GPU SMITH
COMPUTE ENGINEERING

GPU SMITH / RESEARCH & PRACTICAL GUIDANCE

GPU Heat Reuse Worksheet: Demand, Heat Pumps, Backup

> Plan and build private AI compute with GPU Smith. We connect workload requirements with GPU hardware, networking, deployment and operating decisions so your infrastructure fits the work it must perform.

Published by GPU Smith · 2026-11-21 · gpu heat reuse · gpu waste heat recovery · data center heat reuse · heat pump sizing · liquid cooled gpu · district heating · waste heat calculator · fallback cooling

Original article: https://gpusmith.com/articles/en/gpu-heat-reuse-dispatch-worksheet



In this article

1. Executive Summary
2. Introduction and Background
3. Key Changes From Annual Calculators to an Interval Ledger
4. Interval Source and Sink Ledger
5. Heat-Pump Electricity and Temperature Lift
6. Fallback Cooling and Outage States
7. Data Analysis and Evidence
8. Implementation Considerations and Process Changes
9. Implications and Future Directions
10. Frequently Asked Questions (FAQs)
11. Conclusion

Executive Summary

GPU heat reuse is a time-matching problem, not a percentage problem. A private GPU facility that wants to export heat has to keep **three energy ledgers** apart: heat extracted from the IT cooling stream, electric energy supplied to the heat pump and auxiliaries, and useful heat accepted at the contractual delivery meter. This report, current as of November 2026, publishes a screening worksheet method, with a column schema and formulas, that builds those ledgers interval by interval. It is a reference for professional design review, not a forecast, and it contains no universal recovery factor, coefficient of performance (COP), heat price or payback period.

The public evidence explains why annual shortcuts mislead. The Commission's September 2026 report counts only **67 sites** with some reported heat reuse (data.consilium.europa.eu) out of 770 reporting data centres (data.consilium.europa.eu), and warns that its aggregate "may not be fully representative" (data.consilium.europa.eu). It names demand seasonality as a major barrier (data.consilium.europa.eu). Fortum's May 2026 update says its plants currently run on ambient-air heat pumps and electric boilers (www.fortum.com), with data-centre heat phased in later (www.fortum.com).

Capture is a measured, load-dependent quantity. In the one measured liquid-cooling study reviewed here, Lawrence Berkeley National Laboratory (LBNL) found **67 to 75 percent** capture, expressed as a share of cord power, at its best low-temperature results (datacenters.lbl.gov) and 57 and 59 percent at 50 and 100 percent power with 30 degree supply water (datacenters.lbl.gov). The authors state they did not explore net economic gain (datacenters.lbl.gov). Vendor liquid-share claims vary from 85 percent of a rack (lenovopress.lenovo.com) to **up to 98 percent** (www.supermicro.com) and must be replaced by site measurements.

The heat pump adds energy. With a verified heating COP above one, compressor electricity is W = q/(COP-1) for source heat q, so a delivery cap D at the heat pump outlet limits source heat to q <= D(COP-1)/COP. Published COP evidence is scoped: a European survey reports 3 to 4 (vbn.aau.dk), the Danish Energy Agency says 3 to 5 (ens.dk), and an operating Odense data-centre plant averaged 4.2 including pumps (danskfjernvarme.dk). None is a default.

A fictional 8,760-hour example shows the mechanics. Heat accepted at the meter was **93 percent** of extracted heat but only **63 percent** of customer demand, a monthly-bin screen overstated acceptance by **5.8 percent**, and a 72-hour customer outage forced 97.2 MWh of rejection at up to 1.555 MW. The IT cooling path must carry that duty alone. Buyer-supplied meter data, equipment curves and tariffs decide the build, stage, gather-data or decline question.

Introduction and Background

Search interest in "gpu heat reuse" and "gpu waste heat recovery" is usually served by two kinds of material: annual calculators that multiply installed power by a recovery fraction, and guides that explain district heating in general. Both are useful early screens. One public waste-heat calculator states that it does not include capital cost, financing, maintenance or backup boilers (agentcalc.com), and a specialist guide uses a planning COP of **45 percent of the Carnot limit** (dcblueprint.com). Neither is a dispatch ledger. The decision this report supports is narrower: whether an owner, a heat-network counterparty and a lender should fund a heat-export connection and heat-pump capacity, what service to contract, and what cooling capacity must remain available when the heat customer cannot take heat.

GPU Smith is an independent private AI infrastructure engineering consultancy that advises on power and cooling for enterprise GPU deployments. It is not a heat-pump manufacturer, heat-network operator or cloud provider, and nothing here is a client result. The worksheet is a transparent method. Its example data are invented and labelled as such.

Three reasons make GPU facilities a distinct case. First, modern racks mix liquid-cooled compute with air-cooled networking and storage. NVIDIA's DGX GB200 guide puts rack power at **approximately 120 kW** (docs.nvidia.com) and says networking and storage components are air cooled (docs.nvidia.com), so liquid capture is never the whole rack. Second, coolant temperature is a design variable: ASHRAE's liquid-cooling classes run from **W17 through W45, plus the W+ class** (www.ashrae.org), and a consortium white paper notes that W45 and above can enable reuse even in warm climates (www.opencompute.org). Third, GPU load is not constant. Training ramps, idle periods and maintenance change the source term hour by hour while the sink, a heat network, follows weather and building behavior.

