Waste Heat, Warm Cities

Turning Waste Heat into Community Benefit

Publication Date
January 2026

Discipline
COOLING

Estimated Reading Time
20 minutes

Geography
GLOBAL

Prepared By
Jeramy Utara
Volterra Advisors

Perspective
TECHNICAL

Sections

Executive Summary

Data centers are increasingly measured by what they require — megawatts, substations, water, land, transmission upgrades — because those requirements are what the public can see. But there is another physical reality that is just as constant, and usually far less discussed: data centers create a continuous thermal output that must go somewhere. In most regions it is treated as exhaust. In a growing number of places, it is being treated as infrastructure.

The most consistent output of compute isn't digital. It's heat.

This report examines how waste heat moves from an unavoidable byproduct to a usable product — and why that shift has happened unevenly across the world. In mature district-heating cities, the conversation is no longer theoretical: data centers can be integrated into municipal heating supply, and in some cases have become meaningful contributors to local heat systems. In regions without heat networks, the same idea is possible but rarely repeatable, because it depends on custom infrastructure, bespoke deals, and the willingness of utilities or municipalities to build a market for heat.

The central conclusion is intentionally balanced. Waste heat reuse is not a universal requirement that can be bolted onto every facility. But it is also not a novelty. Where the right conditions exist — nearby heat demand, a thermal network (or a credible plan to build one), and a commercial structure that treats heat like a utility product — reuse can turn an otherwise invisible externality into a tangible local benefit. That benefit matters in a moment when data center growth is becoming more visible, more debated, and more tightly coupled to public expectations.

Why This Matters Now

Waste heat has always existed. What's new is the way data center growth is colliding with the realities of energy systems and local permitting. In many markets, the constraints shaping new data center builds are no longer internal engineering challenges — they are external: power availability, grid upgrade timelines, community sentiment, and planning friction. As data centers scale, their "system footprint" becomes part of public discussion rather than an industry detail. The industry can't rely on efficiency metrics alone to explain its role; it increasingly needs outcomes that people can recognize beyond the fence line.

At the same time, heating has moved from background infrastructure to a central decarbonization battleground. Cities and utilities in many regions are being pushed — by fuel costs, climate targets, and electrification strategies — to rethink how heat is produced and distributed. That creates a new kind of demand: not just for cleaner electricity, but for cleaner heat. In that context, a large, steady thermal source located near dense demand becomes strategically interesting in a way it wasn't a decade ago.

Policy and disclosure are also adding momentum. The EU is moving toward standardized sustainability reporting for data centers and an EU-wide framework intended to make performance more transparent and comparable — including elements relevant to energy reuse. Germany's Energy Efficiency Act is an even clearer signal: summaries of the law describe explicit Energy Reuse Factor requirements for certain new data centers beginning operation from mid-2026 onward. Regardless of whether similar rules arrive elsewhere, these frameworks matter because measurement changes behavior.

What is reported becomes managed, and what is managed becomes designed for.

Finally, the technical direction of the industry is making heat reuse easier to contemplate. Higher-density workloads are accelerating adoption of cooling architectures that can support warmer, more controllable thermal loops — conditions that tend to improve the feasibility of exporting useful heat. In other words: the scale is rising, the scrutiny is rising, and the design space is improving — all at once.

Background & Context

To talk about "waste heat" usefully, it helps to separate two ideas that are often conflated: the existence of heat, and the usability of heat. Data centers do not optionally create heat; it is the unavoidable physical shadow of computation. What varies is whether that heat can cross the boundary into another system in a form that is valuable to someone else.

That boundary is the real point of complexity. Heat reuse is less about a clever component inside the data center and more about integration between different worlds: digital infrastructure, thermal networks, municipal planning, and utility economics. Crossing that boundary requires a place to send heat, a mechanism to measure it, a way to maintain reliability on both sides, and a long-lived agreement that defines who owns what and who takes which risks.

This is why heat reuse has never scaled evenly. In many regions, there is no large, ready-made market for imported heat because district heating is limited or absent. Where a heat network exists, integrating a new heat source can be a familiar utility task. Where a heat network doesn't exist, the "reuse project" is not just a connection — it is the creation of an entire local market for heat: pipes, permits, customers, governance, backup systems, and a financing story that can survive changes in development cycles and political leadership.

That's why Europe has become the proving ground for data center heat reuse. It contains a higher concentration of cities where thermal networks already exist, where utilities are accustomed to long-horizon heat planning, and where policy has increasingly prioritized decarbonizing heat alongside electricity. Europe didn't make heat reuse possible — it made it integratable.

Heat Fundamentals: How Data Centers Make Heat, Move It, and Lose It

A data center is often described as a digital factory, but thermally it behaves more like a quiet industrial plant: it runs steadily, and it continuously produces heat. That heat is not a side effect — it is the physical outcome of moving electricity through dense electronics. From the perspective of a community, this matters because it means the facility is not just a large electrical load; it is also a potential, predictable heat source operating year-round.

Inside the building, heat is created at the smallest scale — chips, memory, and power delivery — and then moves outward through layers of infrastructure designed to keep equipment within operating limits. In air-cooled rooms, heat rides on moving air and is transferred into chilled-water or condenser loops; and in liquid-cooled environments, it is captured more directly as warm liquid near the source. The details vary by design, but the goal is consistent: move heat away from IT continuously and safely.

That "safely" is the key constraint. Cooling systems are engineered around reliability first, which means every facility must always be able to reject heat even if something outside the fence line fails. Heat reuse, when it exists, is therefore typically engineered as a parallel path: heat can be exported when a customer and network are ready to accept it — but the data center still needs a complete, independent rejection pathway as a fallback. In practical terms, exporting heat is not replacing cooling; it is adding an additional destination for the same heat.

Where you tap the heat also matters. Capturing heat from air is possible, but it is usually less elegant — air is bulky, variable, and harder to move long distances without big equipment. Capturing heat from water loops is typically cleaner: water carries more heat in less volume, is easier to meter, and interfaces more naturally with district heating and heat pump plants. This is one reason heat reuse projects tend to look like "utility interconnects" rather than rooftop add-ons. Because a heat pump moves heat rather than creating it, the electricity you spend on "temperature lift" can yield several times that amount of delivered heat — which is why lift (temperature gap) and power pricing end up deciding the economics.

