owns and funds the heat recovery infrastructure from the data centre offtake through to the heat pumps and associated equipment. The data centre operator does not charge for the waste heat. There is no guaranteed minimum heat

OFFICIAL

supply; in the event of reduced availability or shutdown, the network operator maintains resilience through alternative heat sources, including the potential use of ambient air using the same heat pumps and infrastructure.

Why the scheme works (key lessons)

- Integration into a large, established heat network with diversified supply sources reduces reliance on a single asset;
- A network-led investment model simplifies commercial arrangements and reduces complexity for the data centre operator;
- The system is designed for flexibility and resilience, recognising uncertainty in long-term data centre operation. To support operational stability the network has a large buffer vessel installed between the heat pump and data centre.

A-2 – Old Oak and Park Royal Energy Network (OPEN)

Project Context

The Old Oak and Park Royal Energy Network (OPEN), led by the Old Oak and Park Royal Development Corporation (OPDC) in West London, is a large-scale heat network planned for delivery between 2027 and 2040. The network is designed to serve a mix of new developments and existing buildings, including a major hospital, with an expected annual heat demand of over 90 GWh at full build-out. The area benefits from a cluster of existing and proposed data centres, forming the basis of the network's primary heat source strategy.

The scheme will use an ambient loop to connect multiple data centres, extracting low-grade waste heat (around 25°C) and transferring it to centralised energy centres for upgrading to network temperatures (75–80°C). Data centre sizes range from 2 MW to 56 MW IT load, with a combined contribution of around 15 MW of recoverable heat. Heat transfer substations use dual heat exchangers, defining the interface between data centre and network systems.

