DESIGN FOR MORE EFFICIENT DATA CENTERS

How Automation and Analytics throughout a Data Center Lifecycle Can Help Reduce Energy Use and Environmental Impact

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Increased demand for information technology is driving a rapid expansion in global data center capacity. It has been estimated that data centers could account for up to 10% of global electricity demand growth by 2030, so sustainable design and operation are becoming increasingly urgent priorities for data center operators.

There are two main strategies for managing the energy use of data centers through more sustainable lifecycle design:

Consider thermal management lifecycle in the design and strategic planning phase and understand options for waste heat recovery and reuse. A narrow focus on minimizing the power usage efficiency (PUE) can miss opportunities for waste heat use that can defray operating costs and substantially reduce the overall carbon footprint of the data center and surrounding community. Our lifecycle analysis shows that there are situations where the overall carbon footprint (including scopes 1-3 under the GHG protocol) can be reduced by 69% using commercially proven technologies. In the future, heat recovery using heat pumps can enable carbon-negative operation of data centers that use waste heat to drive carbon dioxide removal systems.

Deploy all available automation and analytics technologies in the operational phase to minimize the amount and carbon footprint of energy consumed. This includes shifting heat from high-cost, energy-inefficient conditions to lower carbon intensity conditions, accelerate detection and mitigation of defective operations that are consuming excess power and maximize reliability and uptime to prevent outages.

Our analysis suggests that the greatest opportunities for managing energy use in data centers can come from the following activities:

- **Continuously upgrading IT** hardware to take advantage of improvements in technology and remain close to state-of-the-art efficiency of electronic components.
- **Increasing the use of digital control systems and automation** to integrate data from both IT and OT systems as well as any co-located power generation, transmission and distribution equipment, enabling:
  - Development of analytical tools (deterministic, AI or hybrid) and control strategies that exploit the full set of data available in an integrated automation system to optimize energy consumption, asset utilization and power source C-intensity with increasingly high-time resolution to achieve lowest possible carbon footprint of instantaneous energy use without compromising system availability.
  - Deployment of the full range of automation and analytics tools to maximize reliability and uptime of assets and prevent outage conditions that can damage assets (requiring repairs that increase embodied C footprint) and lead to spikes in use of energy or increased use of high C-intensity energy from backup power systems such as generators.
  - Early recognition and remediation of compromised equipment that is running inefficiently and using more power and/or causing a greater power draw from other systems compensating for the compromised equipment. Proactive detection of declining asset health is important for resilience as well as managing energy use.
- **Maximizing the supply of firm low C-intensity power** either by choice of location, co-location with renewable power assets or firm power purchase agreements.
- **Deploying battery energy storage systems** to store variable renewable energy and support resiliency of supply, meet power backup requirements with lower C-intensity than fossil-fueled generators and exploit opportunities for daily price arbitrage while avoiding high C-intensity peak grid power.
- **Replacing legacy high global warming refrigerants** in CRAC and DX cooling systems with low global warming potential refrigerants to reduce the embodied carbon footprint (scope 3 impact) of the data center.
- **Using thermal energy storage systems to shift cooling** loads away from times when refrigeration systems are inefficient (peak daily heat) or electricity prices are high (peak power hours) and thereby reduce the overall C-intensity of power consumed.
- **Using heat pumps to boost the temperature of waste heat** from the data center and allow energy reuse for district heating in nearby communities of other low-medium grade heat applications.
- **Integrating data centers in more remote locations with direct air capture plants** for removing carbon dioxide from the atmosphere for geological sequestration, using the data center waste heat to offset roughly 40% of the energy needed for DAC and achieving overall carbon-negative operation.

INTRODUCTION

THE NEED TO MANAGE ENERGY USE IN DATA CENTERS

We need more data centers and they need to be designed to use less energy. The emergence of cloud computing and storage, expanding e-commerce, streaming entertainment and smart, connected devices together with broadened global access to the internet drove growth in global data center capacity to 149 zettabytes in 2024, and is projected to grow over the next five years to 394 zettabytes in 2028 (Statista, 2025). Increased demand for Internet of Things, Industry 4.0, autonomous mobility systems, advanced communications (such as 5G, 6G, autonomous vehicles V2V communication, etc.) and computationally intensive artificial intelligence in almost every application are only likely to accelerate this trend.

The global electricity demand for data centers was estimated to be 524TWh in 2023 (Farman et al., 2024), representing roughly 1-1.5% of global electricity use and 1% of global anthropogenic greenhouse gas emissions (Rozite, 2023), though the proportion of electricity used in data centers is higher in more developed economies and is approaching 4% of electricity use in the United States (Bloom Energy, 2025). It has been suggested that data centers could account for up to 10% of global electricity demand growth by 2030 (Poudineh, 2025).

Nonetheless, data centers are an important driver for sustainable development as they enable a range of efficiency improvements versus distributed computing. The scale of data centers allows them to rapidly exploit advances in novel algorithms and chip architecture, design and fabrication technology. Siddik et al. (2021) report that between 2010 and 2018 U.S. data center computing workloads increased nearly 550%, while energy use only increased 6% - pointing to significant enabling optimizations at every level of high-performance compute workloads, from transistor design to data center cooling techniques and control algorithms. Centralizing compute and storage resources also creates the potential to recover waste heat from IT operations for other uses, possibly enabling net carbon-negative future operation of data centers on an overall lifecycle basis.

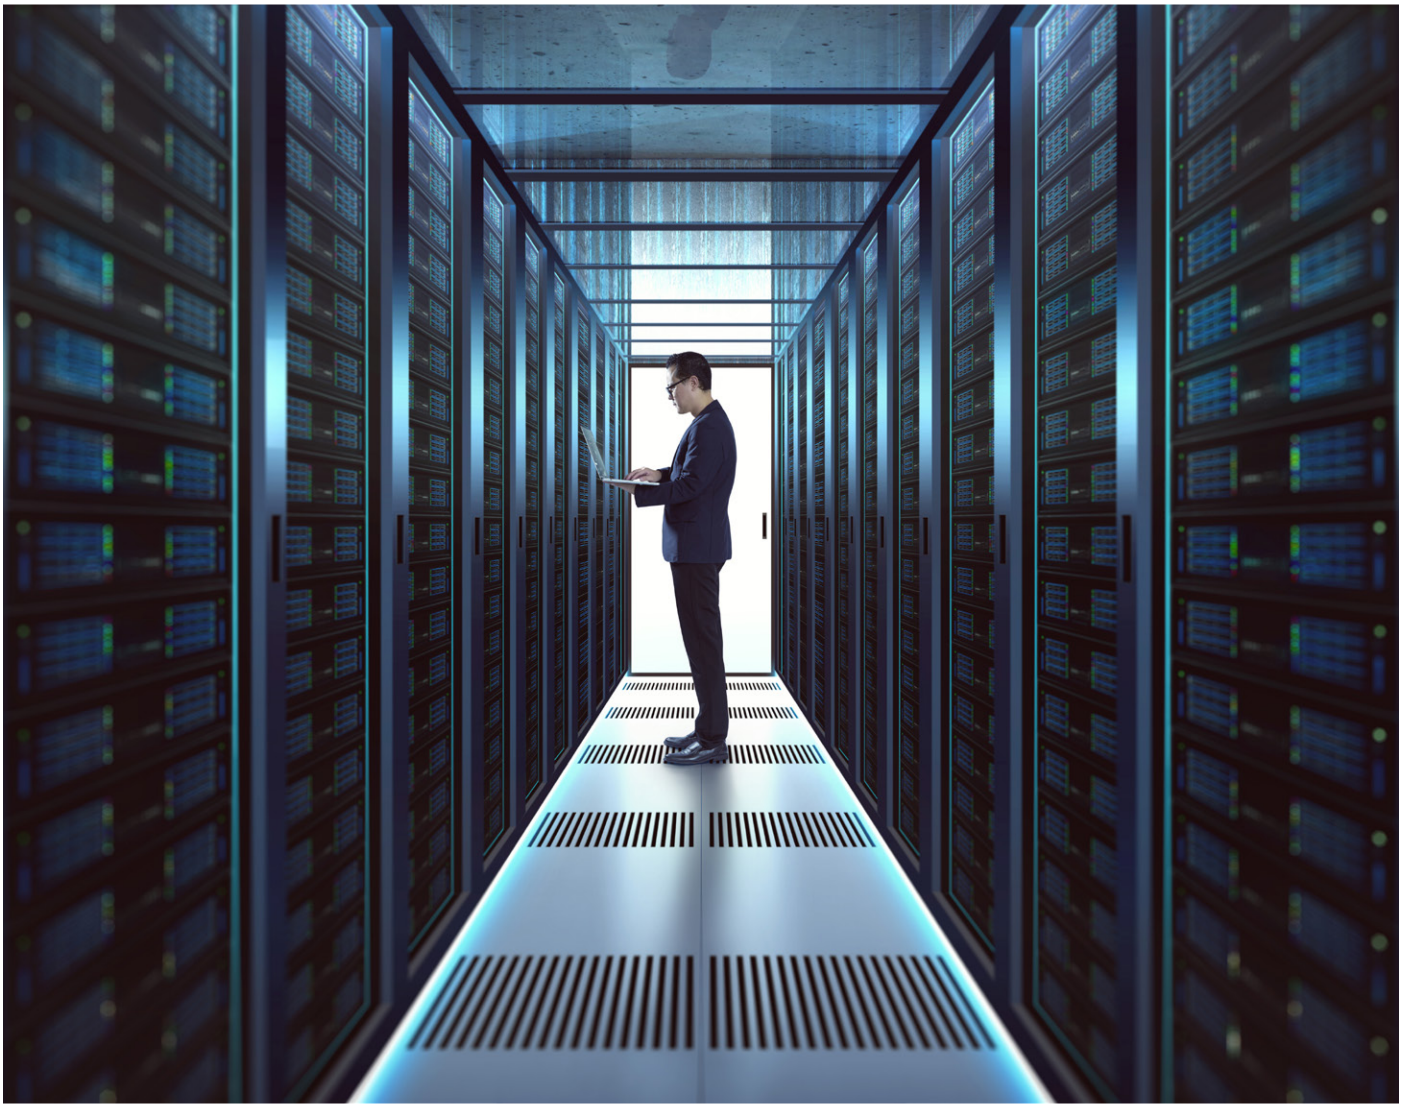

THERMODYNAMIC DESIGN OF DATA CENTERS

All the energy supplied to a data center is ultimately dissipated as waste heat, usually to the surrounding air. The large amounts of heat rejected by data centers make them potentially attractive for waste heat recovery when suitable applications for low-medium temperature heat can be found. The range of applications that can be served can also be increased by raising the temperature of the heat using heat pumps, though this increases the capital cost and electric power consumption of the data center.

The electric power consumption of a data center is set by the power required to run the computer hardware, known as the IT power, together with the power required for cooling systems and building ancillary services such as lighting, security systems, etc. Figure 1.1 shows a typical distribution of data center energy use, taken from Luo et al. (2019).

The energy efficiency of data centers is usually expressed in terms of the power usage efficiency (PUE), which is the ratio of the total electricity consumed by the data center to the electricity consumed for IT operations. A lower PUE is more efficient in use of electricity for DC operations; however, PUE does not account for any use of waste heat, so fixation on PUE can disincentivize energy reuse and lead to designs that have worse overall environmental performance. This is discussed further in Section 1.4.

Initial approaches to data center combined heat and power (CHP) focused on using gas turbine engines to deliver power to the DC with heat recovery from the exhaust gases for district heating or other low-grade heat applications. Darrow and Hedman (2009) provides a good review of such schemes and claimed 8-20% reduction in GHG emissions compared to running on grid power at that time, though we believe the impact is now lower due to the reduction in grid power C-intensity, as discussed below. Darrow and Hedman (2009) also discussed the barriers to adoption of CHP, particularly the redundancy requirements of Tier III and IV DCs and resulting cost implications. It is important to note though that redundancy is only actually required in the DC power supply and the cost of achieving redundancy for the district heat supply is much lower.

More recently, interest has increased in using data center waste heat for district heating, particularly in regions such as Europe where district heating is widespread (Acton et al., 2020). District heating schemes typically only operate at full capacity for four to eight months per year, so an alternative path for heat rejection is required during the summer months. District heating schemes typically require hot water to be supplied at 95°C, so a heat pump is needed to boost the temperature of the DC waste heat. The heat pump replaces the conventional data center chiller, but both capital and operating cost are increased and the additional cost must be recovered from the district heat customer(s). The coefficient of performance (CoP) of a heat pump decreases (becomes less efficient) as the temperature gap between source and target temperatures is increased. We estimate a CoP of 2.2 can be attained for a typical data center exporting heat at 95°C, implying that the PUE would increase to at least 1.5; however, the overall lifecycle impact can still be favorable, as discussed below.

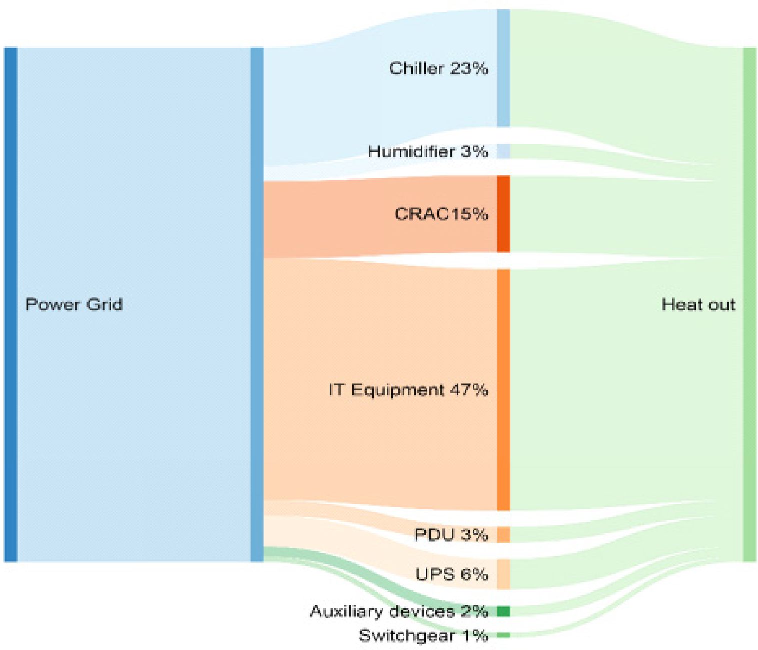

Figure 1.1: Data center energy use (Luo et al., 2019)

SITE SELECTION FACTORS

Site selection is the single most important factor for data center environmental footprint, as it creates the boundary conditions that govern the viability of many sustainable design options:

Power sourcing and low C-intensity power co-location: The easiest way to achieve a low carbon footprint is to locate data centers in countries or regions that have low carbon intensity (C-intensity) electric power, sometimes referred to as a "Go Where the Grid is Greenest" strategy. Among large economies Brazil (89.9% non-fossil power), France (88.8%) and Canada (82.1%) all have a substantial lead in getting to net zero electricity (all data from Energy Institute (2023)), but even within countries there can be regions that have low carbon intensity electricity supply, such as Himachal Pradesh (100%), Uttarakhand (100%) and Kerala (70%) in India or Washington, Vermont and New Hampshire in the United States (see Figure 1.2).

