



Best Practices for Energy-Efficient Data Center Design

Four (4) Continuing Education Hours
Course #EE1645

Approved Continuing Education for Licensed Professional Engineers

EZ-pdh.com
Ezekiel Enterprises, LLC
301 Mission Dr. Unit 571
New Smyrna Beach, FL 32170
800-433-1487
support@ez-pdh.com



Course Description:

The Best Practices for Energy-Efficient Data Center Design course satisfies four (4) hours of professional development.

The course is designed as a distance learning course that provides a foundational guide to designing sustainable data centers with optimal energy performance. It covers IT hardware efficiency, air and cooling system design, benchmarking metrics, and methods for improving energy reuse, water conservation, and carbon reduction.

Course is based from Best Practices Guide for Energy-Efficient Data Center Design published by the Department of Energy, National Renewable Energy Laboratory (NREL).

Objectives:

The primary objective of the course is to equip engineers with knowledge of best practices and design strategies to optimize energy use and sustainability in modern data centers.

Grading:

Students must achieve a minimum score of 70% on the online quiz to pass this course. The quiz may be taken as many times as necessary to successfully pass and complete the course.

A copy of the quiz questions is attached to the last pages of this document.

Table of Contents

Best Practices for Energy-Efficient Data Center Design

- Overview..... 1**
- 1 Background 2**
 - 1.1 Key Steps to Sustainable Data Centers 2
- 2 Information Technology Systems 3**
 - 2.1 Cloud and Colocation Computing/Storage 3
 - 2.2 Efficient Servers 3
 - 2.3 Storage Devices..... 4
 - 2.4 Network Equipment 4
 - 2.5 Power Supplies 5
 - 2.6 Consolidation..... 6
- 3 Environmental Conditions 7**
 - 3.1 2021 ASHRAE Equipment Thermal Guidelines for Data Processing Environments 7
- 4 Air Management..... 10**
 - 4.1 Implement Cable Management..... 10
 - 4.2 Aisle Separation and Containment 10
 - 4.3 Optimize Supply and Return Air Configuration 13
 - 4.4 Other Benefits of Raising Temperature Set Points..... 14
- 5 Cooling Systems 15**
 - 5.1 Direct Expansion Systems 15
 - 5.2 Air Handlers..... 16
 - 5.3 High-Efficiency Chilled Water Systems..... 17
 - 5.4 Free Cooling..... 19
 - 5.5 Thermal Storage..... 20

5.6 Direct Liquid Cooling	20
5.7 Humidification	23
5.8 Controls.....	23
6 Electrical Systems	25
6.1 Power Distribution	25
6.2 Demand Response	27
6.3 DC Power	27
6.4 Lighting	28
7 Use of Waste Heat for Energy-Efficient Design	29
7.1 Use of Waste Heat.....	29
8 Benchmarking.....	30
8.1 Power Usage Effectiveness (PUE)	30
8.2 Energy Reuse Effectiveness (ERE)	31
8.3 Water Use Effectiveness (WUE).....	31
8.4 Carbon Use Effectiveness (CUE)	31
8.5 Related Metrics.....	32
8.6 Airflow Efficiency	32
8.7 Cooling System Efficiency	33
8.8 ISO/IEC 30134	33
8.9 Data Center Infrastructure Management and Monitoring.....	33
9 Conclusion	35
9.1 Key Steps to Sustainable Data Centers	35
9.2 Federal Energy Management Program.....	35
Quiz Questions	37

Overview

This course provides an overview of best practices for energy-efficient data center design which spans the categories of information technology (IT) systems and their environmental conditions, data center air management, cooling and electrical systems, and heat recovery. IT system energy efficiency and environmental conditions are presented first because measures taken in these areas have a cascading effect of secondary energy savings for the mechanical and electrical systems.

This course concludes with a section on metrics and benchmarking values by which a data center and its systems energy efficiency can be evaluated. No design course can offer “the most energy-efficient” data center design but the guidelines that follow offer suggestions that provide efficiency benefits for a wide variety of data center scenarios.