Technical Approach

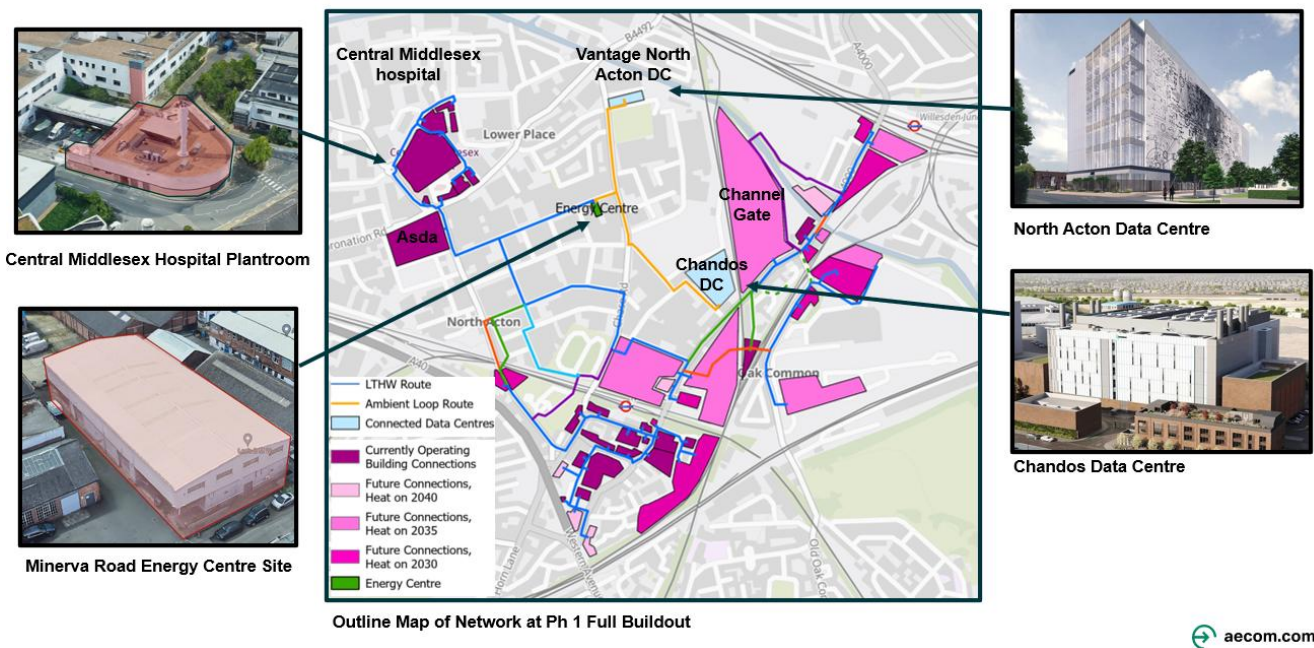


Figure 18. OPDC Heat network map

Planning and Policy Framework

OPDC has established a strong enabling framework through its Industrial Supplementary Planning Document (SPD)²⁰, which builds on the London Plan, OPDC Local Plan and wider Greater London Authority energy guidance. The SPD sets out clear expectations for data centre developments to support heat networks, including requirements to safeguard space and infrastructure for heat export. This includes provision for pipework routes, space for plate heat exchangers, and integration of heat recovery within cooling system design. Developers are also expected to actively explore opportunities to export heat to adjacent developments.

Whilst the SPD does not carry the same weight as formal policy, it provides detailed and practical guidance that can be applied at planning stage, helping to translate high-level policy into implementable design requirements.

Commercial and delivery model

Planning and policy mechanisms have played a key role in enabling delivery. Section 106 agreements for new data centres include provisions requiring “best endeavours” to connect to the network and engage in feasibility assessments. The network is structured around a centralised energy centre model, with clear delineation of ownership and responsibilities between parties.

OFFICIAL

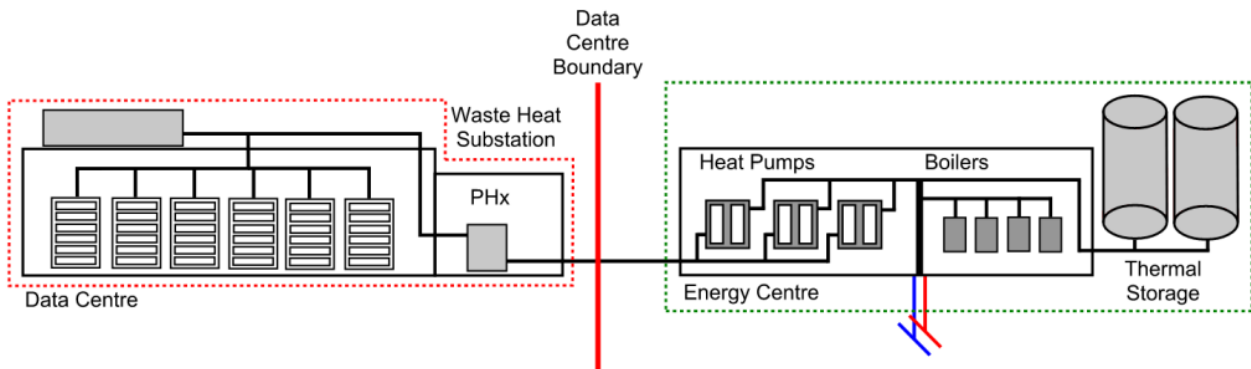


Figure 19. Delineation of responsibility between DC & HN.

Why the scheme works (key lessons)

- A clustered approach, connecting multiple data centres, reduces reliance on a single heat source;
- Strong local authority sponsorship and development from OPDC as a Mayoral development corporation, along with the Greater London Authority and DESNZ;
- Planning obligations (S106) have been used to secure future heat offtake opportunities along with supplementary planning guidance²¹ (SPD) used to translate policy into practical, design-level requirements;
- A significant pipeline of new residential development provides confidence in long-term heat demand, supporting scheme viability.

²¹ <https://www.london.gov.uk/who-we-are/city-halls-partners/old-oak-and-park-royal-development-corporation-opdc/planning-policy/opdc-planning-policy/industrial-supplementary-planning-document-spd>

Appendix C – Data Centre Classifications

Business model options are outlined below – given the variations in the usage in Data Centres, these are generalisations, and there may be exceptions to each case.