Co-location with low C-intensity power usually requires the data center to be somewhat remote from large population centers due to the large land requirements for solar and hydropower and siting issues for wind and nuclear power. Co-location with low C-intensity power production is not required in all countries as many allow operators to claim use of renewable power by purchase of long-term power offtake contracts (purchase power agreements or PPAs in the United States) or renewable energy certificates (RECs) that allow power producers to sell the environmental benefit of renewable power separate from the power itself. RECs currently provide a potential fast path to carbon neutrality but are likely to come under increasing scrutiny and regulation in future as environmentalists push for them to demonstrate additionality on a shorter time increment basis.

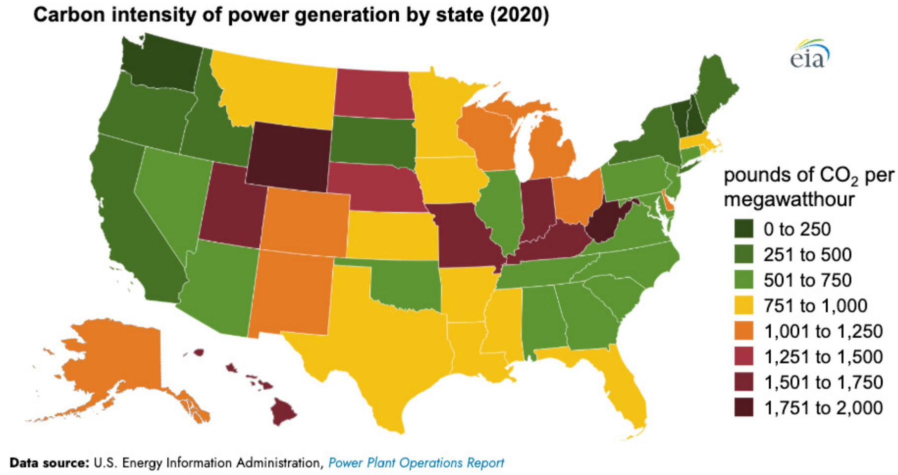

Figure 1.2: Carbon intensity of power generation by state (EIA, 2024)

Availability of free cooling: Colder climates provide opportunities to reduce cooling loads by increasing the efficiency of refrigeration and HVAC systems or incorporating air side economizers to provide free cooling. Free cooling can also be obtained overnight at higher altitudes.

Proximity to population centers: All data centers require a skilled workforce, good power and a communications infrastructure. Location near to regions of high population density is also a critical factor in enabling heat recovery for district heating as it is not feasible to transfer hot water over long distances. Set against this, land is expensive in areas of high population density and noise can be a concern in residential areas.

Resilience to climate change: Data centers are typically not located in regions that have high seismic activity or are prone to other natural disasters; however, the impacts of global warming will increase the frequency and severity of natural disasters in many regions. In addition to increased likelihood of hurricanes and coastal flooding, designers should consider the potential impact of increased rainfall and river flooding (global warming will increase precipitation in many areas) and forest fires.

Taxes, incentives and regulations: Some locations incentivize design of more sustainable energy efficient data centers by promoting investment in renewable power, energy storage systems, district heating or other methods of energy reuse either through grants, tax credits or regulations. These incentives can dramatically improve the economics of designs that would not otherwise pass investment criteria. Conversely, in many locations regulations and permit requirements increase the complexity and cost of building a more sustainable design and favor the simplest approach even if that design is least sustainable.

LIFECYCLE ANALYSIS

While the greenhouse gas (GHG) emissions associated with operational energy use are the primary environmental footprint of data centers, there are also significant environmental impacts from other stages in the lifecycle of a facility. It is important to understand that DCs that are powered by low C-intensity electricity (from nuclear power or renewable sources such as wind, solar or hydroelectric power) still have a significant carbon footprint from the embedded carbon in the facility itself as well as the embedded GHG footprint of the electric power and any impacts of end-of-life disposal activities. To fully weight all these contributions, it is necessary to carry out a life cycle assessment (LCA). Lifecycle assessment can also be used to evaluate the impact of different approaches to waste heat recovery and reuse.

Many LCAs of data centers in the literature study relatively small data centers (< 10MW average power draw) with older computer hardware and cooling systems that may not be reflective of the current state of the art. We therefore performed our own LCA using SimaPro 9.5.0 and the eco-invent 3.8 database and following an ISO 14040/14044 methodology to establish the global warming potential in metric tons equivalent of carbon dioxide per year (t CO_{2e}/y) of different options for heat and power integration.

To make a fair comparison between cases with heat recovery and reuse and cases that do not recover heat, we chose a system boundary that allows for potential energy export from the data center to a local community.

We evaluated the following cases:

Case A (base case): a standalone data center with 50MW average operating power operates on U.S. grid average C-intensity electricity and rejects all heat to the atmosphere. A local community of X homes is heated for six months of the year using natural gas at 80% heater efficiency.

Case B (traditional CHP): the data center of case A operates on electricity from two 25MW gas turbine engines with 45% efficiency. Waste heat is recovered from the turbine exhaust to provide district heat via hot water at 95°C to Y homes in the community, but no waste heat is recovered from the data center IT operations. The remaining X-Y homes in the community are heated with natural gas as in case A.

Case C (renewable power, no heat recovery): the data center of case A operates on dedicated renewable power (either co-located or via a firm power purchase agreement) and rejects all heat to the atmosphere. The local community of X homes is heated with natural gas as in case A.

Case D (renewable power with energy reuse): the data center of case C uses heat pumps with coefficient of performance 2.2 to boost the temperature of the exhaust heat and deliver water at 95°C to heat X homes in the community. This case allows us to determine X, which was found to be 25,370 homes.

Since wind power and solar power have different carbon footprints (due to the embodied C), we ran different versions of cases C and D for solar and wind power. Full details of the calculations, energy flows and breakdowns of the carbon footprint in each case are given in Appendix 2. The results are summarized in Figure 1.3.

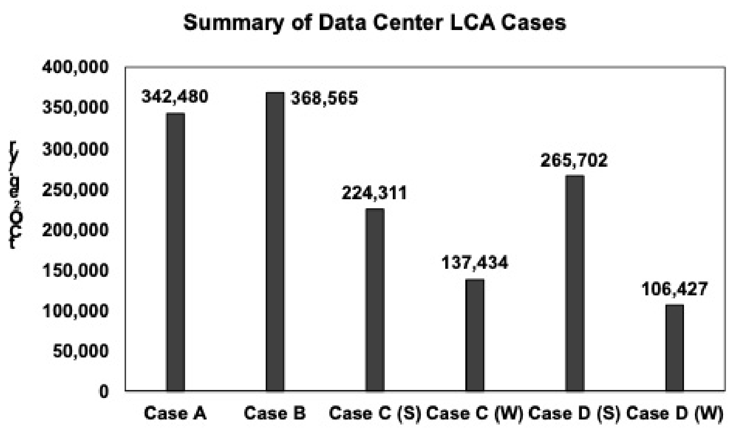

Figure 1.3: Lifecycle analysis of data centers with different power and heat recovery options

In the base case, roughly 20% of the lifecycle GHG impact of the data center is due to the embodied carbon in the facility, hardware and infrastructure and 80% is due to the electricity consumed in operation. Moving the data center to operation on dedicated gas turbine engines (case B) increases the overall carbon footprint, as the electric power C-intensity increases versus the C-intensity of the U.S. grid and this increase is greater than the savings from providing district heat to the community.

Operating the data center on dedicated renewable power (case C) gives a 34% reduction in total system carbon footprint for solar power and 60% reduction for wind power. The higher carbon intensity of solar power comes from the SimaPro LCA database and reflects the high use of coal-fired electric power in the production of polysilicon for photovoltaic cells, which are currently largely made in China. As China decarbonizes its power generation this carbon intensity should fall. Advances in solar power technology should also lead to further improvements in the carbon intensity of solar power. The carbon footprint of wind power is largely due to the large amounts of concrete and steel required for the foundations and structural support of the wind turbines and these GHG impacts will also be reduced as the carbon intensity of the concrete and steel industries is addressed. The carbon intensity of wind power will also continue to drop as technology improvements lead to increased wind plant energy capture and as offshore wind developments with higher capacity factors come onstream (Dykes et al., 2019). We expect to see continued reductions in the carbon footprint of new renewable power installations, but wind power will retain the advantage for at least the rest of this decade. Note that these savings would only be possible with a firm PPA for renewable power supply, and we did not account for additional energy needed to "firm" the supply of renewable power by incorporating energy storage systems into the design or the embodied carbon of energy storage. This is discussed in more detail in Appendix 2.

Operating the data center on dedicated solar power and using a heat pump to provide district heat to the local community (case D – solar) increases the system carbon footprint versus stand-alone operation with solar power. This is because the carbon footprint of the additional solar power needed for heat pump operation for a full year (94 $ktCO_{2e}$/y) is greater than the savings in home heating fuel for the six months that require heating (53 $ktCO_{2e}$/y). This option would therefore not be attractive until the carbon footprint of solar power decreases. The lower C-intensity of wind power means that case D for wind only requires 22 $ktCO_{2e}$/y for heat pump operation and so achieves the lowest overall GHG footprint with 69% GHG savings relative to the base case. It is worth noting that the PUE of case D with wind power would be at least 1.84, showing that PUE disincentivizes heat recovery schemes.

INTEGRATION WITH CARBON CAPTURE AND SEQUESTRATION

The lifecycle analysis showed that one of the disadvantages of heat recovery to district heating schemes is that the waste heat can only be used for half the year. There are relatively few process industries that operate on low-medium grade heat; however, an area of emerging interest is carbon dioxide removal (CDR) from the atmosphere by direct air capture (DAC). One of the advantages of DAC is that it can be carried out anywhere that is suitable for sequestration and so can be co-located with sources of low-cost energy.

The temperatures needed for district heating are at the lower end of the feasible range for providing heat input to DAC plants. Direct air capture technology is still in its infancy with a tiny number of demonstration scale plants in operation, but the separation processes used all require medium-low grade heat to regenerate the solvents or adsorbents that are used to scavenge CO_2, allowing the carbon dioxide to be collected and compressed for geological sequestration. Honeywell UOP has a 70-year history of commercial acid gas capture technologies and estimates that DAC systems could recover carbon dioxide from air with an input of 1500 kWh/t of heat at 95°C and an additional 434 kWh/t of electric power for air blowers and CO_2 compression. This would allow close coupling of a data center to a carbon capture and sequestration (CCS) plant that would remove carbon dioxide from the atmosphere and inject it underground. A schematic of this is shown in Figure 1.4. We performed additional lifecycle analysis cases (Case E) for both wind and solar powered data centers to estimate the potential greenhouse gas impact of coupling a data center and DAC plant.

Our lifecycle analysis calculations suggest that if such a facility is powered with low C-intensity electricity it can be carbon negative. For the 50MW data center studied, we estimate that operation with solar power and DAC could remove a net 190 $ktCO_{2e}$/y, while operation with wind power could remove 396 $ktCO_{2e}$/y. The capital cost is of course dramatically higher as is the PUE when a DAC plant is included in the design. Details of the calculations are given in Appendix 2.

We do not expect to see widespread deployment of DAC plants this decade as the technology is still in demonstration stage and has one of the highest marginal costs of CO_2 abatement compared to options such as fuel switching and electrification using renewable power. Nonetheless, the integration of DAC with data centers would provide a societally beneficial use for the waste heat of the data center and provide about 40% of the energy needed for DAC for free. For a data center operator that has aggressive goals on reaching carbon neutrality, this approach would also have the advantages of eliminating the full scope 3 emissions of the site (which were included in the LCA basis) and providing a measurable and verifiable means of offsetting emissions from other operations (since the CO_2 that is captured and sequestered can be accurately measured and monitored).

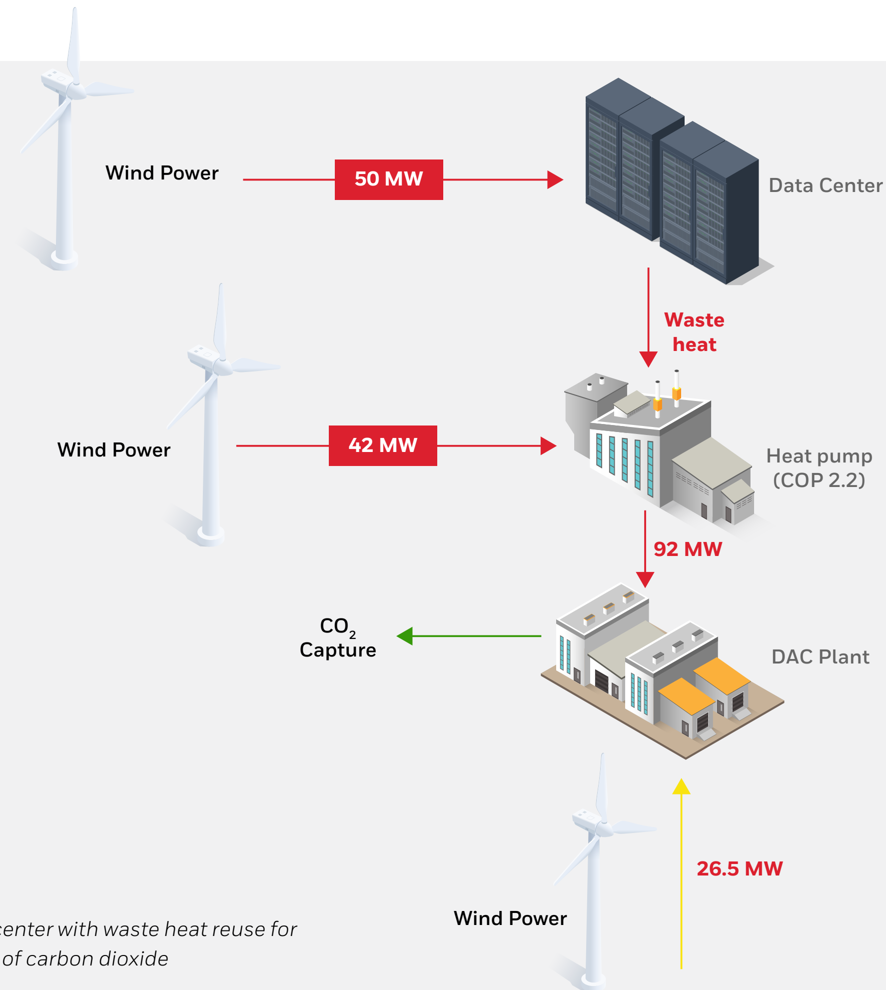

Figure 1.4: Data center with waste heat reuse for direct air capture of carbon dioxide

ENERGY MANAGEMENT

POWER SCHEDULE MANAGEMENT

When relating energy consumption to carbon footprint of electric power, all kWh of electricity is not equivalent. Every region experiences seasonal, weekly and daily variation in power demand with different daily patterns during heating and cooling seasons as shown in Figure 2.1. The highest peaks (which set grid capacity needs) are typically seen during hot summer afternoons when demand for air conditioning and industrial cooling is highest.