The European policy backdrop is described only neutrally. **Delegated Regulation (EU) 2024/1364** applies to data centres with installed IT power demand of at least **500 kW** (eur-lex.europa.eu) and defines the Energy Reuse Factor as **ERF = EREUSE/EDC** (eur-lex.europa.eu). The Energy Efficiency Directive refers to using waste heat unless it is shown not technically or economically feasible (eur-lex.europa.eu). This report offers no compliance conclusion. Reporting boundaries, thresholds and exemptions belong to the governing text and an auditor, and the energy ledger below is an engineering tool, not a regulatory certification.

Key Changes From Annual Calculators to an Interval Ledger

Change 1: Three ledgers instead of one recovery number

The first change is bookkeeping. The extracted heat, the added electricity and the accepted heat are three different quantities with three different meters, and mixing them is the commonest way a heat-reuse case looks better than it is.

- **Extracted heat (q):** thermal energy removed from the IT cooling stream and made available at the heat pump source side. It comes from the liquid-loop measurement, not from nameplate power times installed count.
- **Electric input (W):** compressor, source and sink pumps and controls, metered where possible on its own circuit. It is purchased energy, not recovered GPU heat.
- **Accepted heat (Q_meter):** energy that passes the contractual delivery meter during the interval, after demand limits, equipment limits and network transformations.

Delivery metering has an established physical basis. The OIML heat-meter recommendation describes the measurement as the rate of change of enthalpy between flow and return (www.oiml.org), which is why flow and temperature difference at the billed meter, not the data-centre side, define what was sold. A Danish regulation overview likewise says consumer metering measures actual heat demand (c2e2.unepccc.org).

Change 2: Capture is measured, scoped and load dependent

The one measured liquid-cooling study reviewed here is historical. LBNL's report describes a retrofit of 38 servers with direct cooling in one rack (datacenters.lbl.gov). Its best results at low supply temperature ranged from 67 to 75 percent of cord power (datacenters.lbl.gov). At a 30 degree supply, capture was **57 percent** at half power and **59 percent** at full power (datacenters.lbl.gov), and it fell further at idle. The report states that captured heat was a strong function of return water temperature (datacenters.lbl.gov) and that capture becomes harder as supply or return temperature rises (datacenters.lbl.gov). It also says it did not explore net economic gain or efficiencies of reuse (datacenters.lbl.gov).

That is evidence about how capture varies, not a current GPU cold-plate benchmark. Vendor statements for current hardware are different in kind: Lenovo's GB300 NVL72 tour page states that **85 percent** of the rack is liquid cooled (lenovopress.lenovo.com), Supermicro claims up to 98 percent heat-capture coverage per rack (www.supermicro.com), and trade press describes only the "vast majority" of GB200 heat going to direct-to-chip liquid (www.theregister.com). These are attributed vendor or press statements. An older SuperMUC-NG report quotes 85 to 90 percent heat recovery with the remainder handled by room air conditioning (www.hpcwire.com). The worksheet treats every one of these as a hypothesis to be replaced by logged data.

The captured heat is computed from the loop. The consortium white paper gives the relation as heat equals volumetric flow rate times temperature difference times density times specific heat (www.opencompute.org), which is the same as Q_capture = m_dot x c_p x (T_return - T_supply). Fluid matters, because glycol additives decrease the specific heat capacity of water (www.opencompute.org). A loop that is declared as water but is actually a glycol mix will over-report capture if the water value is used.

Change 3: Demand caps what can be accepted

Annual recovered heat is not annual sold heat. The Commission's report links lower reuse in warmer climate zones (data.consilium.europa.eu) and describes seasonality as a major barrier even where heating demand is high (data.consilium.europa.eu). Without storage, accepted delivery in any interval cannot exceed the demand in that interval. With storage, the balance needs a state of charge, losses and charge and discharge limits. Typical Danish tank storage holds about **12 hours** of full-load production (c2e2.unepccc.org), and seasonal pit storage is a separate solar-oriented concept for storing summer heat for winter (52.iea-shc.org). Neither converts an annual surplus into a saleable one.

Change 4: Heat-pump electricity is an explicit line

A heat pump moves heat and adds compressor energy. The worksheet books that energy separately and caps the source heat that can be upgraded, as developed below. An annual calculator that omits pump and heat-pump electricity, as one public calculator states it does (agentcalc.com), cannot say who pays for the extra electricity.

Change 5: Fallback cooling is part of the energy balance

The final change is that every interval records the heat the IT cooling system had to reject. That residual sizes the fallback path and must never be silently shrunk by a heat offtake schedule.

Interval Source and Sink Ledger

Source ledger

Each row of the source ledger is one interval with a start timestamp (UTC), a duration, separately metered IT and non-IT site electricity, loop mass flow, supply and return temperatures, declared fluid and specific heat, captured heat, and a quality flag for gaps or estimates. Time zones and daylight-saving shifts are a frequent silent error, so the schema fixes UTC and carries a local-time column only for readability.