1 Background

Data center spaces can consume many times as much electricity as standard office spaces. With such large power consumption, they are prime targets for energy-efficient design measures that can save money and reduce electricity use. However, the critical nature of data center loads elevates many design criteria—chiefly reliability and high-power density capacity—far above energy efficiency. Short design cycles often leave little time to fully assess efficient design opportunities or consider first cost versus life-cycle-cost issues. This can lead to designs that are simply scaled up versions of standard office space approaches or that reuse strategies and specifications that worked “good enough” in the past without regard for energy performance. This course has been created to provide viable alternatives to inefficient data center building practices.

An effective organization will consider the total cost of ownership for operational efficiency and cost, utilizing different energy-based metrics and sustainability metrics (water and carbon) to capture a view of the efficiencies at which a data center performs.

1.1 Key Steps to Sustainable Data Centers

The U.S. Department of Energy’s Federal Energy Management Program (FEMP) and the National Renewable Energy Laboratory (NREL) developed the following approach for optimizing data center sustainability, listed in order of importance:

1. Reduce energy use by making systems as efficient as possible – the associated data center metric is Power Usage Effectiveness (PUE).
 - o **Maximize compute entering temperature to maximize energy efficiency while ensuring information technology (IT) equipment thermal guidelines are met to avoid overheating or compromising reliability.**
 - o Use “free” cooling to reduce or eliminate compressor-based cooling (chiller, direct expansion [DX]).
 - o Optimize fan/pump speeds and uninterruptible power supplies.
2. Reuse heat to achieve the lowest Energy Reuse Effectiveness (ERE) metric possible.
 - o **Maximize compute leaving temperature to maximize energy reuse.**
3. Reject as much remaining heat to dry coolers as possible – water savings reflected in data center metric Water Usage Effectiveness (WUE).
 - o **Maximize compute leaving temperature to maximize heat rejected dry to air.**
4. Maximize energy from renewable systems on-site or within grid region – tracked through data center metric Carbon Usage Effectiveness (CUE).
 - o **Work toward the goal of 100% renewable energy 100% of the time.**

2 Information Technology Systems

In a typical data center with a highly efficient cooling system, IT equipment loads can account for over half of the entire facility's energy use. Use of efficient IT equipment will significantly reduce these loads within the data center, which consequently will downsize the equipment needed to cool them. Purchasing servers equipped with energy-efficient processors, fans, power supplies, and high-efficient network equipment; consolidating storage devices; consolidating power supplies; and implementing virtualization are the most advantageous ways to reduce IT equipment loads within a data center.

Efficient algorithms can have big impact on energy use especially in the artificial intelligence/machine learning fields. However, that is outside the scope of this document (that focuses on hardware), but worth exploring depending on application space.

2.1 Cloud and Colocation Computing/Storage

Building and operating an on-premises data center is expensive and requires expert staff. An on-premises data center has finite capacity, must be provided with reliable power and communications, and must provide adequate cybersecurity. If an on-premises data center fails, business operations may be impacted unless a back-up data center, sometimes called a fail-over data center, is available, which adds cost and complexity.

Cloud and colocation computing/storage has lower first cost and may have lower operational cost than on-premises data centers. A cloud data center has potentially unlimited capacity. The cloud vendor is responsible for all operations including infrastructure and cybersecurity. It is essentially computing as a service. Colocation facilities, on the other hand, provide rental of space, power, and cooling to customers along with network service to connect customer owned and managed IT systems. Potential advantages of cloud and colocation services are especially significant compared to small on-premises data centers that may lack around-the-clock expert staff, adequate redundancy, and efficient design.

Depending on mission need, organizations may choose to only have an on-premises data center, or take a "cloud first" strategy where most or all data needs are provided by the cloud, or a hybrid solution where extremely critical data operations are conducted in an on-premises data center but other operations, such as data storage, are cloud provided. Determining what is best for a specific organization is outside the scope of this document.