Peak power demands typically require electric utility companies to generate power using dispatchable assets. Since the daily peak usually occurs from 6:00 -10:00 p.m., the problem of meeting peak demand is exacerbated by the fact that solar power is not available during peak hours, so the utility either needs to deploy large-scale energy storage to balance the grid or else fall back on dispatchable assets such as pumped-storage hydropower and gas turbine engines to meet peak needs. Gas turbine engines operated in peaking mode are expensive (the capital cost is recovered over fewer hours per year), as well as producing electricity of higher C-intensity than the grid average. Utilities therefore generally incentivize large-scale consumers to practice demand reduction during peak hours to reduce strain on the grid, often through a combination of incentives for load shedding and time variable pricing (TVP) or punitive pricing for exceeding demand thresholds.

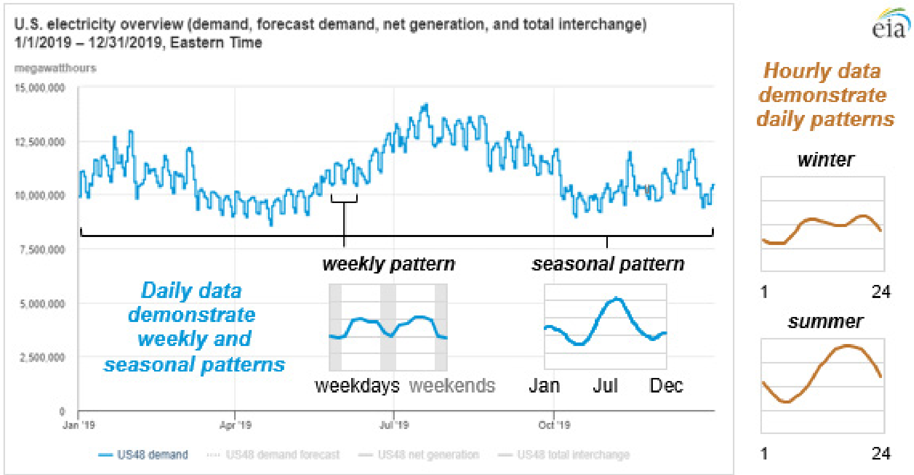

Figure 2.1: Variation in US electricity demand (EIA, 2020)

Figure 2.2 shows varying C-intensity by geography, over daily period across the months of the year (IEA, 2022). California's and Germany's strong solar deployments dramatically reduce the C-intensity during the daylight hours; but winter months have shorter days therefore less solar capture, so the C-intensity is greater during these periods. The UK's stronger dependence on wind power gives a less dramatic seasonal variation, but there can still be significant variation through the day with peak power almost double the C-intensity of overnight power. France's high reliance on nuclear power and solar output during the summer give the most stable supply of low C-intensity power; however, few regions are as accepting of nuclear power as France.

The high price and C-intensity of peak power create an incentive for DC operators to shift as much power demand as possible away from peak hours and to deploy battery energy storage systems (BESS) or thermal energy storage systems (TESS) to further reduce peak grid power draws and take advantage of preferential electricity pricing that may be available for demand reduction during peak hours.

The variation in data center hourly power demand depends on the use application. Co-located DCs and enterprise DCs have demand patterns that reflect the operating hours of the businesses they serve and can often schedule backup and maintenance operations overnight to avoid peak power usage. Hyperscale data centers are more at the mercy of consumer demand and can experience peak activity at the same time as the electric grid peak (and for the same reasons), making peak power avoidance more important for the hyperscale DCs.

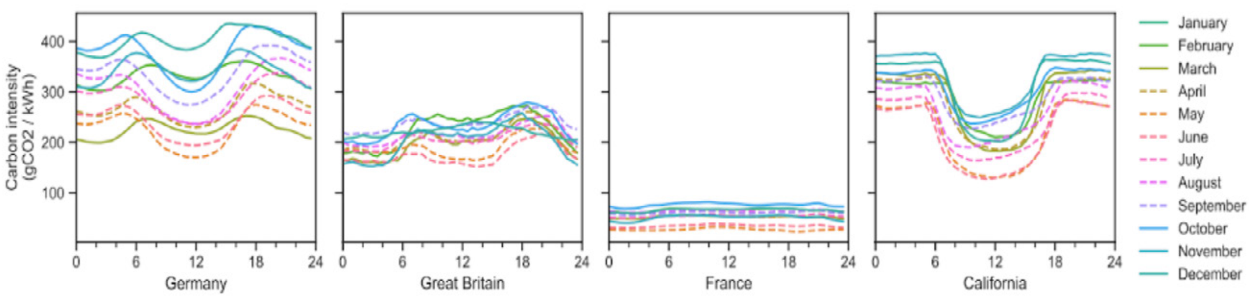

Figure 2.2: Carbon intensity of electricity by month and hour for Germany, Great Britain, France and California (Reply, 2022)

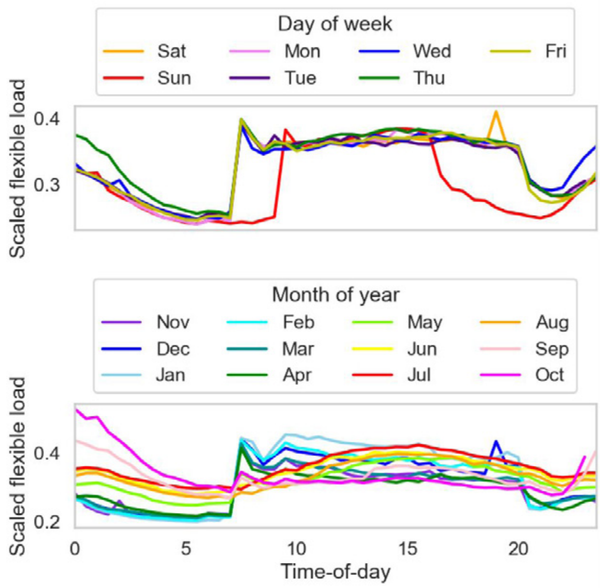

Figure 2.3: Data center hourly usage rates (Krishnadas and Kiprakis, 2020)

Figure 2.3 shows data for typical U.S. hourly data center usage by day of the week and month of the year, based on data from Krishnadas and Kiprakis (2020). Overnight loads are highest in September and October reflecting end of fiscal year activity. The data shows that overnight operations typically run at about half the level of daytime operations, but usage remains relatively flat during the day, with peaks occurring at the start of the business / school day and during the evening peak power period around 7:00 p.m.

Hyperscale DC operators like Microsoft offer spot pricing that can be discounted by up to 90% versus pay-as-you-go pricing. This attractive pricing can be available during times of idle compute capacity, which can also coincide with times of lower demand during cheaper power rate (and lower C-intensity) times of the day. A typical downside with such plans is the risk of workload being evicted on short notice; but certainly, for non-critical workloads, this is a win-win scenario that allows the DC to achieve a higher capacity utilization, exploit lower price off-peak power and maintain a consistent supply of waste heat for heat recovery operations.

CHIP AND RACK LEVEL COOLING

The combination of increase in the size of data centers and higher chip power present three challenges to data centers:

- how to quickly remove the heat from the chips so they can properly function;
- how to improve building design and cooling efficiency to reduce overall data center cooling power usage and power usage efficiency; and
- how to reduce water usage.

All three challenges can be addressed by moving to two-phase liquid cooling. Chip power is expected to reach 1000W TDP (thermal design power) in the 2024-launched AI chips, compared to chips launched in 2020 with less than 400W TDP and we expect the chip power will reach beyond 1500 TDP by 2028.

There are two main cooling technologies, air cooling and liquid cooling. Traditionally, data centers have been air cooled with PUE of ~1.5 and max heat removal capacity of <50kW/rack. It is well known in heat transfer that convective cooling by gases gives low heat transfer coefficients (typically 10-30 W/m2K) compared to convective cooling by liquids (200-800 W/m2K) and evaporative cooling (1000-2500 W/m2K) (Towler and Sinnott, 2022), so higher chip power per rack implies either increasing the available heat transfer surface or upgrading to a more effective heat transfer medium. Figure 2.4 shows how higher rack density can be achieved while improving PUE through deployment of more effective heat transfer mechanisms. Older air-cooled systems often used a cooling water system or chilled water system to cool the air, while more recent designs use direct expansion (DX) in which the air is cooled directly by a refrigeration plant (similar to air conditioning). Higher rack power densities up to 100kW/rack can also be achieved in air cooled systems by using rear-door heat exchangers (RDHx) to increase the rack level cooling capacity.

An additional problem with older computer rack air conditioning (CRAC) and DX systems is that they usually use refrigerants such as R-410A that have very high global warming potential (the GWP of R-410A is 2088). These older refrigerants contribute to the embodied carbon footprint of the data center and should be replaced with newer low global warming potential refrigerants. In the case of R-410A, it can be substituted with R-454B, which has a GWP of 466. Honeywell is also working on a next generation of refrigerant blends that will operate efficiently over the same temperature range with GWP < 150.

Liquid cooling can be further categorized into four types, single-phase direct to chip, single-phase immersion cooling, two-phase direct to chip, and finally two-phase immersion cooling. Single-phase liquid systems typically use treated city water or water glycol as the cooling media, or mineral oil in the case of immersion cooling. Single-phase cooling benefits from design simplicity and significantly improves PUE compared to air cooled systems, with typical PUE around 1.05; however, the cooling power is limited to about 1.5 kW TDP per chip or about 150kW/rack. The use of once-through city water also poses significant environmental concerns, leading to more complex designs in which the liquid coolant is recirculated between the racks and a chiller system that cools the water and rejects heat to atmosphere.

Two-phase liquid cooling uses a coolant that partially evaporates, giving higher heat transfer coefficients and enabling 1.5-2.5kW TDP or 250 kW/rack with PUE close to 1.02. Two-phase cooling can be used in both cold plate and immersion cooling designs. Two-phase immersion cooling also has the potential to eliminate the use of the chiller, tubes and other supporting equipment used in air cool and direct to chip cooling technology but requires substantially modified rack designs and is not yet widely adopted.

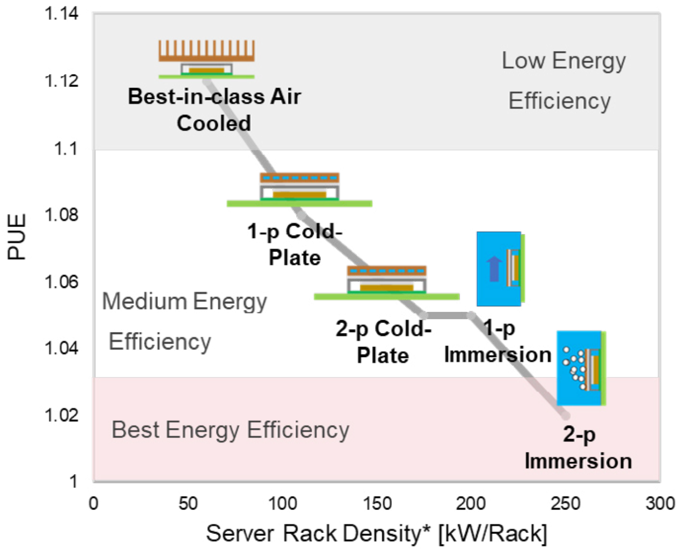

Figure 2.4: Typical power usage efficiency of different cooling technologies

Honeywell, as a leading refrigerant and thermal management innovator, offers solutions across the full spectrum of cooling technologies:

Refrigerant
Vapor-compression cycle chillers will continue to exist in most air- and liquid-cooled facilities, using a refrigerant to chill the air or liquid that is used to cool the data center. Chillers need to have high efficiency and use refrigerants that have low global warming potential (GWP). The Honeywell Solstice® 1234ze refrigerant is an energy-efficient alternative to traditional refrigerants in air-cooled and water-cooled chillers with >90% reduction in GWP and 3-4% energy savings versus R-134a.

Coolants for two-phase cooling
Once chip power reaches ~1.5kW, two-phase cooling is required. Honeywell is working on offerings for both direct-on-chip and immersion cooling:

For direct on chip, Honeywell refrigerants with various boiling points, such as R-515B, R-1233zd etc., are commercially available and in active pilot test with leading cooling solution providers.

For both direct-on-chip and immersion cooling, concerns about the long-term impact of per-fluorinated alkane substances (PFAS) are leading to increased regulation of fluorine-containing compounds. Honeywell is actively pursuing solutions that are non-PFAS (per EPA definition), low GWP, low dielectric and non-flammable.

Heat recovery
In ideal scenarios, waste heat rejected by data center can be fully recovered and reused as discussed in Section 1. Honeywell offers a range of Solstice refrigerants for heat pump applications that enable the conventional chiller to be replaced with a heat pump that delivers heat at high enough temperatures for use in district heating or carbon dioxide recovery by direct air capture. These are discussed in Section 2.4.

THERMAL ENERGY STORAGE

Thermal energy storage systems (TESS) offer a novel approach to data center cooling that tackles both energy efficiency/demand response/ smart grid integration and reliability/ resilience. TESS shift cooling energy use to non-peak times, reducing the consumption of high C-intensity peak electricity, see Figure 2.5. Two approaches are possible: Sensible or chilled thermal energy storage (TES) media such as water or latent energy storage in a phase-change material. Phase change material systems (PCM-TES) use heat to melt an inventory of a thermal storage material, which can later be re-solidified either using off-peak power or by natural cooling overnight in regions where there is a large daily temperature swing. This stored energy can be used for cooling purposes, significantly reducing the energy consumption at peak demand times. By spreading thermal energy production over 24 hours, this solution can reduce chiller demand charges by 30 to 70% (Trane, 2023).