Why district heating changes everything: it creates a market for heat

District heating turns heat from a local nuisance into a local commodity. Instead of each building solving heating on its own, a district network aggregates demand and creates a shared "route to market" for heat sources that would otherwise be stranded.

That matters for data centers because waste heat is real everywhere — but it only becomes a community benefit when there is infrastructure that can buy it, upgrade it, and deliver it. A district heating system provides the pipes, the customers, the billing model, and the operational muscle to treat heat like a utility product rather than a one-off engineering experiment.

In places without district heating, reuse can still happen — but it usually looks like private loops, campuses, redevelopment zones, or a single anchor customer. District heating is different: it makes heat reuse scalable because it makes heat tradable.

All of this leads to the central hinge in heat reuse: if a data center's heat is going to become a community benefit, it must cross a boundary — from a reliability-driven cooling system into a utility-style delivery system — and it must do so in a form that a real customer can use. That brings us to the question that decides almost every project's feasibility: temperature.

Exhibit A: Two Paths for Data Center Heat: Default Rejection vs Reuse

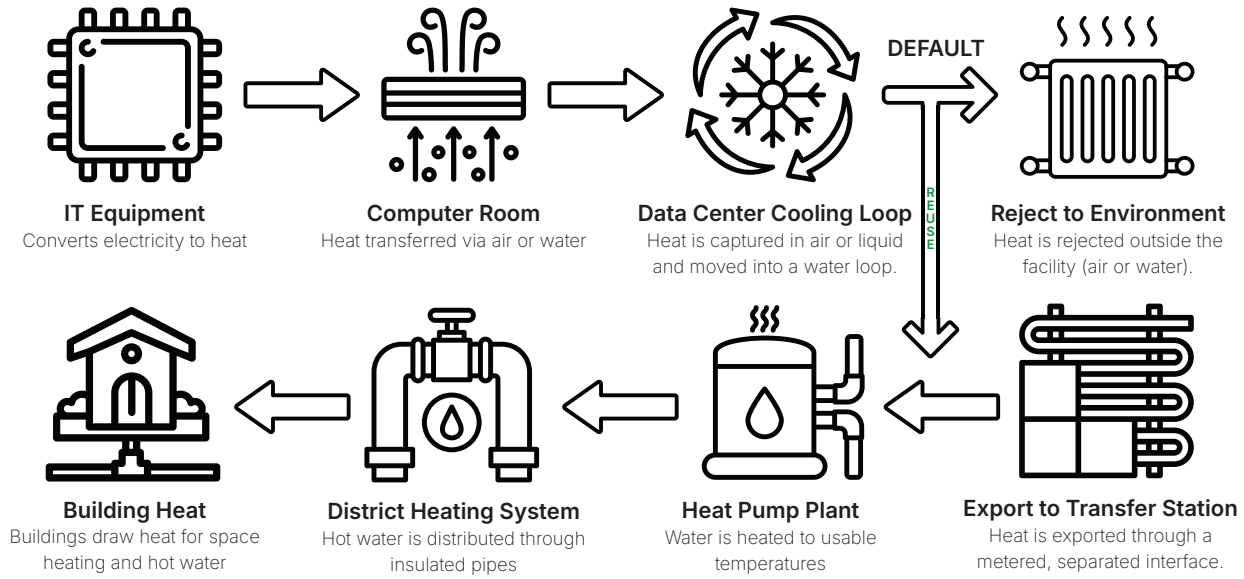

Data center cooling systems must reject heat safely by default. Heat reuse adds a metered transfer station that exports recovered heat to a heat pump and district heating network — while rejection remains the fallback path.

The Temperature Ladder: Why "Having Heat" ≠ "Having Usable Heat"

All waste heat is not created equal. A megawatt of heat leaving a data center at 28°C (82°F) is a very different product than a megawatt leaving at 70°C (158°F). The difference isn't volume — it's temperature, which is what determines whether heat can be delivered into a municipal network without heroic upgrades.

In many conventional facilities, the heat that's easiest to capture sits in a "low grade" band. Municipal planning guidance in Amsterdam, for example, describes typical data center residual heat in the range of roughly 25–35°C (77–95°F). That temperature is not worthless — it's warm water — but it usually cannot replace a boiler loop directly. Most district heating systems were built around supply temperatures high enough to serve space heating in winter and domestic hot water reliably, which is why low temperature heat almost always needs a boost.

That boost is the quiet workhorse of nearly every serious heat reuse project: the industrial heat pump. In Amsterdam, warm water from data centers is delivered to an "energy station," where

a heat pump lifts it from about 25°C (82°F) up to ~70°C (~158°F) before injection into a district heating network. This "temperature lift" step is what turns waste heat into a traded commodity — and it is also where economics, carbon intensity, and electricity prices start to matter.

Two things are making this easier than it used to be. First, district heating itself is evolving: modern low temperature networks are designed to run cooler than legacy systems, which reduces losses and makes lower grade sources easier to integrate. Second, data center cooling is slowly drifting toward warmer operating loops. Industry literature on liquid cooling notes the rise of "warm water" approaches (for example, ~40–45°C (~104–113°F) coolant ranges in some mainstream applications), which increases the quality of recoverable heat and reduces reliance on cold chilled water set points.

Exhibit B: The Temperature Ladder: Why Heat Pumps Enable Reuse

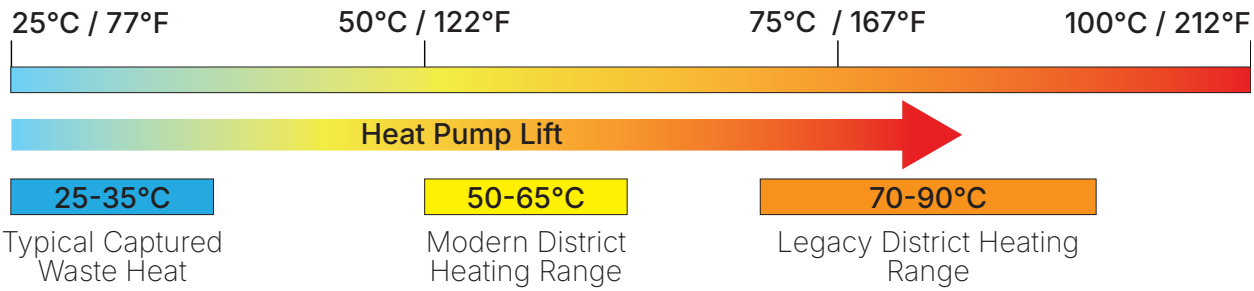

Data centers often deliver "low-grade" heat (about 25–35°C / 77–95°F) at the point of capture, while district heating networks typically require higher supply temperatures (roughly 50–65°C / 122–149°F for modern systems and 70–90°C / 158–194°F for legacy systems). Heat pumps provide the temperature lift that bridges the gap; real network temperatures vary by design and season.