Sink and plant ledger

The sink ledger holds demanded delivered heat at the meter, required supply temperature, heat-pump capacity at the interval's actual source and sink temperatures, equipment availability, network capacity and loss basis, electricity price including demand charges, delivered-heat price and the contracting party for each. *Table 1* below lists the columns of the blank CSV schema and the unit check on each.

Column	Unit	Meaning and unit check
ts_utc, dt_h	ISO 8601, h	Interval start and duration. Energy (MWh) equals power (MW) times dt_h, applied exactly once.
e_it_mwh, e_site_other_mwh	MWh	Electricity metered separately for IT and for non-IT site loads.
mdot_kg_s, t_sup_c, t_ret_c, cp_kj_kgk	kg/s, C, C, kJ/kg/K	Loop flow, temperatures and declared fluid property. Q_capture in kW = mdot x cp x (t_ret - t_sup).

Column	Unit	Meaning and unit check
q_avail_mwh	MWh	Capture computed from the loop row times dt_h, with a quality flag.
d_meter_mwh	MWh	Customer demand or contractual call at the delivery meter, same interval.
hp_cap_mw, cop_h, avail_flag	MW, 1, 0/1	Heat-pump outlet capacity and heating COP at that interval's source and sink temperatures; availability state.
eta_net	1	Fraction of heat-pump outlet heat reaching the delivery meter (loss basis declared and owned).
w_hp_mwh, w_aux_mwh	MWh	Heat-pump and auxiliary electricity, separate circuits where possible.
q_hot_mwh, q_meter_mwh	MWh	Heat at heat-pump outlet and at delivery meter.
q_used_mwh, q_rej_mwh, unmet_mwh	MWh	Source heat actually upgraded, heat the IT cooling plant had to reject, and demand not served.

Table 1 is deliberately plain. The point of the interpretive step is the reconciliation: the sum of q_avail over all rows must equal q_used plus q_rej, and q_hot must equal q_used plus w_hp, within metering tolerance. A row set that does not close is a data problem, and it should be visible in the quality flag rather than smoothed. Missing intervals, partial seasons and monthly bins must stay visible: monthly averages are a lower-resolution screen and are not equivalent to hourly dispatch, a point the fictional example quantifies below.

The accepted-delivery cap

Let D be demand measured at the heat-pump outlet in the same interval. The idealized heating balance before external losses is Q_hot = q + W_HP and COP_h = Q_hot/W_HP. For a verified COP_h above one, W_HP = q/(COP_h - 1) and Q_hot = q x COP_h/(COP_h - 1). The source heat that can be upgraded without exceeding the demand is therefore **q <= D x (COP_h - 1)/COP_h**. Most of the delivered energy at a COP of 3.2 is not GPU heat: about 69 percent of Q_hot is source heat and 31 percent is compressor electricity (a fictional COP used only to show the arithmetic).

If the demand D is instead measured at a downstream meter, apply the transformation explicitly. With a fraction eta_net of heat-pump outlet heat reaching the meter, Q_hot <= min(D/eta_net, HP capacity), and q <= (D/eta_net) x (COP_h - 1)/COP_h. The cap must not be applied unchanged at a different meter. Heat losses in pipes are site specific: a Danish analysis puts pipe heat loss typically between **5 and 10 percent** (ens.dk), but the contract decides who owns that loss.

The unit-checked formulas used throughout are:

- **Capture:** Q_capture (kW) = m_dot (kg/s) x c_p (kJ/kg/K) x (T_return - T_supply) (K), then multiplied by dt_h for kWh.
- **Heat balance:** Q_hot = q + W_HP, and COP_h = Q_hot/W_HP.
- **Electricity for a given source heat:** W_HP = q/(COP_h - 1).
- **Delivery cap at the meter:** Q_meter = min(D, eta_net x min(HP capacity, q x COP_h/(COP_h - 1))).
- **Rejected heat:** q_rej = q_avail - q_used, never negative.
- **Seasonal COP:** sum(Q_hot)/sum(W_HP) over the declared boundary, not the mean of interval COPs.

The storage case needs its own balance: SOC(t+1) = SOC(t) + (charge x efficiency_in - discharge / efficiency_out) x dt_h - standing loss, bounded by capacity and charge and discharge rates, with initial and final state constraints declared. Storage defers delivery within a window. It does not create demand.

Heat-Pump Electricity and Temperature Lift

Source and sink temperatures

Heat pump work depends on how far the heat is lifted. A European survey of large district-heating heat pumps states the rule plainly: the smaller the temperature increase, the higher the COP (vbn.aau.dk). In that survey, units delivering below 70 degrees averaged a COP of **4.6** against 3.5 to 4 for the rest (vbn.aau.dk), and it reports a COP between 3 and 4 as an annual average practical value (vbn.aau.dk). The survey's COP excludes the electricity of auxiliary source pumps (vbn.aau.dk), so system level COP is lower. It counts **28** industrial waste-heat heat pumps in Europe with 129 MW total (vbn.aau.dk), and a text search found no data-centre sources in it, so it is context for lift and COP only.