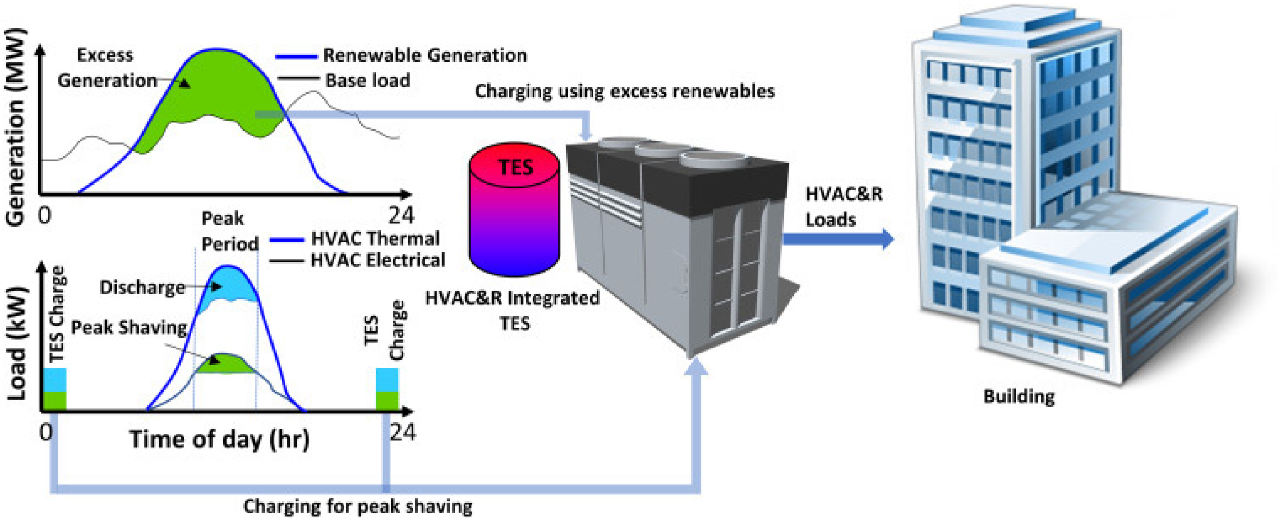

Figure 2.5: Chiller integrated TES for load shaving and shifting (Ragoowansi et al. 2023)

Both sensible (chilled water) and PCM-TES are applicable in DCs, but certain criteria must be met for economic feasibility. A system can be appropriate when maximum cooling load is significantly higher than average load. High demand charges, and a significant differential between on-peak and off-peak rates, also help make TESS economic. They may also be appropriate where more chiller capacity is needed for an existing system, or where back-up or redundant cooling capacity is desirable. Besides shifting load, TESS may also reduce energy consumption, depending on site-specific design, notably where chillers can be operated at full load during the night. Both forms of TESS can be effective for power reduction in regions such as deserts where three is a large daily temperature swing, as the coefficient of performance of refrigeration equipment is much higher when the heat can be rejected to a lower ambient temperature overnight.

Moreover, thermal energy storage systems enable a high-density data center to survive a power outage without damage to IT equipment. TESS are a cost-effective way to provide temporary cooling in high- and medium-density data centers, potentially preventing millions of dollars of damage to IT equipment. In the case of new construction, plant expansion or rehabilitation of an existing cooling system, TESS also have the benefit of reducing capital costs since adding TESS can cost much less than installing equivalent new chiller plant capacity. TESS generally reduce the required investment in conventional chilling equipment by allowing it to be sized to the average cooling load. For those situations, TESS can achieve a rapid payback and possibly an immediate net capital cost saving.

In sensible heat storage (SHS), thermal energy is stored by raising the temperature of a material, typically solid or liquid such as water. Latent heat storage (LHS) is achieved using phase change materials (PCMs), i.e. materials characterized by high latent heat of fusion, which through melting or solidification can store or provide heat respectively. Common PCM are e.g. ICE, paraffins, fatty acids, sugar alcohols, and salts, as a pure material or as a mixture. Organic PCM are often called bio-based, if produced from biological sources. Different PCMs are used according to their working temperature ranges and application temperatures range from -20 °C to +200 °C. The PCM-based TES is a superior way of storing thermal energy due to their large latent heat with a relatively low temperature or volume change (Du et al., 2018).

As an example, to shift 1000 r-ton of cooling for four hours using TES requires 38,000 gallons of PCM. Assuming chiller plant efficiency of 0.60 kW/ton on peak load day, de-energizing the chiller equipment reduces the electric load by 600 kW. For the same application, a chilled water storage option will require over 290,000 gallons of water (8X compared to PCM-based TES) requiring larger footprint area as shown in Figure 2.6.

Figure 2.6: Chiller integrated TES for load shaving and shifting (Ragoowansi et al. 2023)

BUILDING HVAC AND CONTROLS

Current data centers largely rely on air cooling at the rack level and so HVAC systems play a critical role in heat management. ASHRAE Technical Committee 9.9 sets HVAC standards for mission critical facilities, data centers, technology spaces and electronic equipment. Best practices in air flow and HVAC design are extensively documented, see for example Van Greet (2010), Memarzadeh et al. (2013) and Acton et al. (2020). Air circulation systems flow cold air over the server racks and collect warm air that is then cooled and recirculated. Several approaches are used to reject heat from the warm air to atmosphere:

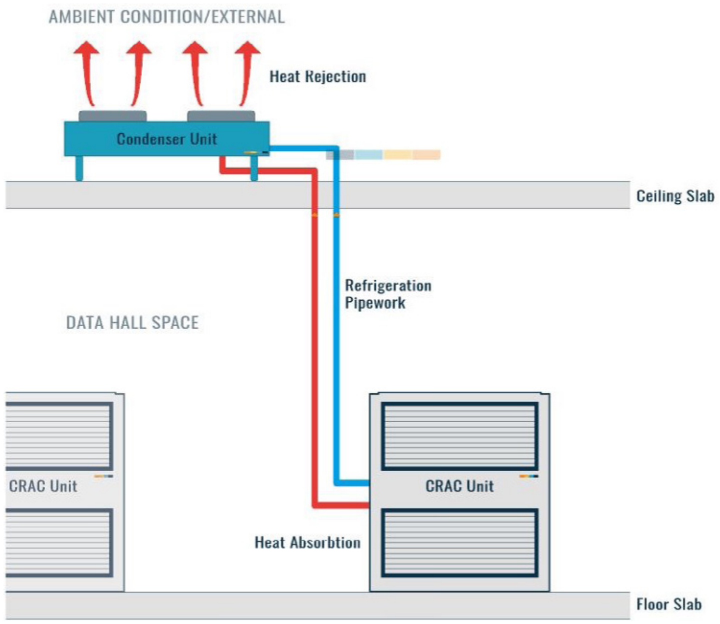

Figure 2.7: Direct expansion cooling system

Direct Expansion: Direct expansion or DX cooling system is a type of air-conditioning system that removes heat from a space through evaporation and condensation of a refrigerant. A DX system operates on the same principle as a home air conditioner, see Figure 2.7. In a DX system the evaporator is placed inside the space to be cooled. The refrigerant enters the DX cooling coils, where it absorbs the heat from the air, such as the heat generated by critical equipment, and transforms to a gas. The refrigerant is then compressed and sent to the condenser located outside, where the heat is released. An expansion valve exists between the condenser and the evaporator to further cool the refrigerant before it is returned to the evaporator, and the entire DX cooling system offers a closed loop solution. DX systems can be effective in smaller data centers but generally lead to high PUE as the coefficient of performance of the refrigerator applies to the full cooling load.

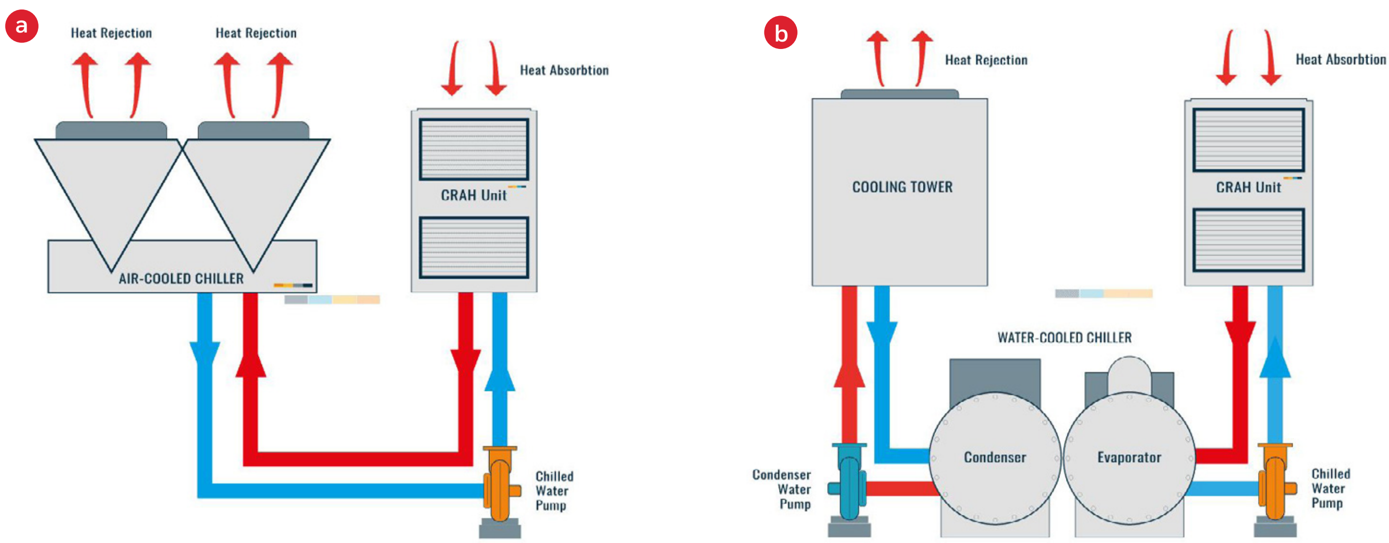

Figure 2.8: Chilled water-cooling systems: (a) air cooled, (b) water cooled

Chilled water: High efficiency chilled water systems cool the racks using recirculating chilled water. The chilled water then rejects heat to the outside. Chilled water systems can exhaust heat to the atmosphere or can allow for heat recovery using plate and frame heat exchangers for heating occupied space (offices typically). Air-cooled chillers are almost always located outside of a building and remove heat from the chilled water by exhausting the heat directly to the surrounding air, Figure 2.8(a). Water-cooled chillers typically use a refrigeration unit to generate the chilled water and reject heat from the hot (condenser) side of the refrigerator to outside air or to a cooling water system, Figure 2.8(b).

Chilled-water cooling systems are energy efficient; however, due to their complexity and many different parts, they are often more expensive to install and maintain. For this reason, they are usually only deployed in large buildings where the energy savings outweigh the cost of installing and maintaining the system.

Free cooling: in colder climates (or regions where there are cold temperatures overnight) chilled water systems can reduce energy use by rejecting all or part of the heat from the water to outside air through an air-side economizer. Economizers can be used in combination with chiller systems, reducing the load on the chiller when cold air is available.

Heat pumps: use of a heat pump allows the waste heat from a chilled water system (or two-phase liquid coolant system) to be rejected at a higher temperature for use in district heating or other energy reuse applications. A typical heat pump arrangement is shown in Figure 2.9, where the heat pump takes heat from the DC chilled liquid at 40°C and rejects heat to deliver water for district heating at 95°C.

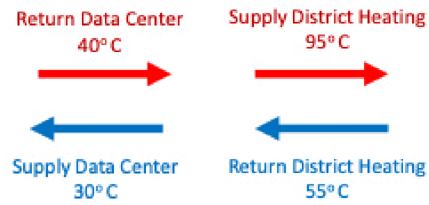

Figure 2.9: System configuration for DC chilled water to district heating hot water cascade.

Honeywell has experience in design and operation of heat pumps over a wide range of temperatures and works with heat pump OEMs to customize the refrigerant (or refrigerant blend) to optimize thermodynamic performance while meeting other safety and cost objectives. For these applications, we would recommend a 1234ze(E) refrigerant, which would give a CoP of 2.2, as shown in Table 2.1. For comparison, a single stage CO2 system operating over the same temperature range would require more compression work and have a CoP of 2.0. Since the CoP of a heat pump is the ratio of energy delivered to work required, that means a R1234ze heat pump would require 17% less energy than a CO_2-based heat pump to remove the same amount of heat.

SYSTEM CONFIGURATION/ REFRIGERANT	SINGLE STAGE R1234ze(E)	SINGLE STAGE CO_2
Evaporator Temperature First Stage (°C)	28°C	28°C
Condenser Temperature First Stage (°C)/Gas Cooler Pressure (kPa)	100°C	22000 kPa
Compressor Efficiency First Stage (%)	60%	60%
Total Waste Heat (MW)	100	100
Heating Capacity (MW)	181	197
Overall Heating COP	2.2	2.0

Table 2.1: Heat pump performance with Solstice refrigerant versus carbon dioxide refrigerant

Data center HVAC systems, chillers and heat pumps are generally controlled by programmable logic controllers (PLCs) as part of the DC operational technology (OT) system. The OT system is distinct from the information technology (IT) system that refers to an extensive collection of data storage devices like servers and software, network hardware such as cables and switches, and communication devices and protocols. The primary difference between IT and OT is how data is used. IT is more focused on broad business needs. This means it deals with transactions, voice communication, data storage – often in unstructured databases – and other meta-level data needs. By contrast, OT deals with machine-driven data meant to be consumed in real time at the user or manager level. This data comes from the control of physical devices through digital technologies such as software with advanced analytics engines dedicated to optimizing processes and is usually structured time sequence data that indicates equipment condition and operational effectiveness.

OPERATIONS MANAGEMENT

Given the exponentially increasing demand for data – as of April 2024, the known pipeline of future hyperscale data centers stood at 440 facilities– data center managers are often forced to do more with less, while at the same time being pressured to improve uptime, reduce costs and minimize energy consumption (Synergy Research Group, 2024). User-friendly tools that automate processes can help reduce the chance for human error, which caused major outages over the last three years among 39% of organizations surveyed by the Uptime Institute, often driven by ignored or inadequate procedures. It can also help improve a data center's overall efficiency.