2.2 Efficient Servers

Rack servers tend to be the main perpetrators of wasting energy and represent the largest portion of the IT energy load in a typical data center. Servers take up most of the space and drive the entire operation. The average server utilization (average to maximum activity) is generally in the range of 20% to 40% in enterprise settings. Server efficiency increases by about 50% when processor utilization is doubled from low levels of 20% to 30% (Rahkonen and Dietrich 2023). Server efficiency is defined as transactions per second (work) per watt. A 50% increase in this

server efficiency means that a server can process 50% more IT workload without an increase in energy. Server efficiency continues to improve, for example server efficiency approximately doubled from 2017 to 2021. Essentially all servers on the market today have variable speed fans. With variable speed fans it is possible to deliver sufficient cooling while running slower, thus consuming less energy. The ENERGY STAR® program aids consumers by recognizing high-efficiency servers. Servers that meet ENERGY STAR efficiency requirements will, on average, be 30% more efficient than standard servers (Rahkonen and Dietrich 2023). Furthermore, purchasing ENERGY STAR servers is required for federal data centers.

Multi-core processor chips allow simultaneous processing of multiple tasks, which leads to higher efficiency in two ways. First, they offer improved performance within the same power and cooling load as compared to single-core processors. Second, they consolidate shared devices over a single processor core. Not all applications are capable of taking advantage of multi-core processors. Graphics-intensive programs and high-performance computing still require higher clock-speed single-core designs.

2.3 Storage Devices

Power consumption is roughly linear to the number of storage modules used. Storage redundancy needs to be rationalized and right-sized to avoid rapid scale up in size and power consumption. Cloud storage can provide significant benefits over on-premises storage.

Consolidating storage drives into a Network Attached Storage or Storage Area Network are two options that take the data that does not need to be readily accessed and transports it offline. Taking superfluous data offline reduces the amount of data in the production environment, as well as all the copies. Consequently, less storage and central processing unit (CPU) requirements on the servers are needed, which directly corresponds to lower cooling and power needs in the data center.

For data that cannot be taken offline, it is recommended to upgrade from traditional storage methods to thin provisioning. In traditional storage systems an application is allotted a fixed amount of anticipated storage capacity, which often results in poor utilization rates and wasted energy. Thin provisioning technology, in contrast, is a method of maximizing storage capacity utilization by drawing from a common pool of purchased shared storage on an as-needed basis, under the assumption that not all users of the storage pool will need the entire space simultaneously. This also allows for extra physical capacity to be installed at a later date as the data approaches the capacity threshold.

2.4 Network Equipment

As newer generations of network equipment pack more throughput per unit of power, there are active energy management measures that can also be applied to reduce energy usage as network demand varies. Such measures include idle state logic, gate count optimization, memory access algorithms and input/output buffer reduction.

As peak data transmission rates continue to increase—requiring dramatically more power—increasing energy is required to transmit small amounts of data over time. Ethernet network energy efficiency can be substantially improved by quickly switching the speed of the network links to the amount of data that is currently transmitted.

2.5 Power Supplies

Most data center equipment uses internal or rack mounted alternating current/direct current (AC-DC) power supplies. Historically, a typical rack server's power supply converted AC power to DC power at efficiencies of around 60% to 70%. Today, through the use of higher-quality components and advanced engineering, it is possible to find power supplies with efficiencies up to 92% by choosing 80 PLUS Titanium power supplies (www.80plus.org). Using higher efficiency power supplies will directly lower a data center's power bills and indirectly reduce cooling system cost and rack overheating issues. At \$0.12/kilowatt-hour (kWh), savings of \$2,000 to \$6,000 per year per rack (10 kW to 25 kW, respectively) are possible just from improving the power supply efficiency from 75% to 85%. These savings estimates include estimated secondary savings due to lower uninterruptible power supply (UPS) and cooling system loads.

The impact of real operating loads should also be considered to select power supplies that offer the best efficiency at the load level at which they are expected to most frequently operate. The optimal power supply load level is typically in the mid-range of its performance curve: around 40% to 60%, as shown in Figure 2-1.