GPU liquid return temperature is a design choice and an outcome. A modelling study of the Frontier supercomputer describes its waste heat as low grade at **30 to 38 degrees** (www.osti.gov) and simulates a high-temperature heat pump that delivers hot water at 75 to 85 degrees (www.osti.gov). A companion article in Scientific Data describes the heat pump route to an 85 degree sink (pmc.ncbi.nlm.nih.gov). These figures belong to one facility's coolant, and the worksheet reads them as an illustration of order of magnitude only.

COP: definition and evidence

Which COP is quoted matters. One manufacturer explainer defines COP as an instantaneous value measured under clearly defined conditions (www.mitsubishi-les.com), while a seasonal figure covers many operating points. The worksheet therefore keeps the seasonal COP as sum(Q_hot)/sum(W_HP) over the declared boundary and separately reports the arithmetic mean of interval COPs, and the two differ whenever load varies. Auxiliaries enter the denominator only where the declared boundary says so.

Theoretical limits help sanity checking. The Danish Energy Agency's technology catalogue prefers the Lorenz relation to Carnot for district heating because the sink temperature glides (ens.dk), and its worked example of heating water from 45 to 85 degrees with a cooling source gives a Lorenz COP of **7.2** (ens.dk), after which a real machine reaches some fraction of it. It also states the COP values in its chapter are practically attainable, not theoretical (ens.dk), and that compression heat pumps typically give **3 to 5 units** of heat per unit of electricity (ens.dk). Its waste-heat datasheet is scoped to a 10 MW source cooled from 27 to 15 degrees (ens.dk), which is a reminder that catalogue values carry a stated duty point.

Operating evidence for a data-centre source exists in Odense, where a 2022 Danish catalogue reports a measured average COP of **4.2** for 2021 and 2022 as a total calculation (danskfjernvarme.dk), and states that ammonia heat pumps work poorly where the return temperature is above 40 degrees (danskfjernvarme.dk). A Siemens and Grundfos white paper reports that in most heat reuse cases the COP falls between **3 and 6** (assets.new.siemens.com). The range of published figures is itself the argument against a default.

Why constant COP is only a hypothesis

Research on large heat pumps warns that assuming a constant Lorenz efficiency, exergy efficiency or COP gives errors when operation departs from design conditions (backend.orbit.dtu.dk). The Danish Energy Agency also notes that large heat pumps on the market are not built for very fast start-up, shutdown or load changes (ens.dk). Both points matter for GPU duty. An interval ledger can use a manufacturer curve of capacity and COP against source and sink temperature and part load, including minimum turndown and start-up constraints. A fixed COP scenario is allowed in the worksheet, but it is labelled hypothetical, and the summary must say so.

Fallback Cooling and Outage States

The heat commitment must not depend on disrupting the GPU service level. For each interval the ledger records q_rej, the heat that the IT cooling plant itself had to reject. When the heat customer, the network connection or the heat pump is unavailable, q_rej equals the whole extracted heat for that interval, not only the surplus.

Four states are tested in the worksheet, and the buyer should fix their definitions in the contract schedule.

- **Normal export:** heat pump runs, the customer accepts up to demand, and the fallback plant rejects only the unmatched remainder.
- **Customer outage or demand curtailment:** accepted delivery is zero, the fallback rejection path must carry the whole captured heat plus any air-side heat.
- **Heat pump unavailable (maintenance or trip):** the same, with compressor electricity falling to zero but the source heat now rejected.
- **Partial or degraded temperature:** the customer requires a higher supply temperature than the heat pump can reach, so the interval is derated and the remainder rejected.

Two sizing rules follow. The residual rejection capacity must be at least the maximum interval q_avail plus the air-side share not captured by liquid, and the transition from export to rejection must be engineered with valves and controls demonstrated to hold the supply temperature the IT hardware requires. Dynamic ride-through testing is a separate subject. This report only requires that the evidence exists: valve positions, bypass sizing, controller logic, and a recorded test of sudden heat-customer loss.

Published design guidance supports the same independence principle. An Open Compute Project reference-design paper states that downstream heat rejection is sized so operation continues regardless of heat host demand or availability (www.opencompute.org), orders the options so that, if the heat host has no demand, dry coolers are the next rejection option (www.opencompute.org), and says that on heat-reuse failure the default is a return to the standard rejection system (www.opencompute.org). A US Department of Energy best-practice guide likewise says data centers will have a redundant cooling system able to remove the heat if the heat host is unavailable (www.energy.gov). The EU Code of Conduct best-practice guidelines note that dry-cooler designs may need chiller or computer-room air-conditioner backup (e3p.jrc.ec.europa.eu). Dry coolers also have an ambient limit: effectiveness depends on the ambient dry bulb temperature (www.energy.gov), and an LBNL paper lists a dry-cooler approach temperature of 10 degrees Fahrenheit (datacenters.lbl.gov). The worksheet therefore evaluates fallback capacity at the hottest ambient hours, when the heat customer may also have the least demand, not at the annual average.