DC Type	DC Business Model	Implications for heat reuse/heat offtake
Enterprise	<p>Enterprise data centres are owned, operated, and used by the same organisation for their own purposes. They are typically single data centres.</p> <p>IT equipment. Capacities can range from 1 – 25MW.</p>	<p>The use profile of an enterprise data centre varies substantially as it depends on the owner organisation. Enterprise Data Centres are tied to this organisation’s IT strategy – introducing longer term risks of operational changes or obsolescence. Heat offtake viability should be considered on a case-by-case basis, but the dependency on the parent organisation make them less suitable in general (this risk is mitigated if the parent organisation is also the user).</p>
Hyperscale	<p>Typically, single major user/operator, but ownership and operation may sit with large technology companies and cloud service providers. Hyperscale DCs are being used for AI workloads.</p>	<p>Significant waste heat opportunity due to high utilisation, large scale, and ability to futureproof for heat offtake in design. Significant considerations needed though – particularly around grid impact. Some designed further from urban centres may invoke strategic heat mains.</p>
Colocation	<p>An operator owns these data centres, but racks/rooms are leased to customers, who supply their own IT equipment. Capacities can range from 1 – 50MW.</p>	<p>There are retail and wholesale opportunities – wholesale DCs tend to be larger than retail and can used AI/Cloud workloads. There are contractual considerations, as the limited-time agreements for space utilisation may result in periods of downtime/ no waste heat.</p>
Edge & Micro	<p>Edge and micro-DCs both are designed to process data closer to the source to reduce bandwidth. Capacities between 250kW – 5MW.</p>	<p>Due to their size, Edge Data Centres can be strategically located – particularly near heat offtakers. Edge Data Centres are increasingly using liquid cooling to support AI workloads, offering significant heat output from limited footprints. There is scope for heat offtakers to receive heat for free if the space is supplied to the Data Centre.</p>

Appendix D – Heat Network Zoning

Heat Network Zoning Framework

Heat Network Zones (HNZs) are a key component of Scotland’s strategy for decarbonising heat in buildings. The Heat Networks (Scotland) Act 2021 established a regulatory framework for the development and operation of heat networks and introduced provisions for identifying areas particularly suitable for heat network deployment.

Under the Act, local authorities are required to undertake a review of their areas to identify locations that may be **particularly** suitable for heat network development. Where appropriate, these areas may subsequently be designated as Heat Network Zones (HNZs).

The identification of potential zones is closely linked with the development of Local Heat and Energy Efficiency Strategies (LHEES), which provide the evidence base on heat demand, building characteristics, and decarbonisation pathways across each local authority area. Outputs from the LHEES process are therefore designed to be used to inform the assessment and designation of Heat Network Zones.

According to Figure 20 below, the HNZ guidance document is structured around the completion of a standard proforma. It consists of three parts: Part A – area review decision (supporting the requirements of sections 47(1), 47(3), 47(4) and 47(6) of the Act); Part B – designation consultation (supporting the requirements of sections 48(1) and 48(2) of the Act); and Part C – designation decision (fulfilling the requirements of sections 48(3) and 48(4) of the Act). Note that this process is under review, with work concluding and designation process commencing in 2026/2027.

Potential waste heat sources, including data centres, are typically first identified during the LHEES analysis as described in Section “Integration with Local Heat”. These sources are then considered in Part A of the HNZ proforma under Matter 48(1a) - Renewable and Waste Heat Opportunities, where the presence, scale, and location of renewable or waste heat resources that could supply a heat network are reviewed.

In Part B, the engagement and consultation process may involve stakeholders associated with these heat sources; relevant stakeholders can include data centre operators. Part C subsequently records the final decision on whether the area will be designated as a Heat Network Zone based on the evidence gathered through the previous stages.