Figure 2-1. Efficiencies at varying load levels for typical power supplies

Source: *Quantitative Efficiency Analysis of Power Distribution Configurations for Data Centers*, The Green Grid

Efficient power supplies usually have a minimal incremental cost at the server level and so should be selected. There are also several certification programs currently in place that have standardized the efficiencies of power supplies in order for vendors to market their product.

2.6 Consolidation

2.6.1 Hardware Location

Lower data center supply fan power and more efficient cooling system performance can be achieved when equipment with similar heat load densities and temperature requirements are grouped together. Isolating equipment by environmental requirements of temperature and humidity allow cooling systems to be controlled to the least energy-intensive set points for each location.

This concept can be expanded to data facilities in general. Consolidating underutilized data center spaces to a centralized location can ease the utilization of data center efficiency measures by condensing the implementation to one location, rather than several.

2.6.2 Virtualization

Virtualization is a method of running multiple independent virtual operating systems on a single physical computer. It is a way of allowing the same amount of processing to occur on fewer servers by increasing server utilization. Instead of operating many servers at low CPU utilization, virtualization combines the processing power onto fewer servers that operate at higher utilization. Virtualization will drastically reduce the number of servers in a data center, reducing required server power and consequently the size of the necessary cooling equipment. Some overhead is required to implement virtualization, but this is minimal compared to the savings that can be achieved.

3 Environmental Conditions

3.1 2021 ASHRAE Equipment Thermal Guidelines for Data Processing Environments

The first step in designing the cooling and air management systems in a data center is to look at the standardized operating environments for equipment set forth by the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) or Network Equipment Building System (NEBS). In 2008, ASHRAE in collaboration with IT equipment manufacturers expanded their recommended environmental envelope for inlet air entering IT equipment. The revision of this envelope allows greater flexibility in facility operations and contributes to reducing the overall energy consumption. The expanded recommended and allowable envelopes for Class A1 to A4 data centers are shown in Figure 3-1 and Figure 3-2 and tabulated in Table 3-1 (for more details on data center type, different levels of altitude, etc., refer to the referenced ASHRAE publication, Thermal Guidelines for Data Processing Environments, 5th Edition, 2021).




Figure 3-1. Classes A1–A4 allowable and recommended operating conditions for low level of pollutants (I-P Units)

Source: Thermal Guidelines for Data Processing Environments, ASHRAE




Figure 3-2. Classes A1–A4 allowable and recommended operating conditions for low level of pollutants (SI Units)

Source: *Thermal Guidelines for Data Processing Environments*, ASHRAE

Table 3-1. ASHRAE 2021 Thermal Guidelines for Air Cooling

Source: *Thermal Guidelines for Data Processing Environments*, ASHRAE)

		Recommended	Allowable			
		A1 – A4	A1	A2	A3	A4
Dry-Bulb Temp	Degree F	64.4 to 80.6	59 to 89.6	50 to 95	41 to 104	41 to 113
	Degree C	18 to 27	15 to 32	10 to 35	5 to 40	5 to 45
Humidity	RH	8% to 70%*	8% to 80%*	8% to 80%*	8% to 85%*	8% to 90%*

Notes:

- Refer to ASHRAE for the high-level pollutants and max rate of change for tape storage.
- If testing shows corrosion levels exceed these limits, then drops to 50% max.

It is important to recognize the difference between the recommended and allowable envelopes presented in the ASHRAE guidelines. The recommended environmental envelope is intended to guide operators of data centers on the energy-efficient operation of data centers while maintaining high reliability. Facilities should be designed and operated to target the recommended range. The allowable envelope outlines the environmental boundaries tested by equipment manufacturers for equipment functionality, not reliability.

Another important factor to consider regarding the optimal server inlet air temperature is that variable speed fans in the servers are usually controlled to the internal server temperature.