Operators of heat networks plan for the same issue on their side. Fortum's heating plants currently use ambient-air heat pumps and electric boilers (www.fortum.com), and one of its statements lists **200 megawatts** of electric boiler capacity (www.fortum.com), a reminder that network supply continuity is arranged by the network, not by assuming the server heat source. In Stockholm, the network operator describes cooling the data hall with the help of one or more heat pumps so that the excess heat can be delivered (www.stockholmexergi.se). The Siemens and Grundfos paper states that the normal commercial model has the local heat network operator set up the heat upgrade plant (assets.new.siemens.com), which moves the heat pump boundary and the question of who holds the backup obligation into the contract. A multi-kilometre connection can significantly affect total investment cost (assets.new.siemens.com), and in one cited Stockholm case the data centre operator co-funded a significant share of the heat upgrade plant (assets.new.siemens.com).

Data Analysis and Evidence

Scope-labelled evidence

Table 2 below collects the public evidence used here with its scope. The right-hand column states what the evidence cannot be used for. This is the denominator discipline that the Commission report itself practices: **70.1 percent** of the reported data were deemed reliable (data.consilium.europa.eu), and the report describes the roughly **1.8 percent** aggregate reuse figure as one that may not fully represent the landscape (data.consilium.europa.eu).

Evidence	Scope and date	Do not use it as
Commission report, September 2026	Partial first reporting round: 770 data centres reported (data.consilium.europa.eu), 67 with some reuse (data.consilium.europa.eu).	A site recovery efficiency, a global GPU benchmark or a census.
LBNL-6641E, March 2014	Retrofit of 38 servers in one rack (datacenters.lbl.gov); capture 67 to 75 percent at best (datacenters.lbl.gov), 57 and 59 percent at 30 degree supply (datacenters.lbl.gov).	A current cold-plate capture factor or a net economics study (datacenters.lbl.gov).
Fortum operating update, 6 May 2026	Plants currently on ambient-air heat pumps and electric boilers (www.fortum.com); data-centre heat phased in (www.fortum.com).	Proof that the promised server heat is already delivered.
Heat pump survey, 2017	COP 3 to 4 annual average practical value (vbn.aau.dk), auxiliary source pumps excluded (vbn.aau.dk).	A design COP or a data-centre measurement.
Danish Energy Agency catalogue, 2026 edition	Practically attainable COP; 3 to 5 typical (ens.dk); Lorenz basis (ens.dk).	An equipment curve at your duty point.
Odense data-centre plant, 2021 to 2022	Measured average COP 4.2 (danskfjernvarme.dk); return above 40 degrees works poorly (danskfjernvarme.dk).	A transferable COP.

Table 2 shows that each public number carries a boundary, a date and a stated exclusion. The practical consequence is that none of them can be copied into a model cell. They can bound a sanity check, for example by flagging a modelled seasonal COP far outside 3 to 5 for a water-sourced plant, but the decision rests on the buyer's own meters and the vendor's curves.

Fictional two-season example (Hypothetical Example)

The following example is invented. It is not a measurement, a client result or a project forecast. The generating script was deterministic and produced **8,760 hourly records**, 4,728 winter hours (days 1 to 120 and 289 to 365) and 4,032 summer hours. Parameters: a water loop with mass flow 40 kg/s and c_p 4.18 kJ/kg/K, a temperature difference that varies with IT load from 3 K up to about 9.3 K at the peak (40 kg/s x 4.18 kJ/kg/K x 9.3 K is about 1.555 MW), a heat pump outlet capacity of 3.0 MW, eta_net of 0.95 from heat-pump outlet to customer meter, and fictional COP placeholders of 3.2 in winter and 4.0 in summer. Two events are added: a 72-hour customer outage in winter and a 168-hour heat-pump maintenance week in summer. Source heat in the example is liquid-loop heat only. Air-side heat is not modelled and would add to the rejection duty. *Table 3* shows the reconciled totals.

Ledger line (MWh)	Winter	Summer	Year
Heat extracted from IT stream (q_avail)	6,226.2	5,272.4	**11,498.6**
Source heat upgraded (q_used)	6,129.0	1,754.1	7,883.0
Heat-pump electricity (W_HP)	2,785.9	584.7	**3,370.6**
Heat at heat-pump outlet (Q_hot)	8,914.8	2,338.7	11,253.6
Accepted at customer meter (Q_meter)	8,469.1	2,221.8	10,690.9
Rejected by IT cooling plant (q_rej)	97.2	3,518.4	3,615.6
Customer demand at meter (D)	14,592.0	2,318.4	16,910.4
Unmet demand	6,122.9	96.6	6,219.5
Seasonal COP, sum(Q_hot)/sum(W_HP)	3.20	4.00	**3.34**

Table 3 reconciles: q_avail equals q_used plus q_rej in every row to the rounding shown (totals are rounded from unrounded hourly records, so a sum of rounded cells can differ by 0.1 MWh, as for q_used in the year column), and Q_hot equals q_used plus W_HP. The seasonal split shows two different limits. In winter the system is source limited, since accepted heat covers only **58 percent** of demand and unmet demand is 6,122.9 MWh with almost no rejection. In summer it is demand limited, and 3,518.4 MWh of the 5,272.4 MWh extracted is rejected. The heat-pump capacity limit never binds in the example. The arithmetic mean of interval COPs is **3.56**, higher than the 3.34 annual figure from summed energies, which shows why the two must not be mixed.