OFFICIAL

Completed Prior to Heat Network Zone Designation Proforma

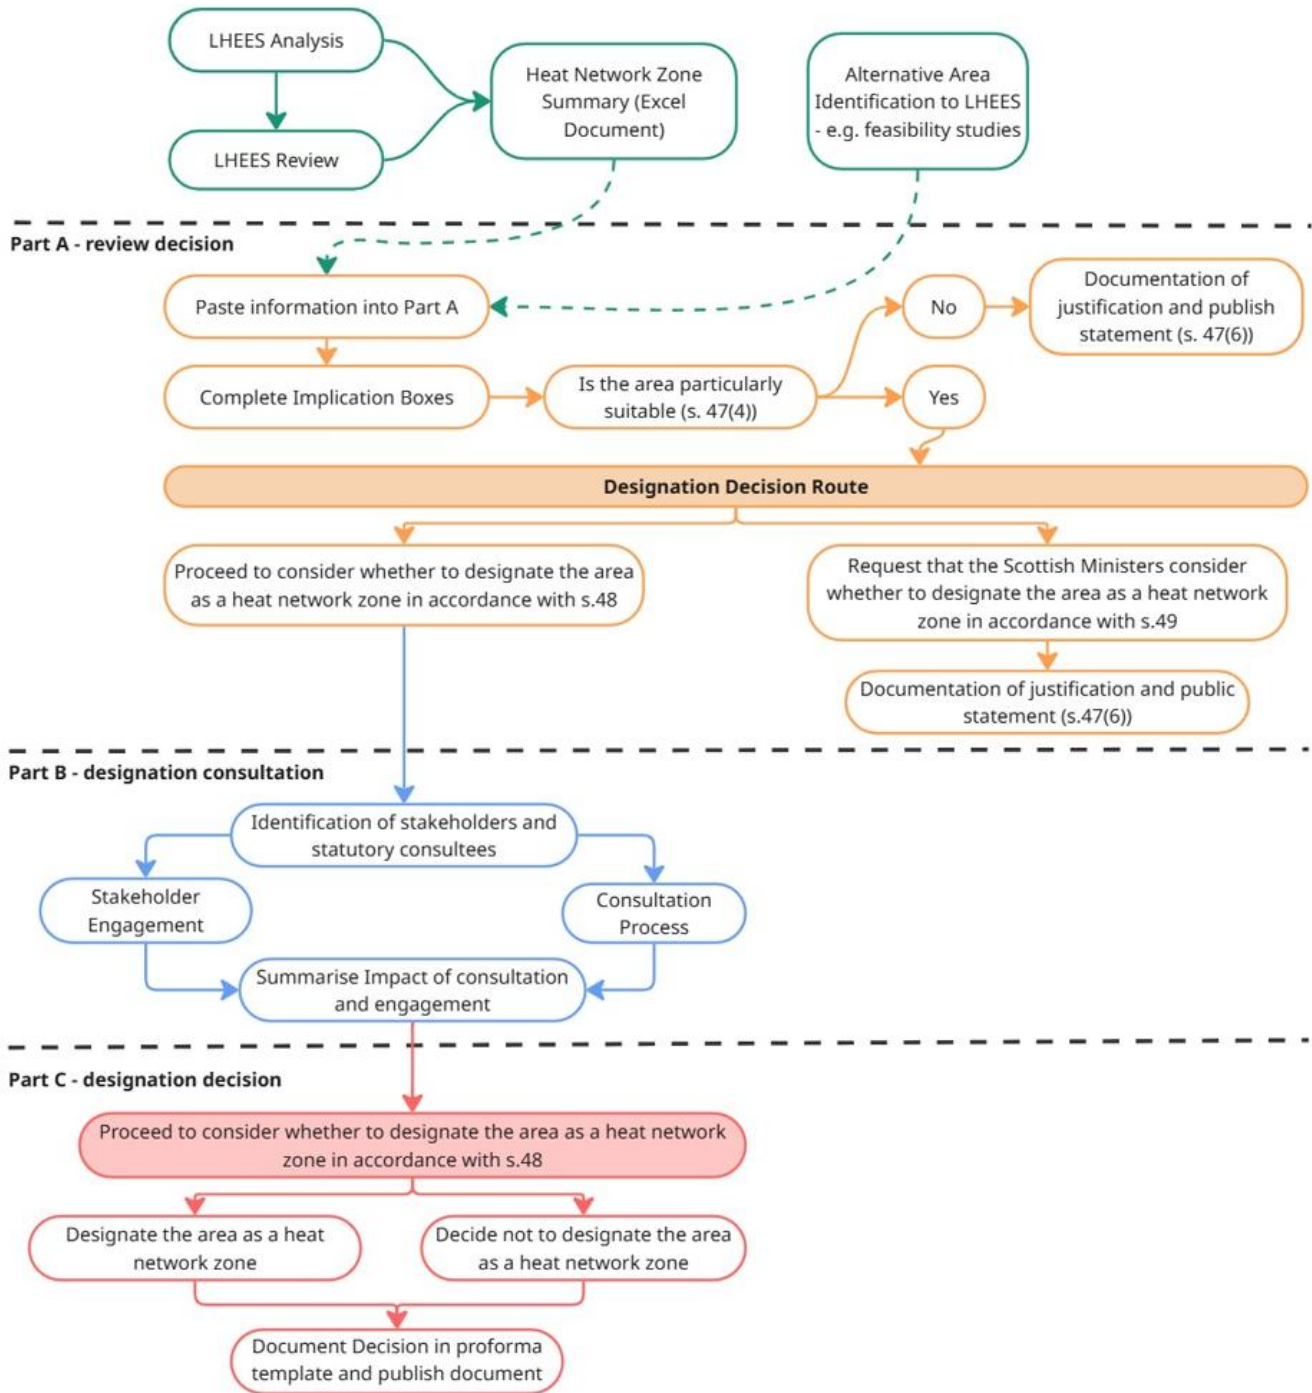


Figure 20. Summary of key stages to designate potential heat network areas as Heat Network Zones²².