Operating the data center at server inlet air conditions above the recommended range may cause these internal fans to operate at higher speeds and consume more power. For example, a server could have a 30% increase in server fan speed with an increase in inlet air temperature from 77°F to 91°F. This increase in inlet air temperature results in more than doubling the server fan power by applying the fan affinity law where fan power increases with the cube of fan speed. The effect of increasing server inlet air temperature on server fan power should be weighed against the data center cooling system energy savings, although in most cases cooling savings are much larger than the increase of server fan power.

ASHRAE has defined a new air-cooling class, Class H1, to address high-density servers that use high-powered components such as CPUs, graphic processing units (GPUs), and memory requiring increased cooling. Data center operators should consider implementing Direct Liquid Cooling of high-density servers (more details in section 5.6).

4 Air Management

Air management for data centers entails all the design and configuration details that go into minimizing or eliminating mixing between the cooling air supplied to equipment and the hot air rejected from the equipment. Effective air management implementation minimizes the bypass of cooling air around rack intakes and the recirculation of heat exhaust back into rack intakes. Air management generally allows increasing the supply air temperature while ensuring the “comfort” for the IT equipment. The latter can be monitored by following the guidelines in the following Berkeley Lab document: <https://datacenters.lbl.gov/resources/thermal-guidelines-and-temperature>. Another document provides a roadmap for accessing the IT equipment intake air temperatures directly from the servers: <http://datacenters.lbl.gov/resources/accessing-onboard-server-sensors-energy>. Either way, the temperature data can be made actionable by using the Rack Cooling Index (RCI), see section 8 Benchmarking. When designed correctly, an air management system can reduce operating costs, reduce first cost equipment investment, increase the data center’s power density (Watts/square foot), and reduce heat-related processing interruptions or failures. A few key design issues include the configuration of equipment’s air intake and heat exhaust ports, the location of supply and returns, the large-scale airflow patterns in the room, and the temperature set points of the airflow.

4.1 Implement Cable Management

Under-floor and over-head obstructions often interfere with the distribution of cooling air. Such interferences can significantly reduce the air handlers’ airflow as well as negatively affect the air distribution. Cable congestion in raised-floor plenums can sharply reduce the total airflow as well as degrade the airflow distribution through the perforated floor tiles. Both effects promote the development of hot spots.

A minimum effective (clear) height of 24 inches should be provided for raised-floor installations. Greater under-floor clearance can help achieve a more uniform pressure distribution in some cases.

A data center should have a cable management strategy to minimize air flow obstructions caused by cables and wiring. This strategy should target the entire cooling air flow path, including the rack-level IT equipment air intake and discharge areas as well as under-floor areas.

Persistent cable management is a key component of maintaining effective air management. Instituting a cable mining program (i.e., a program to remove abandoned or inoperable cables) as part of an ongoing cable management plan will help optimize the air delivery performance of data center cooling systems.

4.2 Aisle Separation and Containment

A basic hot aisle/cold aisle configuration is created when the equipment racks and the cooling system’s air supply and return are designed to prevent mixing of the hot rack exhaust air and the cool supply air drawn into the racks. As the name implies, the data center equipment is laid out

in rows of racks with alternating cold (rack air intake side) and hot (rack air heat exhaust side) aisles between them. Strict hot aisle/cold aisle configurations can significantly increase the air-side cooling capacity of a data center's cooling system.

All equipment is installed into the racks to achieve a front-to-back airflow pattern that draws conditioned air in from cold aisles, located in front of the equipment, and rejects heat out through the hot aisles behind the racks. Equipment with non-standard exhaust directions must be addressed in some way (shrouds, ducts, etc.) to achieve a front-to-back airflow. The rows of racks are placed back-to-back, and holes through the rack (vacant equipment slots) are blocked off on the intake side to create barriers that reduce recirculation, as shown in Figure 4-1.

Additionally, cable openings in raised floors and ceilings should be sealed as tightly as possible. With proper isolation, the temperature of the hot aisle no longer impacts the temperature of the racks or the reliable operation of the data center; the hot aisle becomes a heat exhaust. The air-side cooling system is configured to supply cold air exclusively to the cold aisles and pull return air only from the hot aisles.