Further results. Accepted heat at the meter was 93 percent of extracted heat but 63 percent of demand. A monthly-bin screen using 720-hour averages gave 11,308.2 MWh accepted against **10,690.9 MWh** from hourly dispatch, an overstatement of 5.8 percent. The 72-hour customer outage rejected **97.2 MWh** with a peak of **1.555 MW**, which is also the peak extracted heat in the dataset and therefore the minimum residual rejection duty for liquid heat alone. The 168-hour heat-pump maintenance week rejected **221.2 MWh** and left 96.6 MWh of demand unserved. Net added electricity depends on the counterfactual: with a fictional avoided-cooling credit of 0.04 MWh of electricity per MWh of heat upgraded, incremental electricity was 3,055.3 MWh. With placeholder prices of 50 euros per MWh of delivered heat and 90 euros per MWh of electricity, the arithmetic

gives **534,545 euros** of heat payments, 303,351 euros of heat-pump electricity and 28,379 euros of avoided cooling electricity, a net of **259,572 euros** before capital, operations and maintenance, standby cooling changes and taxes. The implied break-even heat price in that arithmetic is **25.7 euros per MWh**. These are placeholders to test the formulas, and the break-even price is not a market price.

The worksheet reports a fixed set of outputs so that two reviewers can compare cases:

- **Accepted heat at the meter (MWh):** the only heat that can earn delivered-heat payments.
- **Extracted IT heat (MWh):** the source ledger total, with the share of intervals flagged as estimated.
- **Heat-pump and auxiliary electricity (MWh):** metered or computed from curves, with the boundary stated.
- **Curtailed or rejected heat (MWh):** what the IT cooling plant had to reject, by state.
- **Unmet thermal demand (MWh):** demand the plant could not serve, split into source-limited and capacity-limited.
- **Residual rejection capacity (MW):** the maximum interval duty the fallback path must carry.
- **Usable export hours and cash-flow range:** hours above a declared minimum delivery, and a range under explicit assumptions.

Operating comparison formula

The operating comparison is: delivered-heat payments, plus measured avoided baseline cooling electricity cost, minus added heat-pump and auxiliary electricity cost, minus incremental operations and maintenance and any change in standby cooling cost. State the counterfactual, the meter ownership and the tariff basis. Avoid counting both heat revenue and the customer's avoided fuel bill as owner revenue, and keep capital and connection quotes as a separate sensitivity. Electricity price is likewise a stack, not a number. Eurostat describes non-household prices as covering production and supply, network costs and non-recoverable taxes (ec.europa.eu), excluding VAT and some other taxes that these consumers can usually recover (ec.europa.eu), and its H1 2026 EU figures for the 500 to 2,000 MWh band were 0.1466 euros per kWh excluding taxes and 0.1800 euros per kWh including non-recoverable taxes (ec.europa.eu). Those are survey averages and not a tariff for any site. Network charges can be peak or capacity based: Great Britain's system operator describes half-hourly demand charged on use at the three highest system peaks (neso.energy), and Ofgem cites a fixed charge set by grid connection size as one design (ofgem.gov.uk). A winter-peak compressor load can therefore cost more than its energy alone suggests, and the worksheet carries energy price and demand charge as separate columns. Heat price bases are jurisdiction specific. In Denmark, for example, Act no. 738 of 20 June 2025 repealed the separate price cap on surplus heat, and it entered into force on **1 July 2025** (retsinformation.dk). Elsewhere the contract may reference a network avoided-cost basis. UK guidance on negotiating recoverable waste heat was reported as not yet published in October 2026 (cms.law).

Implementation Considerations and Process Changes

Decision framing

Define four outcomes before running numbers:

- **Build now:** data are complete and the case holds across the demand and price ranges.
- **Stage capacity:** the case holds for a first heat-pump block but not for the full connection.
- **Obtain missing data:** capture, demand or tariff data are too thin to separate the scenarios.
- **Decline the connection:** no tested range pays for the added electricity and residual cooling.

The owner, the heat customer, the contract delivery meter and the party who owns heat pump electricity risk must be named. A staged option is often the correct answer when the hourly demand file is incomplete or the capture data come only from vendor sheets.

Evidence to gather first

- **Source evidence:** metered IT and system demand, loop flow, supply and return temperatures, fluid properties, and capture by workload. Gaps must carry quality flags.
- **Sink evidence:** hourly network demand and return temperature for at least a full year, including summer minimum, with the delivery temperature the network needs. Return temperature is decisive: the Odense catalogue says ammonia heat pumps work poorly above 40 degrees (danskfjernvarme.dk), and almost all Danish district heating systems have one or several heat storage tanks that shape when heat can be absorbed (ens.dk).
- **Equipment evidence:** heat-pump capacity and COP curves at actual source and sink temperatures and part load, minimum turndown, start-up behavior (ens.dk), availability and maintenance schedule.
- **Commercial evidence:** electricity price components including demand charges, delivered-heat price basis, outage and availability clauses, loss ownership and connection quotes.