²² Scottish Government, *Heat Network Zone (HNZ) guidance*, 30 May 2023.

Appendix E - Supply Chain Opportunities

The market for data centre heat offtake in Scotland is currently limited and immature. However, growing interest in both data centre development and heat decarbonisation could drive market maturation over the coming years.

This presents opportunities for suppliers and third parties to develop capabilities across the value chain. Key areas of opportunity within the heat offtake market are highlighted in Figure 21.

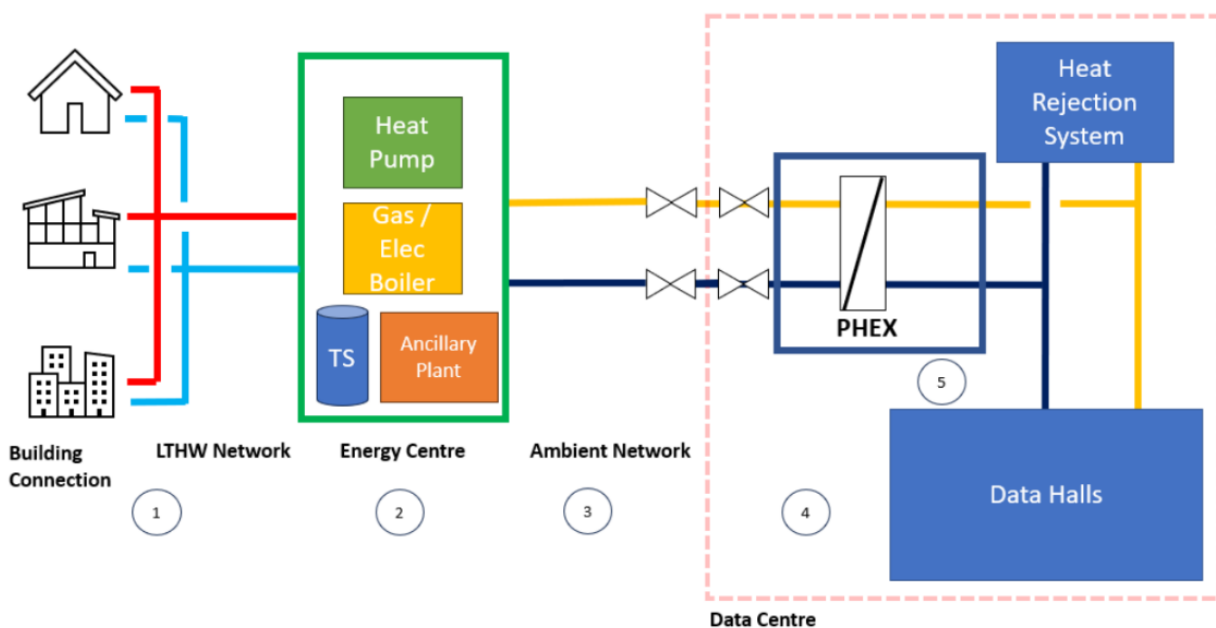


Figure 21. Schematic of DC/HN interface – Areas for Supply Chain Input

Areas of the Scottish Supply Chain that could develop to support the increase in equipment and jobs include:

- Heat Pump manufacturers;
- Heat Exchanger manufacturers;
- Boiler and heating system manufacturers;
- Manufacturers of auxiliary equipment – controls and BMS;
- Civil construction companies – heat network pipes and trenching;
- Construction and maintenance contractors/Energy companies;
- Building and land surveyors, Owners, and developers of future heat offtakers – NHS estates, Universities, schools, housing developers etc.



Scottish Enterprise

Atrium Court 50, Waterloo Street, Glasgow, G2 6HQ

www.scottish-enterprise.com