A Quick Lens for the Case Studies: Why Some Projects Scale and Others Stall

Before we step into the European examples, it's worth naming the quiet truth behind nearly every successful heat-reuse project: the limiting factor is rarely "Can we capture heat?" The limiting factor is whether heat can be treated like a product — something with a buyer, a route to market, and an operating model that survives the real world.

Heat doesn't behave like electricity. It is bulky, local, and impatient. It doesn't travel far without infrastructure. It is hard to store at scale without purpose-built thermal storage. And demand is seasonal: the need for heat peaks when outdoor temperatures fall, while the data center produces a steady output year round. The projects that work solve these mismatches explicitly — with network design, storage, backup supply, and commercial structures that don't pretend seasonality and distance don't exist.

With those constraints in mind, Europe is less an exception than a preview: it's where heat networks, policy pressure, and utility planning already exist — and where data center heat can become part of the municipal system rather than a one-off experiment.

The Heat Reuse Test: When Waste Heat Becomes a Community Benefit

Proximity
Dense, durable heat demand nearby

Examples:
Housing • Hospitals • Campuses • Pools • Industry

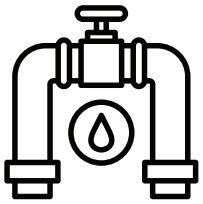

Infrastructure
A heat network exists (or will)

Examples:
District Heating • Utility Corridor • Redevelopment Loop

Commercial
Metering + terms that make heat bankable

Examples:
Price • Uptime • Fallback Heat • Responsibilities

Heat reuse becomes repeatable when three conditions align: (1) dense heat demand close enough to reach, (2) a district energy network (or a credible plan to build one), and (3) metering and contract terms that make heat a real product. Miss one cue and reuse tends to become a bespoke project; hit all three and it starts to behave like infrastructure.

Europe's Playbook: Proof That This Can Scale

Europe didn't become the leader in data center heat reuse because European data centers are different. It did so because many European cities already have district heating networks that can absorb heat, and because policy has increasingly pushed utilities and municipalities to decarbonize heating.

A good place to start is Denmark, where the Odense region demonstrates what "industrialized" heat reuse looks like. The Tietgenbyen Energy Center recovered 215,000 MWh of energy from Meta's hyperscale data center and distributed it through the Odense district heating network, delivering enough heat for more than 12,000 homes in 2023. The system relies on electrically driven ammonia heat pumps that lift the temperature to around 70–75°C (158–167°F), with total heat production described as about 45 MW. The story here is not just technical: Denmark abolished a tax on using surplus heat (for certified businesses) in January 2022, improving the business case — a reminder that policy can decide whether a project is feasible.

Finland offers the next step in ambition: not just a single project, but a broader strategy to make data center heat a meaningful component of a metro-scale heating mix. Microsoft's Espoo and Kirkkonummi data centers are expected to provide roughly 40% of the district heating needed in the area, with about 75% of the data centers' waste heat used annually for district heating (lower in summer when demand falls). Project details for Kirkkonummi show what

"making it utility-grade" requires: a plant with water-to-water and air-to-water heat pumps, electric boilers, and a very large heat accumulator tank (about 20,000 m^3 / 5,283,441 gallons), designed so district heating can be produced reliably even as the data center heat supply ramps and demand fluctuates. Finland also shows how quickly this can move from pilot to pattern. Google has described an off-site heat recovery project at its Hamina data center — a reminder that in markets with mature heat networks, reuse becomes an interconnect problem, not a science project.

Ireland is valuable precisely because it is not a classic district-heating market. Tallaght demonstrates a model for places trying to build district heating almost from scratch: start with a public-facing anchor network, operate it like a municipal utility, and expand. Tallaght is Ireland's first district heating scheme of its kind, operated by a not-for-profit utility (Heat Works) owned by the municipal council, and using waste heat from a nearby data center. Since April 2023 the system has delivered heat to a cluster of public buildings and has expanded to include residential connections; SDCC reports over 1,264 tons of CO_2 saved in 2024. What matters here is the institutional lesson: in new markets, someone has to be willing to be the "first utility." Tallaght shows a municipality can play that role.

Stockholm shows another scaling pathway: rather than treating every heat reuse project as a bespoke negotiation, create a standardized route into the network. Stockholm's Open District Heating (introduced in 2014) is a mechanism enabling industries — including data centers — to feed excess heat into the district heating system. By 2022, it reported 20 suppliers and enough recovered heat to warm about 30,000 modern apartments annually. This is heat reuse as a platform: if you make interconnection normal, supply diversity grows.

France provides the most "narrative-friendly" example – Equinix's PA10 facility near Paris was designed to recover and export heat to the Plaine Saulnier zone and the Olympic Aquatic Centre, with heat recovered at around 28°C (82°F) year round and then raised to roughly 65°C (149°F) using heat pumps before entering the heating network. The arrangement is described as delivering heat free of charge for 15 years. You don't need the Olympics for heat reuse to make sense, but you do need something like what the Olympics provided: a clear public anchor, a concentrated heat demand, and a visible "benefit story" that lands outside the data center industry.

Taken together, these European examples share a pattern: the data center isn't the hero — the heat network is. Odense, the Helsinki region, Tallaght, Stockholm, and Paris all work because there is a utility (or municipality) willing to buy heat as a product, upgrade it to usable temperatures, and deliver it through pipes that already exist or are being built. Europe looks like the center of gravity not because waste heat is uniquely abundant there, but because the receiving system is mature enough to make that heat tradable. Outside Europe, the physics is the same — but without standardized heat networks, reuse tends to cluster where ownership and geography simplify the problem: campuses, redevelopments, and one-off district loops built around a single anchor customer.