Contract and meter evidence

A heat offtake schedule should state the delivery meter and its standard basis (flow and enthalpy difference (www.oiml.org)), minimum and maximum delivery, supply temperature band, availability and curtailment rules, notice periods for customer outage, who bears pipe losses, and what happens to the IT cooling duty during each state. For regulatory reporting the Delegated Regulation states that reused energy is measured at the boundary of the data centre where the energy is handed off (eur-lex.europa.eu), and that data centres measure heat used or reused outside the boundary (eur-lex.europa.eu). The passages read did not address heat pumps, so whether heat-pump electricity is counted in any reported ERF must be confirmed with the governing standard and the auditor.

Data that could reverse the decision

- A liquid return temperature well below the network need, forcing a large lift.
- Hourly demand showing summer minimum below the heat-pump turndown, leaving long rejection periods.
- Customer availability or curtailment terms that require the heat pump to ramp faster than its documented capability.
- Residual cooling quotes that make full independence from the heat customer more expensive than the heat income.
- Electricity demand charges that make compressor electricity at winter peak costly.

Implications and Future Directions

Public heat-reuse numbers will stay aggregates with caveats: the Commission's quality assessment deemed 70.1 percent of reported data reliable (data.consilium.europa.eu). For a private GPU owner the more useful development is internal: interval logs of loop flow and temperatures are cheap to collect, and the same dataset serves cooling commissioning, tenant energy accounting and heat export. A worksheet like this one turns those logs into a decision file that a heat-network counterparty can audit.

Hardware changes will move the source term. NVIDIA describes the Vera Rubin NVL72 MGX compute tray as 100 percent liquid cooled (blogs.nvidia.com), which raises the liquid share, but by LBNL's measurements on older hardware capture varies with temperature (datacenters.lbl.gov), so each platform needs its own measurement.

For advisers such as GPU Smith, the practical guidance is to refuse a single number. A recovery factor, a COP or a heat price can each be correct for one plant and wrong for the next. Readers should consult qualified thermal, electrical and legal professionals; this report was not reviewed by any such professional and is not a substitute for one.

Frequently Asked Questions (FAQs)

How is GPU waste heat recovery calculated?

Compute captured heat per interval from the loop (mass flow times specific heat times temperature difference) and cap it by the customer's demand at the same interval. Do not multiply installed GPU thermal design power by hours. Section "Interval Source and Sink Ledger" gives the schema and caps.

What temperature does GPU server waste heat have for a heat pump?

It depends on the coolant design. ASHRAE classes run W17 through W45, plus W+ (www.ashrae.org), and the Frontier example describes 30 to 38 degree waste heat (www.osti.gov). The measured return temperature of your loop is the input, and LBNL found capture depends on it (datacenters.lbl.gov).

How big should the heat pump be, and how much electricity will it use?

Size to the interval demand distribution and the source heat available, not to peak source heat. Electricity is W = Q_hot/COP_h from a verified curve at the actual lift. Published COPs of 3 to 5 (ens.dk) are references. Auxiliary pumps are additional (vbn.aau.dk).

Can liquid-cooled GPU heat go into district heating?

Yes, in principle, through a heat pump or direct exchange when temperatures match, but the network must accept it. Existing examples include Stockholm, where the operator describes cooling the data hall with heat pumps (www.stockholmexergi.se), and Odense, where the measured COP was 4.2 (danskfjernvarme.dk).

What fallback cooling is needed if the heat customer stops?

Enough to reject the entire IT heat duty, liquid and air, in every state in which the customer, the connection or the heat pump is unavailable, with tested controls.

Is there a data center waste heat reuse calculator that covers all this?

Public calculators are useful first cuts but state exclusions such as seasonal mismatch, pump electricity and capital (agentcalc.com). The interval worksheet here covers the dispatch and electricity boundary only, not capital.

Conclusion

Heat reuse from a GPU facility is decided in the hours, not in the annual total. Extracted heat, compressor electricity and accepted heat are separate ledgers. Demand caps acceptance, the heat pump adds energy and sets a limit on how much source heat can be upgraded, and the IT cooling plant must still be able to reject everything when the customer cannot take heat. The fictional example illustrates the structure: a high ratio of accepted to extracted heat can coexist with low coverage of demand, a coarse monthly screen can overstate acceptance, and a short outage defines the minimum rejection duty.

The public evidence is scoped and dated: LBNL's 2014 measurements, a partial Commission reporting round, one operator's phased start and several catalogues of COP values. None transfers to a specific installation. The next step for a buyer is to collect interval source data, obtain the network's hourly demand and temperature requirement, request equipment curves at the real duty points, and run the worksheet with ranges. Staging or declining the connection is a legitimate result.