Enabling reliability and uptime are critical to sustainable operation of data centers, as well as having significant financial implications. The Uptime Institute data center resilience survey reported that in 2023 55% of data centers had experienced an outage in the past three years, with 54% of the outages costing >$100,000 and 16% costing >$1 million (Donnellan and Lawrence, 2024). Power losses played a role in 52% of the reported outages. Aside from the obvious financial impact of outages, data center outages and power disruptions have several impacts on carbon footprint:

- Before a downtime event, the compromised equipment is often running at an inefficient operating point and using more power and/or causing a greater power draw from other systems compensating for the compromised equipment. These scenarios unnecessarily increase the energy consumption, operating cost and carbon footprint of the data center. Proactive detection of declining asset health is important for resiliency as well as sustainability.
- Loss of cooling systems while the IT systems are safely shutting down on UPS power exposes IT hardware to potential thermal damage, as heat is still being released during shutdown operations but is no longer being efficiently removed.
- Any damaged equipment must be replaced ahead of the planned service life, increasing the embodied carbon footprint of the data center.
- Repairs and system backup restore functions unnecessarily increase energy use.
- Operation on backup power usually involves use of relatively inefficient small fossil-fueled generators and has much higher C-intensity than operation on grid power.

It is important that the IT, OT and power systems all be designed for high reliability as well as the redundancy required for Tier III and IV specifications and that operators use the full range of automation and analytics tools to anticipate, avoid and mitigate situations that could cause an outage.

Honeywell Forge Sustainability+ for Buildings | Power Manager
Honeywell Forge Sustainability+ for Buildings | Power Manager is a turnkey end-to-end solution for optimizing on-site supply side resources and building assets from project design and execution to ongoing operation and maintenance for commercial facilities. It enables orchestration and optimization using ML algorithms of building demand side and supply side assets based on grid consumption, utility rates, and building demand. Supply side assets include on-site energy generation (Solar PV and traditional fuel generation) as well battery energy storage. The Power Manager solution reduces operational and utility costs, increases site resiliency and uptime, and helps customers meet their sustainability goals.

POWER MANAGER CAN

- Track and analyze carbon emissions by asset and reduce use of conventional fuel generation with renewables
- Help reduce hidden charges from energy services and orchestrate energy across both supply and demand side
- Help increase revenue streams with market participation in demand response programs and Virtual Power Plants
- Help use cleaner sources of alternate power to provide backup power along with remote monitoring
- Integrate with Experion controls for automated peak shaving, frequency and voltage regulation as well as microgrid controls

DISTRIBUTED ENERGY RESOURCE MANAGEMENT SYSTEM (DERMS)

Honeywell has a range of offerings that can be applied to control and automation of distributed energy assets, such as solar, wind, pumped hydro, green hydrogen, and battery energy storage, enabling geographically distributed energy resources to be operated from centralized facilities and bringing data from multiple assets together to allow optimization of charge and discharge cycles for ES assets and creation of microgrids and virtual power plants (VPPs) that aggregate power from a range of consumers and generators to increase overall dispatchability. These systems can also be used when integrating or co-locating renewable power assets with a data center.

SCADA systems

Our proven Honeywell Ionic Control & Energy Management System SCADA solution, provides a single panel for viewing and controlling single sites or fleets of new and existing assets such as renewable power generation sites and BESS systems, with potential to bring all the data from multiple sites to a single location. Honeywell Ionic Control & Energy Management System SCADA can be deployed onsite or cloud hosted providing a scalable and cybersecure solution to control one or hundreds of sites.

Microgrid controls

Honeywell Ionic Control & Energy Management System optimizes selection of energy sources based on priorities for generator efficiency curves, dynamic grid power pricing, start/stop maintenance costs, weather forecasts, and carbon footprint reduction. Honeywell's microgrid controls are based on the proven ControlEdge™ RTU and PLC controllers, which are powerful, modular and scalable devices capable of all remote automation and control applications. The microgrid controls provide improved management of field assets through simplified and efficient remote monitoring, diagnostics, and management. ControlEdge RTU and PLC come with an extensive library of control algorithms for renewable energy and can be configured to provide stable high-availability edge control of assets during communication outages, while storing data in onboard memory for uploading when communications are restored. Cybersecurity is built into ControlEdge RTU and PLC with ISASecure EDSA Level 2 certification protecting the safety of the system, personnel and data.

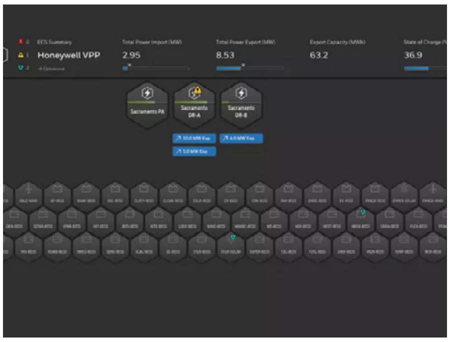

Virtual power plants

The Honeywell Virtual Power Plant solution within Honeywell Ionic Control & Energy Management System enables users to dispatch a network of distributed energy resources such as energy storage systems through a centralized control process, see Figure 4.5. By centralizing the dispatching process, users can aggregate distributed assets enabling them to participate in a variety of electricity markets, while helping to stabilize the grid. VPP functionality comes integrated with the Honeywell Ionic Control & Energy Management System. Asset owners can subscribe to the service for a monthly fee.

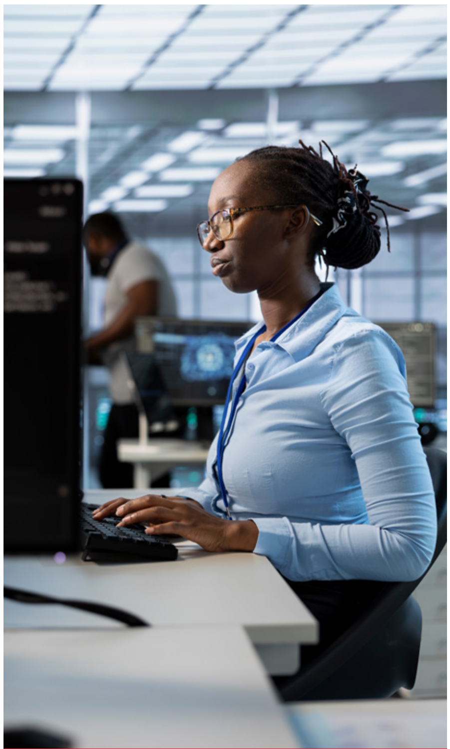

Honeywell has guaranteed $9.5B in energy and operational cost savings through more than 3,400 projects for customers around the world including federal agencies such as the U.S. Department of Defense and several branches of the U.S. military, U.S. Department of Veterans Affairs, NASA, General Services Administration, U.S. Department of Energy, Food and Drug Administration, and others. Many of these sites deploy microgrids with a combination of on-site generation, energy storage and demand load management.

DESIGN FOR RELIABILITY AND RESILIENCE

All Tier III and IV data centers require redundant power supply systems to enable continuity of critical operations in the event of grid outages. Low C-intensity backup power can be supplied by using diesel generators powered by renewable fuel or by storing electricity using battery energy storage systems (BESS). BESS systems have an advantage versus backup generators in that they can also be used for daily price arbitrage, allowing the DC operator to charge the BESS overnight (with wind power) or during the middle of the day (with solar power) and then use part of the charge to offset power demand during the evening peak power time when prices are highest while still meeting system availability requirements. This price arbitrage offsets the higher initial cost of BESS and makes BESS systems the overall most economical option.

Honeywell Renewable Diesel technologies

Honeywell UOP has been a pioneer in developing and licensing technologies for producing renewable fuels from a wide range of renewable feedstocks. While Honeywell does not directly produce any renewable fuels, our licensees are able to supply diesel fuel (hydrotreated vegetable oil or HVO) with lifecycle GHG impact 60-100% lower than fossil fuel-based diesel, depending on the feedstock used to make the renewable fuel. Honeywell has licensed 50 renewable fuels plants globally since commercially demonstrating this technology in 2013.

Honeywell Ionic™ Modular BESS

Honeywell Ionic™ Modular is Honeywell's second-generation BESS, following the first generation containerized 1-hour and 3-hour energy storage offerings released in 2022 and 2023. Honeywell Ionic Modular is a compact, end-to-end modular battery energy storage system (BESS) and energy management tool that delivers a significant reduction of installation costs, scalable modular architecture provides an optimized energy outcome, improves uptime and allows electricity market participation to help our customers increase their use of renewable electricity and meet corporate sustainability goals. Honeywell Ionic Modular is currently available with (LFP type) lithium-ion-based batteries.

Honeywell Ionic Modular includes Honeywell Ionic Control & Energy Management System and a chemistry-agnostic Battery Management System (BMS). Honeywell Ionic Control & Energy Management System helps users to manage and optimize energy use by improving uptime, maximizing arbitrage potential from peak shaving and providing the ability to create a Virtual Power Plant. The BMS provides insight into performance at the cell level and is configurable with advances in battery chemistry, insulating the end user from future supply-chain risks.

Key features of the Honeywell Ionic Modular BESS include:

- Scalable architecture allows you to right size the system for both front of the meter and behind the meter use cases.
- Proven lithium-ion-based cell chemistry with 730kWh modules scalable to any capacity.
- Compliant to energy storage standard UL9540.
- Optional, industry-leading off-gas detection which can enable earlier mitigation actions to prevent thermal runaway and fires.
- Integrated Honeywell controls to support all use cases.
- Turnkey installation from utility engagement, engineering, procurement, construction, commissioning, start-up, operations and maintenance. EPC scope is evaluated case by case.
- The batteries come pre-installed to reduce the on-site hours.
- The forklift-able design allows for fast installation without the use of expensive cranes.

Figure 3.1: Honeywell Ionic™ Modular BESS

DEMAND SCHEDULING AND LOAD MANAGEMENT

Power management without Automation (Electrical Power Monitoring System)

Electrical power monitoring systems (EPMS) record and provide data about power systems and power-related events. That information is used to manage power generation efficiencies, batteries and capacitor banks, gas or steam turbine relays and other systems in power generation stations and power substations. EPMS can visually display real-time or historical data. Supervisory control and data acquisition systems (SCADA) systems often use EPMS, especially those used in power plants. EPMS that include generator protection and control (GPC) relays and those that are integrated with SCADA can automate many power-related relays. This control helps increase power efficiency, especially in times of high draw. Better power management is helpful in terms of smoothing power demands. Smoothing out peak and low demand is often beneficial and lower in cost as the problem in energy systems is often not that average power is too high but that peak power draws exceed momentary power production. Most EPMS systems can be categorized into two types. Electrical Power Monitoring System and Electrical Power Management System. The Electrical Power Monitoring System monitors but takes no action to self-heal versus the Electrical Power Management System which through automation will self-heal and operate breakers by taking action from monitored data.

The trend is that more data center companies are moving to dedicated electrical power monitoring systems. This is to help further the ability to find and implement tighter tolerances in the overall facility operations gaining efficiencies not previously found with building management systems (BMS) alone.

Power Management with Automation (Electrical Power Management System)

In a large electrical power management system, there can exist automatic monitoring and control systems Tier III and IV data centers are designed with high equipment redundancy, creating opportunities to shift electric loads to duplicate systems and achieve improved uptime and energy efficiency.

The Honeywell Experion® PKS control system is a robust and cyber-secure distributed digital control system originally developed for control of large-scale critical assets such as chemical plants and oil refineries that typically have up to ~106 I/O requirements. Experion control systems are deployed in more than 5,000 sites globally connected to tens of millions of assets and field data shows that redundant systems have 99.9996% availability, with ability to carry out on-line updates and system migrations without interrupting operations. As these sites typically include substantial electrical generation and distribution systems, Experion has been developed to be compliant with IEC 61850. Honeywell Ionic™ Control & Energy Management System can be used to manage renewable power generation assets and local microgrids and can be extended to incorporate electric power system automation functionality.

DESIGN FOR CIRCULARITY AND NATURE IMPACT

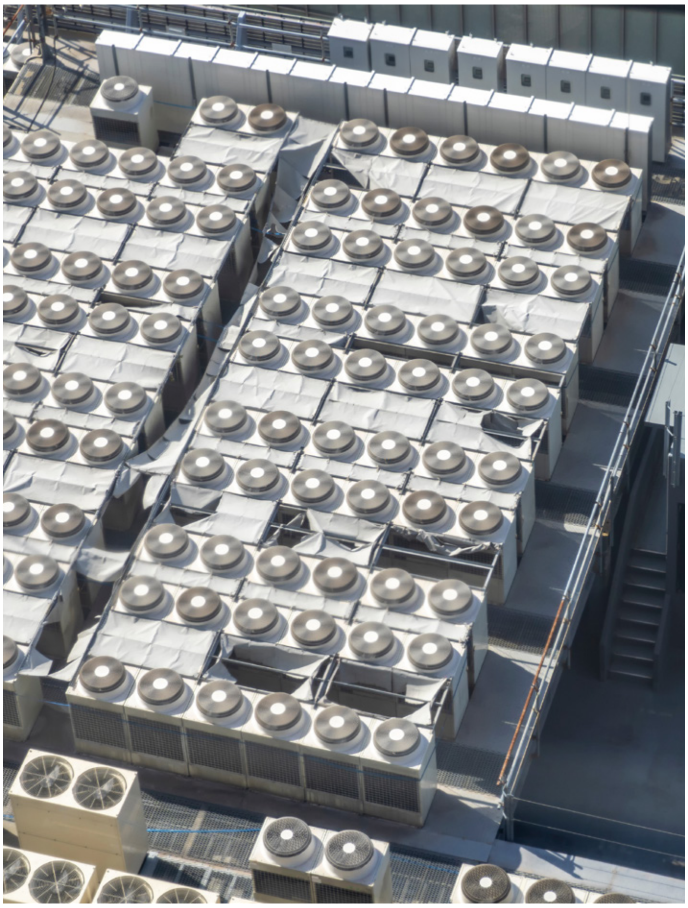

MATERIALS REUSE EFFICIENCY

Modular design and design for re-use plays an important role in design for circularity. Maximizing the use of components that are interchangeable and reusable between racks and servers allows DC operators to upgrade the chips to take advantage of advances in electronics, without requiring full replacement of server and stack hardware, and hence obviates the embodied carbon footprint of a full hardware refresh.

All refrigerants and refrigerant blends that are used in rack cooling, HVAC and chillers can be recovered at end of system life and reprocessed into either the pure refrigerant components or raw materials to make similar molecules, usually referred to as "reclaiming" the refrigerant. Materials that have potential PFAS concerns can be safely destroyed by incineration such that no PFAS material is released into the environment. For the purposes of lifecycle impact assessment of refrigerants in well-managed, large facilities the refrigerant loss and make-up rate is typically less than 5% per year, meaning that more than 50% of the total refrigerant charged (including make-up) is recovered at the end of a 20-year system life.