Beyond Europe: Rarer, but Not Absent

Outside Europe, heat reuse is real but less standardized, and therefore less common.

In the United States, Amazon has reported using excess heat from a Seattle data center (Westin Building Exchange) to help heat parts of its headquarters campus since 2019 — essentially a campus-scale district energy approach. This is a crucial proof point: you don't need a Scandinavian district heating backbone to reuse data center heat. But it also reveals the pattern: where district heating is uncommon, the easiest wins are often private or semi-private networks (campuses, redevelopments, multi-building portfolios) where a single owner can justify the pipe.

Canada's Québec region offers another emerging model in which energy distributors and utilities are treating waste heat as part of a broader decarbonization toolkit. QScale and Énergir Development have partnered to enable waste heat recovery, with early projects targeting deliveries and longer-term ambitions framed in the tens of megawatts and household-equivalent heating potential. The exact pace will vary by permitting, network build-out, and customer demand — but the direction is clear: once heating becomes a serious infrastructure priority, utilities start looking at data centers not only as loads, but as potential sources.

There is also a parallel path entirely: instead of moving heat to people, move compute to the places that already need heat. 'Computing heat' concepts treat servers as controllable thermal devices — placing workloads inside buildings so the heat is useful by default. Qarnot's distributed 'computing heaters' and building-integrated approaches like Cloud&Heat sit in this category. These models won't displace hyperscale, but they matter because they bypass the hardest variable in heat reuse: distance.

Why it Hasn't Spread Everywhere

If heat reuse were only a question of technical possibility, it would already be common. The reasons it isn't common are more mundane and more decisive.

First, in many regions there is simply no heat network to connect to. Without distribution infrastructure, there is no market. Building district heating is slow and politically complicated because it is street-level infrastructure: digging, permitting, easements, construction disruption, long-term governance. A data center can be built faster than a city can decide where a heat main should run.

Second, the heat often leaves "too cold." As the Amsterdam guide implies, typical residual heat temperatures are often in the 25–35°C (77–95°F) band. That makes data centers excellent heat pump sources, but it means heat reuse is inseparable from heat pump investment and the electricity needed to run those pumps.

Third, siting priorities historically ignored heat demand. Data centers go where power, fiber, land, tax structures, and permitting align. Heat reuse adds a new siting constraint: the facility has to be close enough to dense, long-lived heat demand to justify a pipe — and close enough to existing or planned networks to avoid becoming an isolated thermal island.

Without pipes, there is no market.

Fourth, utilities and municipalities need certainty. Heating customers expect heat. Data centers are built for reliability, but their heat output is tied to IT load, phasing, and operational constraints. That can be solved with storage, backup heat sources, and redundancy — as Finland's heat pump plant design illustrates — but it requires intentional design and additional capex.

Finally, until recently, the industry didn't measure heat reuse with the same seriousness it measures uptime and PUE. That is changing, and it matters.

The European Commission has now adopted a framework that requires reporting of defined data center KPIs into a European database and is establishing an EU-wide sustainability rating scheme. Germany goes even further: its Energy Efficiency Act describes explicit Energy Reuse Factor requirements for new data centers beginning operation from mid2026, with thresholds that rise over time and a structure intended to push operators toward systematic reuse rather than one-off pilots. When heat reuse becomes a disclosed KPI — or a compliance item — it stops being "optional ESG" and becomes part of design diligence.

ERF: The Metric That Turns Heat Reuse From a Story Into a Requirement

For years, the industry's dominant efficiency headline has been PUE — a measure of how much total facility energy is required for each electron delivered to IT. PUE mattered, and still matters. But it is blind to one critical question: whether the heat leaving the building is simply rejected, or actually reused.

Energy Reuse Factor (ERF) is built to answer that question. In The Green Grid's formulation, ERF is the fraction of a data center's total energy that is reused outside the data center boundary. A data center with no heat export has an ERF of 0. A facility that exports a meaningful share of its energy as useful heat moves toward 1 (though in practice real systems remain far below that).

ERF also connects to a related metric: Energy Reuse Effectiveness (ERE). The Green Grid defines ERE in a way that essentially "credits back" reused energy:

ERE = (Total Data Center Energy − Reused Energy) / IT Energy, which can be expressed as ERE = (1 − ERF) × PUE.

What makes ERF more than a new acronym is that it has started showing up in policy. Germany's Energy Efficiency Act (EnEfG) requires new data centers to achieve a planned ERF of at least 10%, rising to 15% for those starting from 1 July 2027 and 20% from 1 July 2028.

If a 20 MW facility exports 2 MW as metered heat, ERF ≈ 0.10. If PUE is 1.25, ERE ≈ (1−0.10)×1.25 = 1.125.

In parallel, the EU's reporting and rating framework is bringing heat into the measurement boundary. Data center operators above the threshold (installed IT power demand ≥ 500 kW) must report sustainability KPIs — including waste heat utilization and temperature set points — into the European database on a defined annual schedule.

The deeper implication is that heat reuse stops being a one-off "good news" feature and becomes a design constraint — the same way PUE once pushed the industry toward more efficient cooling and power distribution.

How it Could Scale Elsewhere: Making Heat Reuse Repeatable

Heat reuse scales when it becomes boring — and that requires a repeatable playbook that begins before site selection is final.

The first step is to treat heat the way we treat fiber and substations: map it early. Where are the concentrated, year-round heat demands that can act as anchor customers? Hospitals, campuses, pools, multifamily housing zones, industrial parks, municipal buildings — these are the loads that can justify pipe. A city that wants heat reuse at scale should be able to publish a "heat demand atlas" as readily as it publishes a zoning map.

The second step is to make proximity to heat sinks a siting factor. Europe's best projects are not miracles of engineering; they are outcomes of geometry. Odense works because there is a network ready to accept heat and the economics support it. If a region wants more of these projects, it has to stop assuming that any data center location is equally compatible with reuse.

PUE vs WUE vs ERF:
three metrics, three different stories

PUE is the industry's classic efficiency metric — it measures how much total facility energy is required per unit delivered to IT. But PUE is an internal metric: a data center can drive PUE down while exporting zero heat to the outside world.