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. https://data.consilium.europa.eu/doc/document/ST-13443-2026-INIT/en/pdf
2. https://www.fortum.com/media/2026/05/fortum-has-started-heat-production-two-large-data-centre-sites-finland
3. https://datacenters.lbl.gov/sites/default/files/direct_liquid_cooling.pdf
4. https://lenovopress.lenovo.com/3dtours/gb300nvl72
5. https://www.supermicro.com/en/pressreleases/supermicros-dlc-2-next-generation-direct-liquid-cooling-solutions-aims-reduce-data
6. https://vbn.aau.dk/ws/files/281265126/energies_10_00578_v2.pdf
7. https://ens.dk/media/8614/download
8. https://danskfjernvarme.dk/media/lcgnphch/inspirational-catalogue-for-large-heat-pumps.pdf
9. https://agentcalc.com/data-center-waste-heat-reuse-calculator
10. https://dcblueprint.com/insights/data-centre-waste-heat-reuse-europe
11. https://docs.nvidia.com/dgx/dgxgb200-user-guide/hardware.html
12. https://www.ashrae.org/File%20Library/Technical%20Resources/Bookstore/Emergence-and-Expansion-of-Liquid-Cooling-in-Mainstream-Data-Centers_WP.pdf

13. https://www.opencompute.org/documents/ocp-acf-reference-design-guidance-white-paper-r1-pdf
14. https://eur-lex.europa.eu/legal-content/EN/TXT/HTML?uri=OJ:L_202401364
15. https://eur-lex.europa.eu/legal-content/EN/TXT/HTML?uri=CELEX:32023L1791
16. https://www.oiml.org/en/files/pdf_r/r075-1-e02.pdf
17. https://c2e2.unepccc.org/wp-content/uploads/sites/3/2023/02/regulation-and-planning-of-district-heating-in-denmark.pdf
18. https://www.theregister.com/on-prem/2024/03/21/a-closer-look-at-nvidias-120kw-dgx-gb200-nvl72-rack-system/912087
19. https://www.hpcwire.com?p=47689
20. https://52.iea-shc.org/article?NewsID=265
21. https://ens.dk/media/6558/download
22. https://www.osti.gov/pages/servlets/purl/2372986
23. https://pmc.ncbi.nlm.nih.gov/articles/PMC11450168
24. https://www.mitsubishi-les.com/en-de/knowledge/what-are-the-cop-and-scop-values-for-heat-pumps-14912.html
25. https://assets.new.siemens.com/siemens/assets/api/uuid:ddbf25c8-95c5-4843-861b-6361c4289e4a/DataCenterWhitePaperCustomerReferenceGrundfos.pdf
26. https://backend.orbit.dtu.dk/ws/files/215617397/ECOS2019_energy_all_final.pdf
27. https://www.opencompute.org/documents/2025-03-18-ocp-heatreuse-wp-referencedesigns-v0-1-pdf
28. https://www.energy.gov/sites/default/files/2024-07/best-practice-guide-data-center-design.pdf
29. https://e3p.jrc.ec.europa.eu/sites/default/files/2024-04/JRC136986_2024_best_practice_guidelines.pdf
30. https://datacenters.lbl.gov/sites/default/files/Hot%20for%20Warm%20Water%20Cooling.pdf
31. https://www.stockholmexergi.se/en/heat-recovery
32. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Electricity_price_statistics
33. https://neso.energy/industry-information/charging/charges-using-great-britains-electricity-system
34. https://ofgem.gov.uk/blog/spreading-costs-networks-fairly
35. https://www.retsinformation.dk/eli/lta/2025/738/pdf
36. https://cms.law/en/gbr/legal-updates/waste-heat-and-heat-networks-opportunity-or-obligation
37. https://blogs.nvidia.com/blog/gigawatt-ai-factories-ocp-vera-rubin

THE PUBLISHER

About GPU Smith

GPU Smith provides independent engineering and research for private AI infrastructure. We help technical buyers and operators reason about GPU workload sizing, hardware selection, procurement, deployment and inference operations, from the compute system to the networking, power and cooling requirements around it.

Start with the workload

Useful infrastructure decisions begin with the models, data, throughput, latency and operating constraints a team actually has. GPU Smith helps connect those requirements with system design and deployment choices. The goal is a privately controlled AI environment whose capacity and operational demands are understood before a purchasing decision.

Hardware and supplier research

Our public reference library covers GPUs, complete systems and networking components. The server-vendor directory and manufacturing research help teams investigate suppliers, compare options and follow sources behind technical claims. We distinguish manufacturer specifications from measured benchmarks and advertised capabilities from tested configurations. Reference pages are research resources, not a live inventory listing or a binding equipment quote.

Deployment and operations

GPU Smith's engineering scope includes infrastructure integration, networking, power, cooling and the practical operation of inference workloads. Our published articles explain the tradeoffs behind deployment, procurement and ongoing operations so teams can ask better questions and document decisions.

Work with GPU Smith

Explore the hardware library, GPU references, networking references and vendor research. Contact GPU Smith with your workloads and deployment constraints to discuss a project and confirm current scope, pricing and availability.

Inclusion of a manufacturer or product in our research does not imply a partnership, certification or endorsement.

Website: https://gpusmith.com

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.