WATER USE EFFICIENCY

Once-through use of city water as coolant was never practical for large scale DCs and is problematic because of cost as well as the environmental impact of rejecting large volumes of warm water to the municipal water treatment system. Closed-loop cooling water systems also have a high environmental footprint. Cooling water systems achieve their cooling by evaporating part of the water into ambient air and therefore require a make-up stream of fresh water to compensate for that lost by evaporation. Any dissolved solids brought in with the fresh water will accumulate in the circulating water, and so a portion of the water is taken as a purge stream (known as "cooling water blowdown") to control the level of solids at a level that does not cause excessive fouling of heat transfer surfaces. Additional make-up water is needed to compensate for the blowdown. Cooling tower water is usually treated with a range of chemicals such as algaecides and biocides (to prevent biological fouling), corrosion inhibitors and scale inhibitors (to prevent mineral fouling) all of which are present in the cooling water blowdown. Cooling water blowdown therefore contains chemicals that can significantly impact the environment if not treated adequately.

The water footprint of data centers can be eliminated by moving away from water-based cooling systems. Refrigerant-based cooling systems (dual exchange air-cooled or liquid-cooled) avoid the use of cooling water and eliminate the resulting environmental impact on water distressed regions. Systems that cool ambient air also condense moisture from the air in areas with high humidity, allowing some DCs to even operate water negative, producing more fresh water than they consume.

Eliminating water discharges from the data center has the added advantage of eliminating nature impacts due to potential discharge of biocide compounds and other cooling water additive chemicals into river systems.

CONCLUSIONS AND PATH FORWARD

Our analysis suggests that the greatest opportunities for increasing the sustainability of data centers come from the following activities:

- Continuously upgrading IT hardware to take advantage of improvements in technology and remain close to state-of-the-art efficiency of electronic components.
- Increasing the use of digital control systems and automation to integrate data from the IT systems and OT systems as well as any co-located power generation, transmission and distribution equipment, enabling:
  - Development of analytical tools (deterministic, AI or hybrid) and control strategies that exploit the full set of data available in an integrated automation system to optimize energy consumption, asset utilization and power source C-intensity with increasingly high-time resolution to achieve lowest possible carbon footprint of instantaneous energy use without compromising system availability.
  - Deployment of a full range of automation and analytics tools to maximize reliability and uptime of assets and prevent outage conditions that can damage assets (requiring repairs that increase embodied C footprint) and lead to spikes in use of energy or increased use of high C-intensity energy from backup power systems such as generators.
  - Early recognition and remediation of compromised equipment that is running inefficiently and using more power and/or causing a greater power draw from other systems compensating for the compromised equipment. Proactive detection of declining asset health is important for resiliency as well as sustainability.
- Maximizing the supply of firm low C-intensity power either by choice of location, co-location with renewable power assets or firm power purchase agreements.
- Deploying battery energy storage systems to store variable renewable energy and enable firmness of supply, meet power backup requirements with lower C-intensity than fossil-fueled generators and exploit opportunities for daily price arbitrage while avoiding high C-intensity peak grid power.
- Replacing legacy high global warming refrigerants in CRAC and DX cooling systems with low global warming potential refrigerants to reduce the embodied carbon footprint (Scope 3 impact) of the data center.
- Using thermal energy storage systems to shift cooling loads away from times when refrigeration systems are inefficient (peak daily heat) or electricity prices are high (peak power hours) and thereby reduce the overall C-intensity of power consumed.
- Using heat pumps to boost the temperature of waste heat from the data center and allow energy reuse for district heating in nearby communities of other low-medium grade heat applications.
- Integrating data centers in more remote locations with direct air capture plants for removing carbon dioxide from the atmosphere for geological sequestration, using the data center waste heat to offset roughly 40% of the energy needed for DAC and achieving overall carbon-negative operation.

Honeywell expects to see continued innovation and improvements in all these areas and will continue to co-innovate with customers to accelerate development and demonstration of new technologies that help reduce the environmental impact of data centers of the future.

REFERENCES

Acton, M., Bertoldi, P., Booth, J. (2020) 2020 Best Practice Guidelines for the EU Code of Conduct on Data Center Energy Efficiency, European Commission, Ispra, JRC119571.

Bloom Energy (2025) From Gridlock to Growth: How Power Bottlenecks Are Reshaping Data Center Strategies, 2025 Data Center Power Report. Bloom Energy. https://www.bloomenergy.com/wp-content/uploads/2025-Data-Center-Power-Report.pdf (Accessed: March 6, 2025).

Darrow, K. and Hedman, B. (2009) Opportunities for combined heat and power in data centers. Prepared by ICF International under contract 4000021512 for Oak Ridge National Laboratory ORNL/TM-2009/052.

Donnellan, D. and Lawrence, A. (2024) Annual outage analysis 2024, Uptime Intelligence. Uptime Intelligence. https://datacenter.uptimeinstitute.com/rs/711-RIA-145/images/2024.Resiliency.Survey.ExecSum.pdf?version=0&mkt_tok=NzExLVJJQS0xNDUAAAGSPCeKfdv0kYTrLS-6 (Accessed: March 6, 2025).

Du, K., Calautit, J., Wang, Z., Wu, Y., and Liu, H. (2018), A review of the applications of phase change materials in cooling, heating and power generation in different temperature ranges, Appl. Energy 220, 242–273.

EIA (2021a) "Electric Power Monthly, Capacity Factors for Utility Scale Generators Primarily Using Non-Fossil Fuels" https://www.eia.gov/electricity/monthly/epm_table_grapher.php?t=epmt_6_07_b

Energy Institute (2023) Statistical review of world energy. (Energy Institute). Available at https://www.energyinst.org/statistical-review/about

Engel-Cox, J. (2020) "Renewable Energy for Resilience & Adaptation: Workshop on Adaptive Actions Toward a Sustainable and Resilient Future" 10 January 2020, NREL/PR-6A50-75738. (National Renewable Energy Laboratory)

Farman, A. et al., (2024) Data Centers & AI: Powering the Future, Jefferies. Thematic Research. https://javatar.bluematrix.com/pdf/uHPRqtoR (Accessed: March 6, 2025).

Hirschey, J.R., Kumar, N., Turnaoglu, T., Gluesenkamp, K.R. and Graham, S. (2021) "Review of Low-Cost Organic and Inorganic Phase Change Materials with Phase Change Temperature between 0°C and 65°C," International High Performance Buildings Conference. Paper 389. https://docs.lib.purdue.edu/ihpbc/389.

Kontorinis, V. et al., (2012) "Managing Distributed UPS Energy for Effective Power Capping in Data Centers," Proc. Int'l Symp. Computer Architecture, pp. 488–499.

Krishnadas, G. and Kiprakis, A. (2020) A machine learning pipeline for demand response capacity scheduling. Energies, 13(7), 1848-1873.

Luo, Y., Andresen, J., Clarke, H., Rajendra, M., and Maroto-Valer, M. (2019), A decision support system for waste heat recovery and energy efficiency improvement in data centers. Appl. Energy, 250, 1217-1224.

Memarzadeh, F., Wheeland, D.G. and Ramesh, A. (2013) "National Institutes of Health Sustainable Data Center Design Guide" (National Institutes of Health Division of Technical Resources).

NREL, 2021. Annual Technology Baseline. https://atb.nrel.gov/electricity/2021/utility-scale_pv (National Renewable Energy Laboratory).

Poudineh, R. (2025) Global electricity demand: what's driving growth and why it matters?, Oxford Energy. The Oxford Institute of Energy Studies. https://www.oxfordenergy.org/wpcms/wp-content/uploads/2025/01/Global-electricity-demand.pdf.

Ragoowansi et al. (2023), "Realistic utilization of emerging thermal energy recovery and storage technologies for buildings," Cell Reports Physical Science, volume 4, issue 5. https://www.sciencedirect.com/science/article/pii/S2666386423001613 (digital image accessed 1 February 2024, https://ars.els-cdn.com/content/image/1-s2.0-S2666386423001613-gr2_lrg.jpg)

Reply, 2022. "Making Cloud Computing More Sustainable: Green Cloud", Available at Green web design (whitepaper) (reply.com) [data credited to IEA].

Rozite, V. (2023) Data Centres and Data Transmission Networks. https://www.iea.org/energy-system/buildings/data-centres-and-data-transmission-networks (Accessed: March 6, 2025).

Siddik, M.A.B, Shehabi, A. and Marston, L. (2021) The environmental footprint of data centers in the United States. Environ. Res. Lett. 16, 1-11.

Statista (2025) "Volume of data/information created, captured, copied, and consumed worldwide from 2010 to 2023, with forecasts from 2024 to 2028", available at https://www.statista.com/statistics/871513/worldwide-data-created/

Synergy Research Group (2024) Hyperscale Data Centers Hit the Thousand Mark; Total Capacity is Doubling Every Four Years | Synergy Research Group. https://www.srgresearch.com/articles/hyperscale-data-centers-hit-the-thousand-mark-total-capacity-is-doubling-every-four-years (Accessed: March 6, 2025).

Towler, G.P. and Sinnott, R.K (2022) Chemical Engineering Design: Principles Practice and Economic of Plant and Process Design, (Elsevier).

Trane (2023) Trane Technologies White Paper "Thermal Energy Storage: More Affordable Than Ever", https://elibrary.tranetechnologies.com/public/commercial-hvac/Literature/Sales/IB-SLB003-EN_10182023.pdf

Van Greet, O. (2010) "FEMP best practices guide for energy-efficient data center design". NREL report NREL/BR-7A40-47201.

APPENDIX 1: GLOSSARY

Anthropogenic greenhouse gas emissions: emissions of greenhouse gases due to human activity excluding natural sources but not excluding agriculture and land use impacts.

Assessment Report (AR): periodic reports issued by the Intergovernmental Panel for Climate Change (IPCC) that summarize the consensus state of scientific opinion on the extent, impact and potential mitigation of global warming.

Carbon capture and storage (CCS): collection of carbon dioxide from any source and permanent sequestration of the carbon dioxide in geological storage so that it does not enter the atmosphere.

Carbon capture, utilization and storage (CCUS): collection of carbon dioxide from any source followed either by geological sequestration (CCS) or conversion of the carbon dioxide into durable materials that are not subsequently combusted with re-release of the carbon dioxide to the atmosphere.

Carbon dioxide equivalent (CO2e): the equivalent amount of carbon dioxide that would cause the same global warming impact. This is a measure used to report other GHG emissions on a carbon dioxide equivalent basis and allows for the fact that other GHGs can have stronger warming effects or be more persistent in the atmosphere.

Carbon footprint: shorthand term used for carbon dioxide emissions footprint (more strictly GHG emissions footprint) – the carbon dioxide emissions associated with a given activity.

Carbon intensity of energy (C-intensity): shorthand for carbon dioxide intensity (or more strictly GHG intensity) of energy. The amount of CO2 (strictly CO2e, including actual carbon dioxide as well as other GHG on a carbon dioxide equivalent basis) emitted per unit energy consumed.

Carbon-negative technology: strictly, GHG emissions negative technology. Applies to any technology that permanently removes more GHG from the atmosphere than the entire carbon footprint associated with installation, operation and decommissioning of the technology over the entire service life of the technology.

Carbon neutral: widely used but imprecise term, strictly meaning carbon dioxide emissions neutral. Since all activities that consume energy or materials have some emissions impact, the term carbon neutral strictly applies only to systems that have offset all their GHG emissions footprint with an equivalent amount of permanent carbon dioxide sequestration from the atmosphere.

Clean hydrogen: defined in the U.S. Federal Infrastructure bill and Clean Hydrogen Production Incentives Act of 2021 (S.1017) as "hydrogen produced with a carbon intensity equal to or less than 2 kilograms of carbon dioxide-equivalent produced at the site of production per kilogram of hydrogen produced." Note that steam methane reforming typically produces about 7 kg CO2 per kg H2, so the U.S. definition of clean hydrogen requires at least 72% carbon capture and sequestration if applied to conventional hydrogen production.

Computer rack air conditioner (CRAC): a dedicated air conditioning unit serving a rack of servers.

Computer rack air handler (CRAH): a dedicated air circulation system serving a rack of servers.

Decarbonization: strictly, "removal of carbon from." Generally used in the context of decarbonization of the energy supply. Note that it is correct to say "decarbonization of the energy used for light duty transportation," implying the continued use of light duty transportation with energy sources that do not contain carbon, but it is incorrect to say "decarbonization of gasoline" as gasoline intrinsically contains carbon. Note also that decarbonization describes any level of removal of carbon. We therefore use the term "full decarbonization" to describe the complete removal of carbon from a particular energy supply.

Direct air capture (DAC): strictly, direct air capture of carbon dioxide. CCS or CCUS applied to carbon dioxide that is already in the atmosphere, thereby actually reducing the atmospheric concentration of carbon dioxide.

Energy efficiency: the proportion of energy consumed that is converted into useful mechanical work or required heat as opposed to waste heat or other non-usable forms of energy.

Greenhouse effect: global warming caused by the accumulation of anthropogenic greenhouse gas emissions in the atmosphere.

Greenhouse gases (GHG): gas species such as carbon dioxide, methane, nitrogen oxides and some fluorinated gases that absorb infrared radiation and consequently reduce the ability of the earth to cool itself by radiation to outer space.

Low-carbon energy: strictly "lower carbon energy." Energy sources that have reduced GHG emissions when compared to conventional energy sources used in the same application.

Net-zero emissions: strictly, net-zero GHG emissions. Somewhat stricter than carbon neutral, a net-zero GHG condition applies to a system that has offset all GHG emissions with an equal amount of carbon dioxide sequestration from the atmosphere.

Renewable energy: energy sources that are replenished by solar power or heat from the earth's core over non-geological timescales. This term can be used for wind power, wave power, solar power, hydroelectric power, geothermal power, ocean thermal power and energy from biomass sources that are grown sustainably.

Renewable distillate fuel: a distillate range fuel (kerosene, jet or diesel) derived from sustainable biomass sources.

Zero-emissions process: strictly, a technology that captures and sequesters an amount of GHG emissions sufficient to offset all the GHG emissions associated with installation, operation and decommissioning of the technology.