WUE captures the water side of the equation — how much water a facility uses for cooling per unit of IT energy. WUE can improve dramatically with architectural choices (or climate), even if heat is still rejected as waste.

ERF is different. Energy Reuse Factor measures the fraction of a data center's energy that is reused outside the facility boundary — for example, exported as useful heat into a district heating network. That's why ERF is the metric that turns "nice story" into "measurable outcome," and why policy frameworks like Germany's have begun to set explicit ERF expectations for new builds.

PUE = internal efficiency

WUE = cooling water footprint

ERF = external heat benefit

The third step is design-for-export. Heat reuse is dramatically easier when a facility is planned with the right interfaces: space for heat exchangers and pumping, pipe corridors, metering, and control schemes that let the data center reject heat safely even when export is down. The Paris PA10 example is explicitly framed as a facility designed to recover and export heat rather than "retrofitted later."

The fourth step is to standardize the commercial model. At scale, heat reuse needs something like a Heat Purchase Agreement that is as familiar as a PPA: temperature and flow guarantees, outage terms, metering standards, expansion phasing, and a clear delineation of who owns the heat pump plant and who carries reliability risk. Germany's regulatory approach even reflects the idea that waste heat should be offered and accepted through structured processes, because the "deal mechanics" are central to whether reuse actually happens.

Finally, heat reuse needs its own reliability architecture. The best municipal projects treat heat like a utility product, supported by storage and backup sources — not as a fragile add-on. Kirkkonummi shows what "utility-grade" heat reuse requires: storage, backup, and resilience

For an illustrative list of U.S. cities with district steam/hot-water networks (and a rough compatibility tier), see Exhibit D (page 15).

Exhibit C: The Heat Reuse Blueprint: From Cooling Loops to Buildings

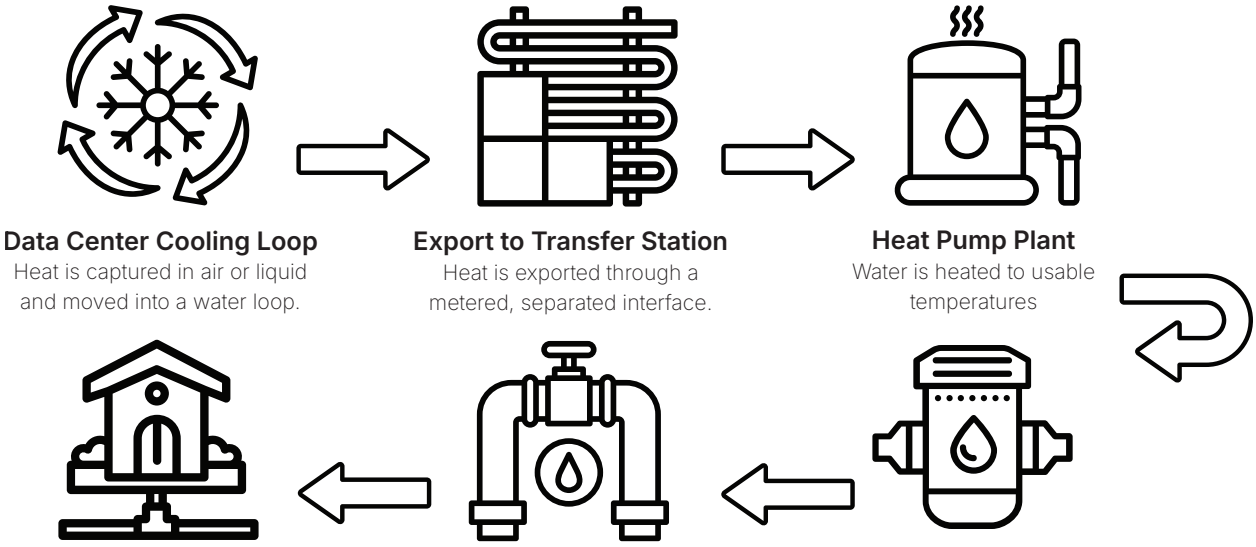

In a reuse system, recovered heat crosses a transfer station into district energy infrastructure, is upgraded by a heat pump plant, buffered with thermal storage (and/or backup heat), and delivered through district heating pipes to building loads.

What Changes the Equation Next: Liquid Cooling, Policy, and Visibility

Three forces are now converging in ways that could make heat reuse more common globally.

The first is cooling architecture. Liquid cooling is expanding with AI densities, and liquid loops often provide warmer, more stable coolant temperatures than air-based approaches. Direct liquid cooling implementations can accept inlet water temperatures in a warm range (commonly cited around 35–50°C (95–122°F) depending on design), which can improve free cooling potential and make exported heat more valuable by reducing the temperature lift needed. Research prototypes like IBM's Aquasar and projects like iDataCool demonstrate even higher-temperature coolant operation and explicit reuse pathways, suggesting that as liquid cooling mainstreams, "heat ready" designs become easier to justify.

The second is regulation and measurement. The EU has moved beyond encouragement into standardized reporting and sustainability rating infrastructure. Germany has moved into explicit reuse requirements for certain new builds. Even where similar laws don't appear immediately, the compliance tooling, reporting standards, and investor expectations created by these policies tend to travel.

The third is public visibility. Data centers are no longer anonymous. When power demand becomes local news, the industry needs benefits that are equally local and legible. Heat reuse is one of the few that meets that test. A pipe that warms a public building is easier to understand than a spreadsheet of efficiency metrics.

Closing: A Positive Story, with Real Constraints

The most honest conclusion is also the most useful: data center heat reuse is not a universal solution, but it is a repeatable one — in the right contexts. It works best where there is (or will be) thermal network infrastructure, where heat demand is dense and durable, where contracts are structured like utility agreements, and where facilities are designed for export rather than retrofitted in desperation.

> If the industry wants a credible positive narrative as data center sentiment hardens, waste heat is one of the few levers that can convert a core physical byproduct into a community benefit.

But it only becomes a dependable story when it is treated like a real product: engineered, metered, contracted, and planned as early as the substation.

Key Takeaways

- Waste heat is not a rounding error. In steady state, a data center's electrical input is overwhelmingly converted into heat; the question is whether that heat is rejected or exported.
- The main barrier is not "can we capture heat?" but "do we have somewhere to send it?" District heating networks (or equivalent heat sinks) are the unlock.
- Temperature is the gating variable. Typical residual heat (~25–35°C / 77–95°F) is valuable but usually needs heat pumps to become municipal-grade heat.
- Europe leads because infrastructure + policy align. EU reporting/rating initiatives and national policy (e.g., Germany EnEfG) are pushing standardized measurement and requirements.
- Large-scale success looks like "heat as baseload." Finland's Fortum/Microsoft model positions data center heat as a major share of a metro district heating system, supported by storage and redundancy logic.
- Municipal ownership can bootstrap heat networks. Tallaght demonstrates a city/county-led model can work even where district heating isn't historically common.
- Heat marketplaces may scale faster than one-off deals. Stockholm's Open District Heating shows how standardizing interconnects and procurement can multiply suppliers.
- "Designed for export" matters. Projects like Equinix PA10 show the advantage of building heat export into the facility from the start.
- Liquid cooling is a strategic accelerant. Warmer coolant loops improve heat quality and reduce lift requirements, making export easier and more consistent.
- Heat reuse is a credible positivity lever — but not universal. It scales where heat demand density, network infrastructure, and long-term contracts align.

Exhibit D - U.S. District Heating Systems and Waste-Heat Reuse Viability

Most U.S. district heating systems are legacy steam networks originating in the late 19th or early 20th century. While these systems demonstrate that urban heat distribution infrastructure exists, their design characteristics materially affect the feasibility of integrating low-grade waste heat from data centers.

The viability tiers shown here reflect structural compatibility rather than intent: hot-water networks and dense, anchor-served steam systems are meaningfully better positioned to absorb recovered heat at scale, while fragmented or purely steam-based systems face higher technical and commercial barriers.

Location	Operator	Medium	Origin	Viability
Akron, OH	Akron Energy Systems	Steam	1918	TIER 3
Baltimore, MD	Vicinity Energy	Steam	1891	TIER 2
Boston, MA	Vicinity Energy	Steam	1889	TIER 2
Cambridge, MA	Vicinity Energy	Steam	1910	TIER 2
Cleveland, OH	Cleveland Thermal	Steam	1893	TIER 3
Denver, CO	Xcel Energy	Steam	1880	TIER 3
Detroit, MI	Detroit Thermal	Steam	1936	TIER 3
Grand Rapids, MI	Vicinity Energy	Steam	1892	TIER 3
Harrisburg, PA	Cordia	Steam	1890	TIER 3
Hartford, CT	Hartford Steam Company	Steam	1854	TIER 3
Indianapolis, IN	Citizens Thermal	Steam	1887	TIER 3
Kansas City, MO	Vicinity Energy	Steam	1888	TIER 3
Lansing, MI	Lansing Board of Water & Light	Steam	1896	TIER 2
Milwaukee, WI	We Energies	Steam	1906	TIER 3
Minneapolis, MN	Cordia	Steam	1890	TIER 2
Morgantown, WV	Vicinity Energy	Steam	1905	TIER 2
New Orleans, LA	CenTrio	Steam	1890	TIER 3
New York City, NY	Con Edison	Steam	1882	TIER 2
Oklahoma City, OK	Vicinity Energy	Steam	1910	TIER 3

Location	Operator	Medium	Origin	Viability
Omaha, NE	Cordia	Steam	1884	TIER 3
Philadelphia, PA	Vicinity Energy	Steam	1887	TIER 2
Pittsburgh, PA	Cordia	Steam	1895	TIER 2
San Francisco, CA	Cordia	Steam	1893	TIER 2
St. Louis, MO	Ashley Energy	Steam	1892	TIER 3
St. Paul, MN	District Energy St. Paul	Hot Water	1983	TIER 1
Seattle, WA	CenTrio	Steam	1893	TIER 2
Trenton, NJ	Vicinity Energy	Steam + Hot Water	1914	TIER 2
Tulsa, OK	Vicinity Energy	Steam	1910	TIER 3
Tucson, AZ	Cordia	Steam	1928	TIER 3

Viability tiers:

TIER 1 Hot-water network or steam system already adapted for low-temperature integration

TIER 2 Dense legacy steam system with anchor loads and conversion potential

TIER 3 Steam system with structural constraints limiting scalable reuse

These tiers reflect structural compatibility, not project intent. Actual feasibility depends on network temperatures, proximity to demand, and commercial terms.

Glossary of Terms

COP (Coefficient of Performance)

COP describes heat pump efficiency as heat delivered ÷ electricity consumed. A higher COP means more usable heat produced per unit of electric input. In data center heat reuse, COP is one of the most important drivers of project economics because it determines how much additional electricity is required to lift low-grade waste heat to district heating temperatures.

District heating

District heating supplies hot water (or steam) from one or more heat production sites to a connected set of buildings for space heating and domestic hot water. The network's supply temperature varies by design generation; newer systems aim to reduce supply temperature to improve efficiency and enable integration of waste heat and renewables.

ERE (Energy Reuse Effectiveness)

ERE is an efficiency metric that adjusts a data center's energy profile by crediting energy that is reused outside the facility boundary. It is commonly expressed as: ERE = (Total Data Center Energy – Reused Energy) / IT Energy. In simplified form, it can be related to PUE and ERF as ERE = (1 − ERF) × PUE, showing how reuse can lower the "effective" facility overhead when exported energy displaces other heating sources.

ERF (Energy Reuse Factor)

ERF quantifies the fraction of a data center's total energy that is reused outside the facility boundary (for example, as exported heat into a district heating network). ERF helps distinguish "efficient" facilities from "useful" ones: a data center can have a strong PUE while exporting no heat, and a modest PUE while delivering substantial community heat benefit.

Heat pump (industrial)

In heat reuse, industrial heat pumps elevate low-grade data center heat (e.g., 25–35°C / 77–95°F) to temperatures compatible with district heating networks (commonly 50–90°C / 122–194°F depending on network). This is the enabling technology behind most municipal-scale reuse projects because it turns "too cool to use" heat into dispatchable, network-grade heat.

Heat Purchase Agreement (HPA)

An HPA is a long-term contract that defines how recovered heat is delivered, measured, priced, and guaranteed — including temperature/flow expectations, outage terms, maintenance boundaries, and expansion phasing. HPAs are the "utility contract" layer that makes heat reuse repeatable rather than bespoke.