APPENDIX 2: LIFECYCLE ANALYSIS

Details of the LCA cases run and breakdowns of the carbon footprints calculated are given below:

CASE A: CONVENTIONAL FULLY DECOUPLED

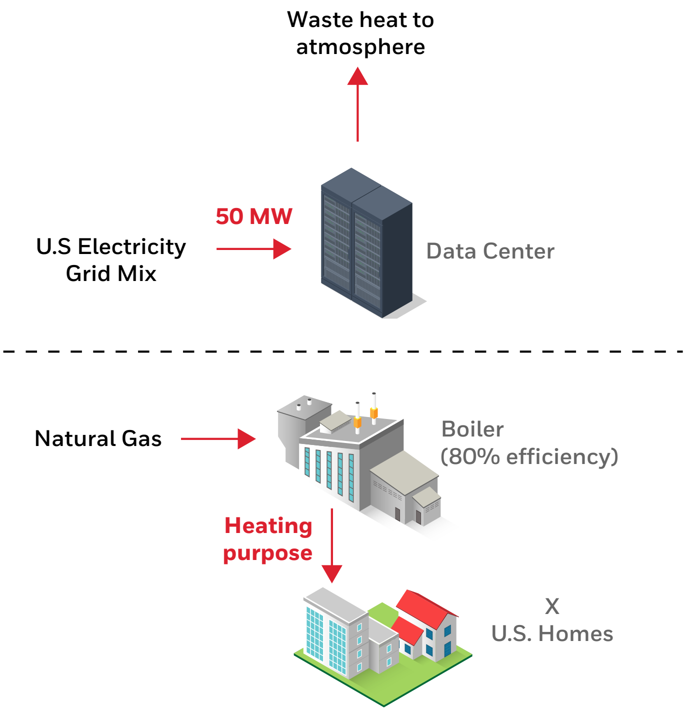

50 MW power draw = **4.38x10⁸ kWh/ yr** [8760 h - basis]

EF (US Elec. Grid mix) = **0.528 kg CO_2 eq./ kWh**

Heat consumption of U.S. Home = **45 MMBtu/ yr**

NG input to Boiler (@ 80% eff.) = **54 MMBtu/ yr = 15825.84 kWh/ yr**

EF (US NG district heating) = **0.133 kg CO2 eq./ kWh**

CFP of heating 1 US Home = **2.10 t CO2 eq./ year**

X = 25,370 homes (from Case-D)

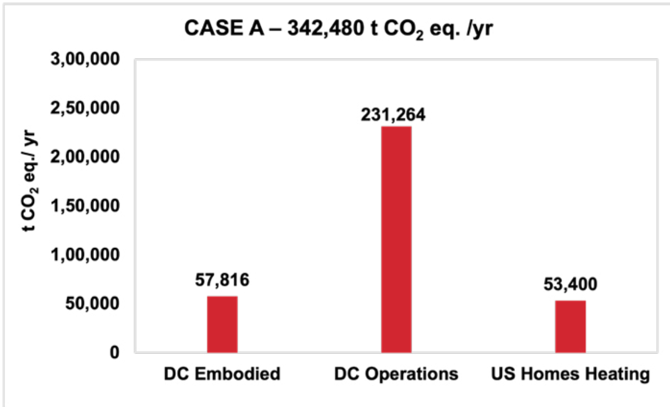

CASE B: CONVENTIONAL WITH HEAT AND POWER INTEGRATION

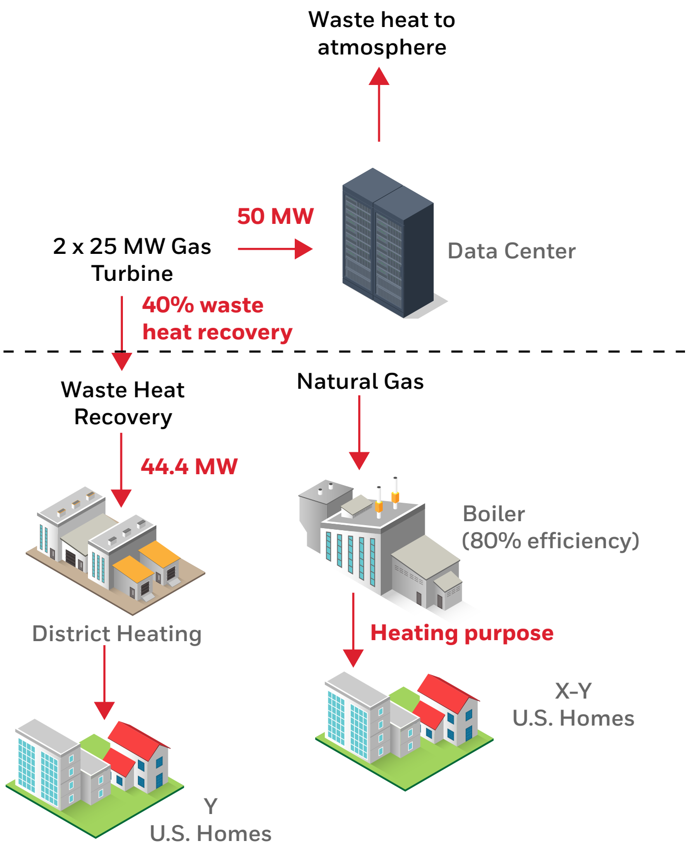

50 MW power draw = **111.11 MW input power (45% GT efficiency)**

111.11 MW input power = **9.73x10⁸ kWh/ yr** [8760 h - basis]

EF (US, NG in GT) = **0.291 kg CO_2 eq./ kWh**

40% heat recovery (111.11 MW) = **44.4 MW useful waste heat**

District heating to 'Y' homes = **1.94x10⁸ kWh/ half year** [8760/2 h - basis]

'Y' = 12,301 homes; 'X' = **25,370 homes**

'X-Y' homes(NG @ 80% efficiency) = **15825.84 kWh/ yr**

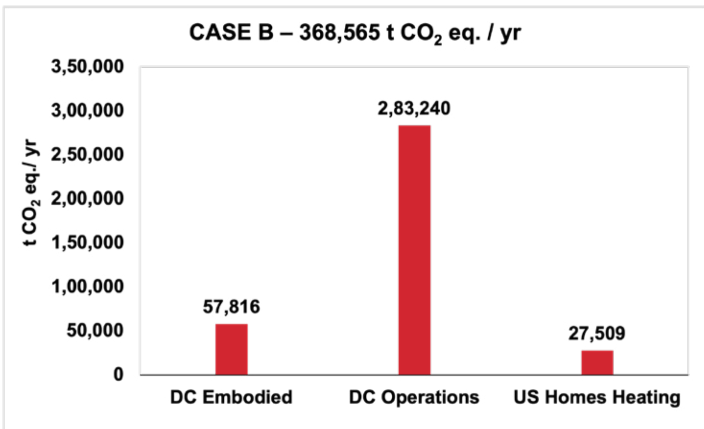

CASE C (SOLAR): DEDICATED RENEWABLES FULLY DECOUPLED

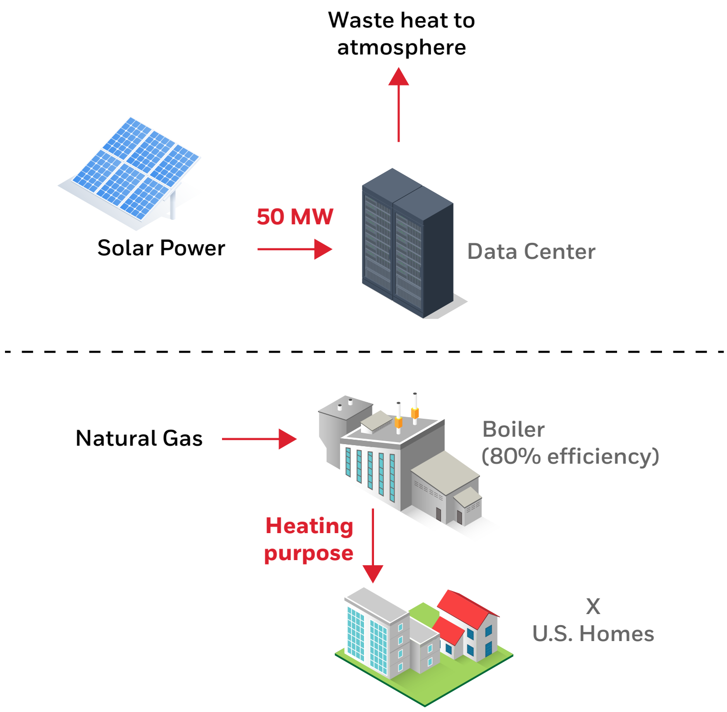

50 MW power draw = **187 MW installed cap**. (26.8% attainment factor)

187 MW capacity = **1.63x109 kWh/ yr** [8760 h - basis]

EF (Solar plant with infrastructure) = **0.0692 kg CO_2 eq./ kWh**

'X' Homes = **25,370 homes** heating with NG @ 80% efficiency

Input heat per home (NG @ 80% efficiency) = **15825.84 kWh/ yr**

EF (US NG district heating) = **0.133 kg CO2 eq./ kWh**

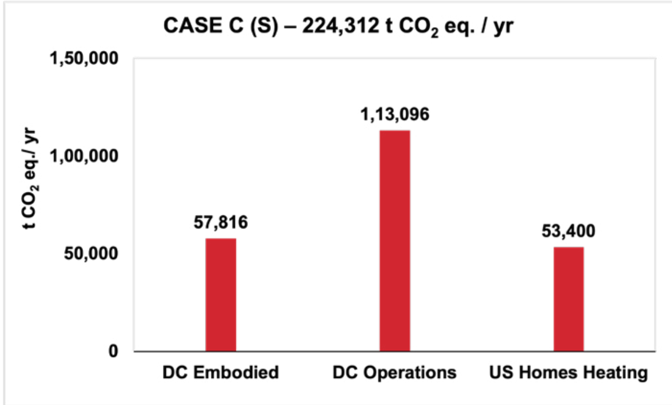

CASE C (WIND): DEDICATED RENEWABLES FULLY DECOUPLED

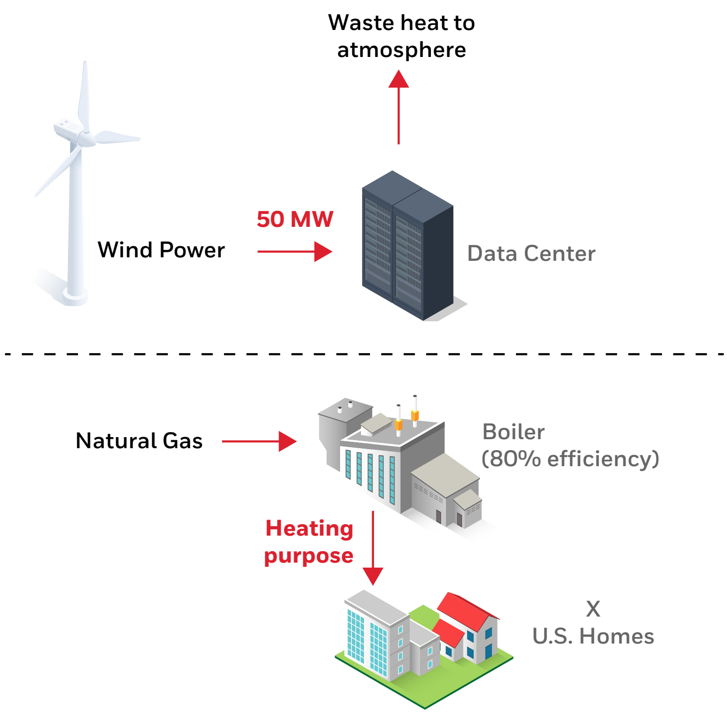

50 MW power draw = **117 MW installed cap.** (42.6% attainment factor)

117 MW capacity = **$1.03x10^9$ kWh/ yr** [8760 h - basis]

EF (Wind plant with infrastructure) = **0.0255 kg CO_2 eq./ kWh**

'X' Homes = **25,370 homes** heating with NG @ 80% efficiency

Input heat per home (NG @ 80% efficiency) = **15825.84 kWh/ yr**

EF (US NG district heating) = **0.133 kg CO_2 eq./ kWh**

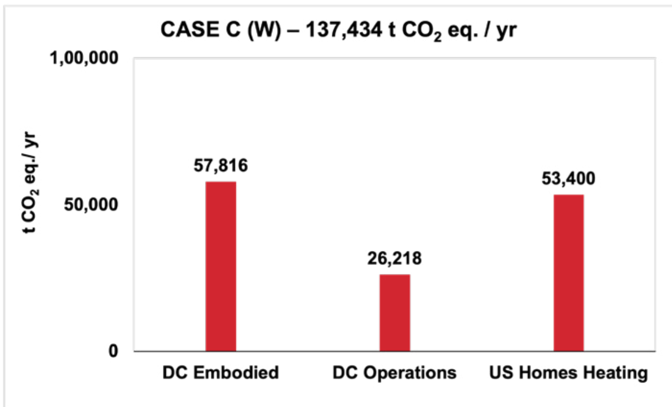

CASE D (SOLAR): DEDICATED RENEWABLES WITH HEAT AND POWER INTEGRATION

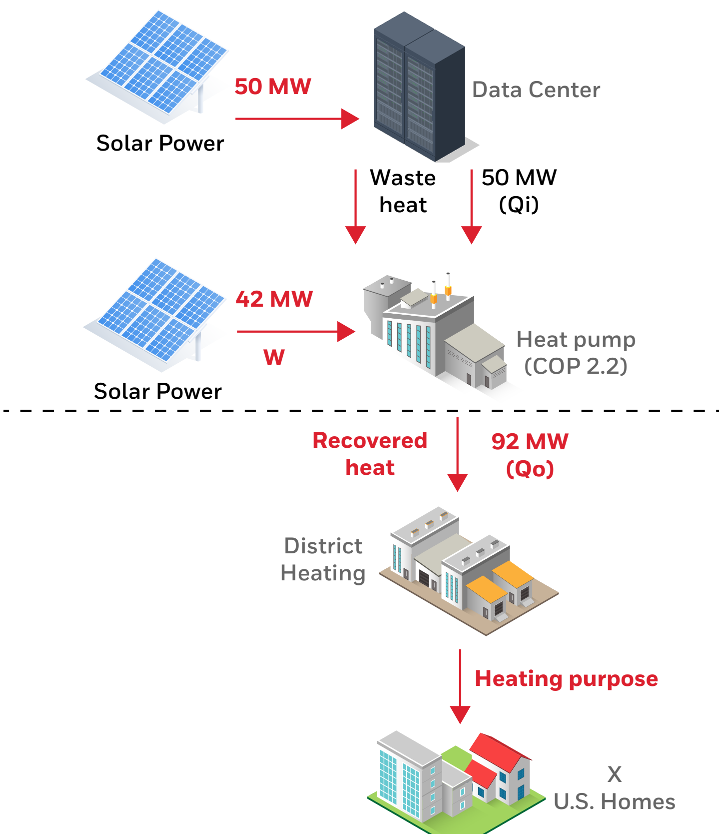

COP = Qo/W = (Qi+ W)/W

50 MW DC with heat pumps (COP 2.2) = **92 MW power draw** requirement

92 MW power draw = **342 MW installed cap.** (26.8% attainment factor)

342 MW capacity = **$2.99x10^9$ kWh/ yr** [8760 h - basis]

EF (Solar Plant with infrastructure) = **0.0692 kg CO_2 eq./ kWh**

92 MW recovered heat to district heating **'X' homes = $4.01x10^8$ kWh/ half yr**

'X' = **25,370 homes**

HP Embodied – calc using scale up factor of 0.85 (30kW HP in SimaPro)