Heat recovery

Heat recovery typically involves extracting thermal energy at a defined interface (heat exchangers, transfer stations) and delivering it—often via a dedicated loop—to a heat pump plant or directly to a low-temperature user. Recovery systems require metering, controls, and reliability design so that the data center can still reject heat safely if export is unavailable.

Heat rejection

Heat rejection is the final stage of cooling: after heat is removed from IT equipment, it must be transferred to an external sink (ambient air via dry coolers, atmosphere via cooling towers,

or water bodies in some configurations). Heat rejection is necessary for reliability, but it also represents an opportunity cost if that heat could have been reused.

Heat transfer station

A transfer station (sometimes called an energy station) typically contains heat exchangers, pumps, valves, and metering that allow heat export while keeping the data center loop and district heating loop hydraulically separated. It is central to contracting and reliability because it's where heat quantity and temperature are measured and guaranteed.

Low-grade heat

Many data center cooling architectures collect and reject heat at temperatures close to ambient comfort levels. This "low-grade" heat is plentiful and steady but lacks the temperature required by many legacy heating systems. It becomes most useful when paired with heat pumps or when connected to low-temperature district heating networks designed to operate at lower supply temperatures.

PUE (Power Usage Effectiveness)

PUE measures how much overhead energy (cooling, power conversion, lighting) a data center uses relative to IT load. It is useful for internal efficiency comparison but does not measure whether energy is reused outside the facility. Germany's EnEfG uses PUE thresholds as part of compliance expectations.

Temperature lift

Temperature lift is the difference between the heat source temperature and the required delivery temperature. In heat reuse projects, lift is what the heat pump must "add." Lower lift generally means better performance and economics — which is why both low-temperature district heating networks and warmer data center cooling loops tend to make reuse easier.

Thermal storage

Thermal storage increases reliability and economic performance by buffering the difference between steady data center heat production and fluctuating heat demand. Fortum's Kirkkonummi plant plan includes a large heat accumulator tank, illustrating how storage is integrated into real municipal-scale projects.

Waste heat

In data centers, nearly all consumed electricity ultimately becomes thermal energy that must be removed to maintain reliability. "Waste heat" refers to that thermal output when it is rejected to the environment rather than reused — and to the opportunity to capture it and deliver it to a useful application such as district heating or nearby buildings.

WUE (Water Usage Effectiveness)

WUE measures how much water a data center uses relative to IT energy consumption. It is commonly expressed as liters (or gallons) of water per kilowatt-hour of IT energy, and is used to compare how "water intensive" different cooling approaches and climates are.

Scope & Methodology

Volterra Reports are based on a combination of publicly available data, industry standards, operator and municipal disclosures, and first-hand market experience. Technical examples and metrics are illustrative, reflecting modern hyperscale and colocation facility designs — and the district energy systems they connect to — under typical climate, load, and operational assumptions. Where ranges are shown, they are intended to clarify relative trade offs rather than prescribe specific designs.

Analysis focuses on facility-level systems and local, on-site impacts unless otherwise noted, and does not attempt to model upstream grid effects, full life-cycle emissions, or site-specific engineering constraints.

Sources & References

European Commission, Directorate-General for Energy. "Commission adopts EU-wide scheme for rating sustainability of data centres."
Accessed January 2026.
https://energy.ec.europa.eu/news/commission-adopts-eu-wide-scheme-rating-sustainability-data-centres-2024-03-15_en

International Energy Agency (IEA). "Energy supply for AI — Energy and AI."
Accessed January 2026.
https://www.iea.org/reports/energy-and-ai/energy-supply-for-ai

Ramboll. "Largest Danish heat pump installation utilizing surplus heat from data center (Meta: surplus heat to district heating)."
Accessed January 2026.
https://www.ramboll.com/en-us/projects/energy/meta-surplus-heat-to-district-heating

Fortum. "Fortum and Microsoft's data centre project spearheads energy efficiency."
Accessed January 2026.
https://www.fortum.com/data-centres-helsinki-region

Fortum. "Construction of Fortum's heat pump plant has started at Microsoft's data centre site in Kirkkonummi."
Accessed January 2026.
https://www.fortum.com/data-centres-helsinki-region/construction-fortums-heat-pump-plant-has-started-microsofts-data-centre-site-kirkkonummi

South Dublin County Council (SDCC). "Development of the Tallaght District Heating Scheme."

Accessed January 2026.
https://www.sdcc.ie/en/climate-action/what-we-are-doing/energy-buildings/energy-buildings-actions/energy-efficiency-renewables/development-of-the-tallaght-district-heating-scheme.html

South Dublin County Council (SDCC). "Spotlight: Tallaght District Heating Scheme."
Accessed January 2026.
https://www.sdcc.ie/en/climate-action/spotlight-tallaght-district-heating-scheme/

Google. "First off-site heat recovery project: Hamina data center waste heat for district heating."
Accessed January 2026.
https://blog.google/inside-google/infrastructure/first-off-site-heat-recovery-project-hamina-finland/

EU Covenant of Mayors. "Stockholm, Sweden: Heat recovery from data centres."
Accessed January 2026.
https://eu-mayors.ec.europa.eu/en/news/stockholm-sweden-heat-recovery-data-centres

Data Center Knowledge. "Paris 2024: Excess Data Center Heat Used to Warm Olympic Swimming Pools."
Accessed January 2026.
https://www.datacenterknowledge.com/sustainability/paris-2024-excess-data-center-heat-used-to-warm-olympic-swimming-pools

Amazon. "Reusing data center heat to warm Amazon buildings in Seattle."
Accessed January 2026.
https://www.aboutamazon.com/news/sustainability/reusing-data-center-heat-to-warm-amazon-buildings

White & Case LLP. "Data center requirements under the new German Energy Efficiency Act."
Accessed January 2026.
https://www.whitecase.com/insight-alert/data-center-requirements-under-new-german-energy-efficiency-act

Bird & Bird. "Rechenzentren & Abwärme: Ein Überblick über die gesetzlichen Vorgaben zur Abwärmenutzung (Germany)."
Accessed January 2026.
https://www.twobirds.com/de/insights/2024/germany/rechenzentren-und-abwaerme-ein-ueberblick-ueber-die-gesetzlichen-vorgaben-zur-abwaermenutzung

City of Amsterdam / Open Research Amsterdam. "The Amsterdam Heat Guide."

Accessed January 2026.
https://openresearch.amsterdam/en/media/inline/2020/12/3/the_amsterdam_heat_guide.pdf

IBM Research. "Aquasar: A hot water cooled data center with direct energy reuse." Accessed January 2026.
https://research.ibm.com/publications/aquasar-a-hot-water-cooled-data-center-with-direct-energy-reuse

arXiv. "iDataCool: HPC with Hot-Water Cooling and Energy Reuse."
Accessed January 2026.
https://arxiv.org/abs/1309.4887

NEC. "NEC Direct Liquid Cooling Technologies (Whitepaper)."
Accessed January 2026.
https://www.nec.com/en/global/solutions/hpc/docs/NEC-DLC-Technology-Whitepaper.pdf

GlobeNewswire. "QScale and Énergir Development Join Forces... Through Waste Heat Recovery."
Accessed January 2026.
https://www.globenewswire.com/news-release/2023/03/16/2629024/0/en/QScale-and-Energir-Development-Join-Forces-to-Drive-Decarbonization-in-Quebec-Through-Waste-Heat-Recovery.html

Vattenfall. "Vattenfall launches pilot data centre project to offer sustainable compute capacity (with Cloud&Heat)."
Accessed January 2026.
https://group.vattenfall.com/press-and-media/pressreleases/2021/vattenfall-launches-pilot-data-centre-project-to-offer-sustainable-compute-capacity

D-Link (Case Study). "Qarnot is the first company in the world to make use of the heat produced by computer processors."
Accessed January 2026.
https://www.dlink.com/fi/fi/resource-centre/case-studies/2020_qarnot_computing_case_study_fr

The Green Grid. "ERE: A Metric for Measuring the Benefit of Reuse Energy From a Data Center (White Paper #29)."
Accessed January 2026.
https://datacenters.lbl.gov/sites/default/files/EREmetric_GreenGrid.pdf

European Commission. "Commission Delegated Regulation (EU) 2024/1364 of 14 March 2024

on the first phase of the establishment of a common Union rating scheme for data centres." Accessed January 2026. https://eur-lex.europa.eu/eli/reg_del/2024/1364/oj/eng

European Commission. "Commission Delegated Regulation (EU) 2024/1364." Accessed January 2026. https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=OJ%3AL_202401364

About Volterra

Volterra Advisors is a boutique advisory platform dedicated to accelerating growth across the digital infrastructure and data center sectors. Founded by Jeramy Utara, a recognized industry veteran with nearly two decades of experience, Volterra helps developers, investors, landowners, and operators navigate the decisions that matter most — from power and land strategy to leasing, platform growth, and market expansion.

Jeramy began his career in 2007 at DuPont Fabros Technology (later acquired by Digital Realty), supporting the development of some of the earliest large-scale hyperscale campuses in North America. He later joined CloudHQ as its first employee, where he helped grow the company from concept to one of the world's most successful privately held data center developers. Over more than a decade at CloudHQ, Jeramy played a leading role in leasing more than 1 gigawatt of capacity and driving over $15 billion in revenue, shaping relationships with the world's largest cloud, AI, and enterprise clients.

Through Volterra, Jeramy now partners with digital infrastructure platforms to define and execute strategies around site selection, powered land, leasing, marketing, and organizational growth. His advisory work emphasizes clarity, momentum, and execution — translating complex market forces into tangible outcomes for emerging and established platforms alike.

How We Can Help

Market Intelligence
We analyze power, permitting, and competitive dynamics to identify where the next generation of data center growth will land.

Site Identification & Readiness
We help landowners and developers qualify, position, and advance properties into premier data center ready sites.

Utility & Power Engagement
We maintain direct relationships with utilities nationwide to confirm capacity, align on timing, and support interconnection strategy.

Marketing & Positioning
We produce investor-grade materials — from decks to data sheets — that clearly communicate technical strengths and value.

Sales & Buyer Engagement
We connect landowners and developers directly with hyperscalers, operators, and investors ready to transact.

Partner & Consultant Network
We bring trusted engineering, permitting, and development partners to every opportunity — ensuring readiness, speed, and credibility.

Capital & Transaction Support
We support capital formation, deal structuring, and JV alignment — keeping incentives tied to successful outcomes.

Strategic Advisory
We serve as an extension of your leadership team — providing guidance, relationships, and insight from concept to close.

About Volterra Reports

Volterra Reports examine the systems, markets, and decisions shaping digital infrastructure — connecting power, land, cooling, policy, and technology into a coherent view of how the sector operates and evolves.

Each publication combines technical understanding with strategic perspective to explain how legacy decisions and emerging trends influence today's investments and tomorrow's platforms.

Average Read Time
15 minutes

Release Frequency
Monthly

Audience
Digital Infrastructure

Distribution Format
PDF + web

Example Reports

Water, Watts, & Workloads DEC 2025

Behind every megawatt of compute sits a cooling system that shapes both the energy profile and water footprint of a data center. Watts, Water, & Workloads follows the evolution of cooling from traditional air systems to modern evaporative and liquid-cooled architectures — revealing how AI-driven densities are rewriting the rules of heat rejection. With communities and utilities paying closer attention to water use, and operators pushing for efficiency at scale, cooling has become a defining design decision. The result is a landscape where water is emerging as a second 'fuel' for compute, and where liquid cooling is enabling dry, high-temperature pathways once considered impractical.

From the Curb to the Computer Room FEB 2026

Data centers are often discussed in terms of megawatts, efficiency metrics, and uptime statistics, but those abstractions obscure the physical reality that ultimately determines reliability. *From the Curb to the Computer Room* walks through the spaces and systems that make a modern data center function — from perimeter security and logistics to power distribution, cooling infrastructure, and the data hall itself. Drawing on first-hand operational experience, the report follows the physical path of people, equipment, electricity, and heat through a facility, explaining how design decisions made early shape safety, operability, and resilience over decades of operation. Rather than prescribing best practices, it offers a practical framework for understanding how reliability is built room by room — and why physical design has become increasingly visible as data centers grow larger, denser, and more scrutinized by communities, utilities, and regulators.

Volterra Reports are written for clarity, momentum, and practical application — bridging technical depth with real-world decisions.