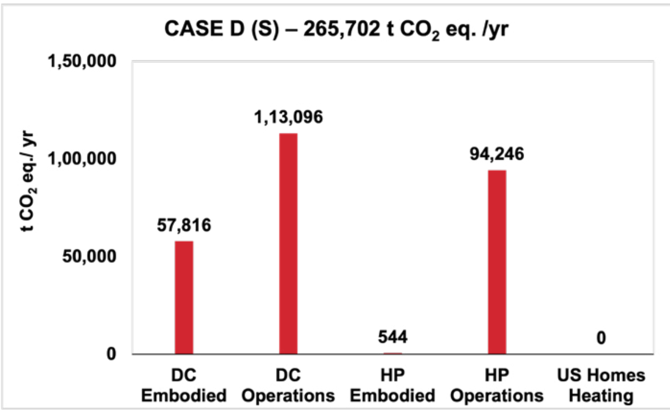

CASE D (WIND): DEDICATED RENEWABLES WITH HEAT AND POWER INTEGRATION

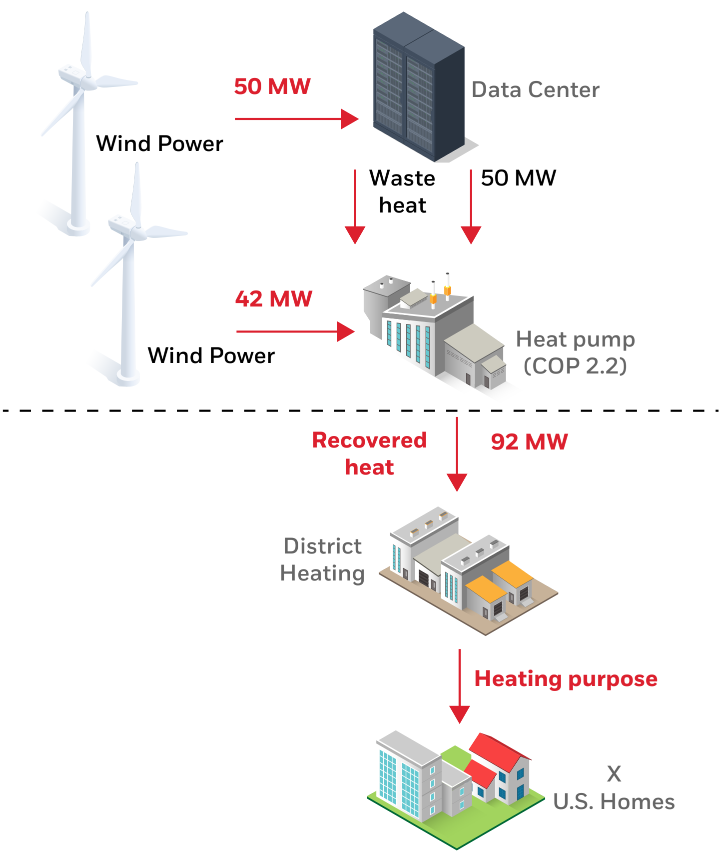

50 MW DC with heat pumps (COP 2.2) = **92 MW power draw** requirement

92 MW power draw = **215 MW installed cap.** (42.6% attainment factor)

215 MW input power = **$1.88x10^9$ kWh/ yr** [8760 h - basis]

EF (Wind Plant with infrastructure) = **0.0255 kg CO2 eq./ kWh**

92 MW recovered heat to district heating **'X' homes = 4.01x108 kWh/ half year**

'X' = **25,370 homes**

HP Embodied – calc using scale up factor of 0.85 (30kW HP in SimaPro)

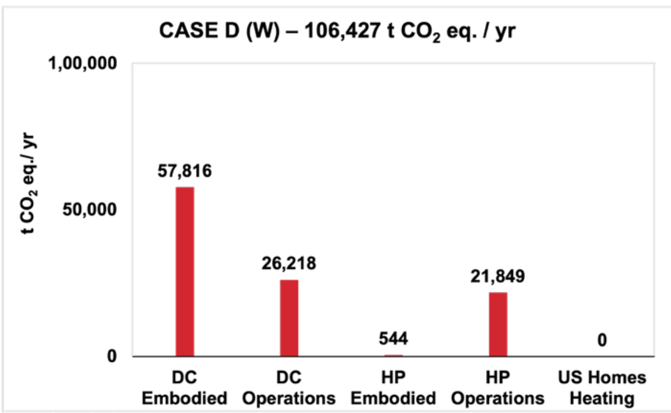

CASE E (SOLAR): DEDICATED RENEWABLES COUPLED TO DAC

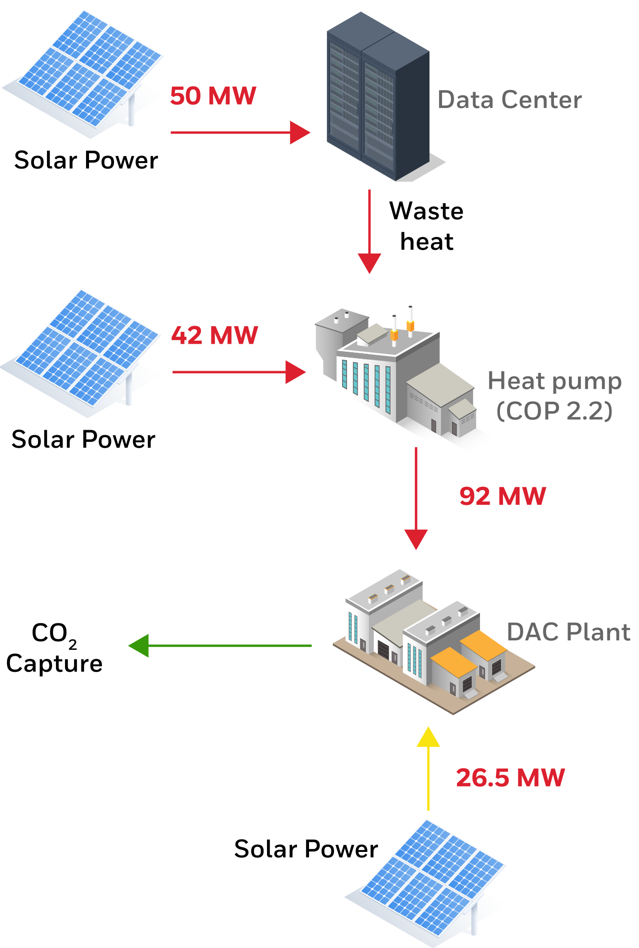

50 MW DC with heat pumps (COP 2.2) = **92 MW power draw**

92 MW power draw = **342 MW installed cap.** (26.8% attainment factor)

342 MW capacity = **2.99x10^9 kWh/ yr** [8760 h - basis]

EF (Solar Plant with infrastructure) = **0.0692 kg CO_2 eq./ kWh**

92 MW recovered heat to **DAC for CO_2 capture = 8.03x10^8 kWh/ yr**

Req: 1500 kWh heat / t CO_2 captured; **434 kWh of electric power** / t CO_2

CO2 capture in DAC @ 102 MW heat and 95oC = **535,333 t CO_2**

DAC solar req. = 8.67x10^8 kWh/ yr

DAC Embodied = **19 kg CO_2eq./ t CO_2**; Adsorbent = **17 kg CO_2eq./ t CO_2**

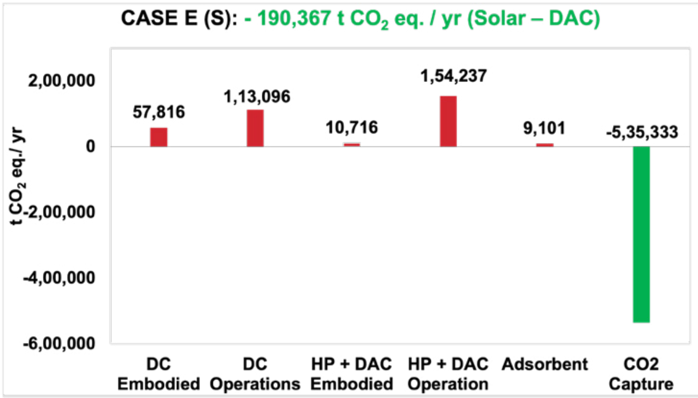

CASE E (WIND): DEDICATED RENEWABLES COUPLED TO DAC

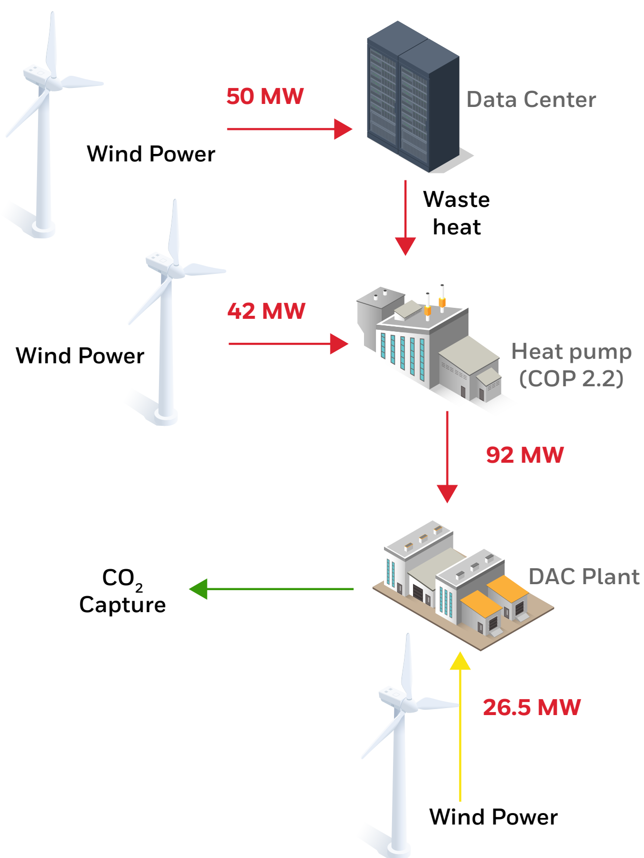

50 MW DC with heat pumps (COP 2.2) = **92 MW power draw**

92 MW power draw = **215 MW input power** (42.6% attainment factor)

215 MW input power = **1.88x10^9 kWh/ yr** [8760 h - basis]

EF (Wind Plant with infrastructure) = **0.0255 kg CO_2 eq./ kWh**

92 MW recovered heat to **DAC for CO_2 capture = 8.03x10^8 kWh/ yr**

Req: **1500 kWh heat** / ton CO_2 captured; 434 kWh of electric power / ton CO_2

CO_2 capture in DAC @ 102 MW heat and 95oC = **535,333 t CO_2**

DAC wind power requirement = **5.45x10^8 kWh/ yr**

DAC Embodied = 19 kg CO_2eq./ t CO_2; Adsorbent = **17 kg CO_2eq./ t CO_2**

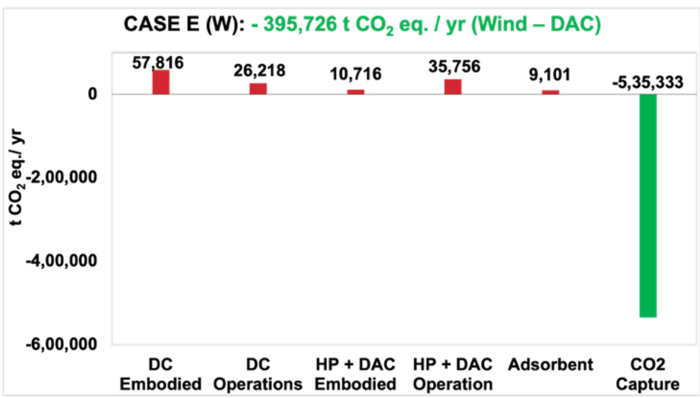

Key Assumptions

- We did not build a detailed model of home heating requirements, which vary hourly, daily and monthly. For a first-pass analysis, we assume the annual heating requirement is spread evenly over six months of the year. A more accurate analysis would allow for local temporal variation in heat demand.
- In the heat pump to district heating cases (case D) we assume that the heat is only required six months of the year. In the remaining six months we assume the heat pump is used to reject heat to atmosphere with the same coefficient of performance (CoP). In practice, depending on the location and climate the CoP could improve dramatically when delivering heat at lower temperature and the heat pump duty would be correspondingly reduced in the non-heating season. A more detailed monthly heat demand/rejection model would take account of this and show lower carbon footprints for case D.
- Wind and solar power are intermittent sources of energy and only deliver power when the wind is blowing or the sun is shining. The capacity factor (sometimes referred to as attainment) is the ratio of the average power produced to the nominal power rating. For newly installed wind turbines in 2020, the average capacity factor was 42.6% (Engel-Cox, 2020). For solar power (class 5 resources) the average capacity factor in 2020 was 26.8% (NREL, 2021; EIA, 2021a). Note that the capacity factor of wind power is higher than that of solar power because the sun does not shine at night, effectively limiting solar power to a maximum of 50% capacity factor.
- A co-located variable renewable energy power source such as wind or solar would require an energy storage (ES) system such as battery energy storage (BESS) to provide firm power during periods when the renewable resource is not available. The renewable power source would also need to be oversized to allow for the round-trip efficiency of the fraction of the power that was drawn from energy storage. We did not include ES GHG footprint in the LCA, as there is a wide range of variation in ES GHG footprint depending on the type of energy storage selected and we assume renewable power can be provided by a firm PPA. This can be explored in future work.
- The embodied carbon footprint of the heat pump assumed a service life of 10 years.
- The embodied carbon footprint of the DAC plant assumed a service life of 20 years.
- The embodied carbon footprint of the adsorbent used in the DAC plant assumes a service life of 1 year. This is probably conservative, but since there are no full-scale DAC plants in operation yet we made a safe-side assumption.

For more information

www.honeywell.com

Honeywell

855 S Mint St Charlotte, NC,
28202-1517 USA
www.honeywell.com

WPR-25-04-EN | 03/25

THE
FUTURE
IS
WHAT
WE
MAKE IT

Honeywell