Figure 4-1. Example of a hot aisle/cold aisle configuration

Source: Rumsey Engineers

The hot rack exhaust air is not mixed with cooling supply air and, therefore, can be directly returned to the air handler through various collection schemes, returning air at a higher temperature, often 85°F or higher. Depending on the type and loading of a server, the air temperature rise across a server can range from 10°F to more than 40°F. Thus, rack return air temperatures can exceed 100°F when densely populated with highly loaded servers. Higher return temperatures extend economizer hours significantly and allow for a control algorithm that reduces supply air volume, saving fan power. If the hot aisle temperature is high enough, this air can be used as a heat source in many applications. In addition to energy savings, higher equipment power densities are also better supported by this configuration. The significant increase in economizer hours afforded by a hot aisle/cold aisle configuration can improve equipment reliability in mild climates by providing emergency compressor-free data center

operation when outdoor air temperatures are below the data center equipment's top operating temperature (typically 90°F to 95°F).

Using flexible plastic barriers, such as plastic supermarket refrigeration covers (i.e. “strip curtains”), or other solid partitions to seal the space between the tops of the rack and air return location can greatly improve hot aisle/cold aisle isolation while allowing flexibility in accessing, operating, and maintaining the computer equipment below. One recommended design configuration, shown in Figure 4-2, supplies cool air via an under-floor plenum to the racks; the air then passes through the equipment in the rack and enters a separated, semi-sealed area for return to an overhead plenum. This approach uses a baffle panel or barrier above the top of the rack and at the ends of the hot aisles to mitigate “short-circuiting” (the mixing of hot and cold air). These changes should reduce fan energy requirements (as long as the fans are on VSDs) and could result in facility energy savings.

There are modular containment systems that are manufactured off site and can be installed quickly and cleanly. These solutions provide good hot aisle/cold aisle isolation. They can receive cooling air from the data center room cooling system or can be cooled by In-Row cooling units that are built into the modular containment system.

Generalized thermal best practices are outlined in the following Berkeley Lab document and summarized below: <https://datacenters.lbl.gov/resources/thermal-guidelines-and-temperature>. These best-practice recommendations are a first step towards temperature management and measurements in data centers, ultimately saving infrastructure energy as well as protecting the electronic equipment.

- Use environmental specifications per ASHRAE or NEBS. Select the default Recommended temperature range of 65°F to 80°F. The temperature should be controlled to be within the Recommended range at the IT equipment air intakes. Warmer intake temperature and supply air temperature require less cooling energy.
- Select Allowable temperature range A2 or higher since IT equipment is often rated for wider temperatures than the A1 range. ASHRAE's allowable range is 59°F to 90°F for Class A1 and 50°F to 95°F for Class A2. Manufacturers are offering extended Recommended and Allowable temperature range equipment.
- Increase supply air temperature to keep the most demanding intake air temperature as close to 80°F as possible. Leaving room for error, a setpoint of 77°F to 79°F may be the most practical approach. Small incremental temperature changes are recommended to avoid local IT overheating and compromised reliability, and only after implementing air management.
- Control supply temperature (and airflow) based on IT equipment intake air temperatures and not on the return temperature. Use wired or wireless external-to-rack temperature sensors or, even better, network data exchange with IT equipment on-board temperature sensors. All ENERGY STAR servers have the latter capability.

- Showing compliance with equipment intake air temperature specifications is the ultimate cooling performance metric in data centers. The Department of Energy’s Air Management Tool uses the RCI for that purpose.

Control of the fan speed based on the IT equipment needs is critical to achieving savings. Variable speed drives on direct expansion (DX) computer room air conditioning (CRAC) unit supply fans are widely available. For more descriptions on CRAC units and common energy-efficiency options, refer to section 5.1 Direct Expansion Systems.




Figure 4-2. Sealed hot aisle/cold aisle configuration

Source: Rumsey Engineers

4.3 Optimize Supply and Return Air Configuration

Hot aisle/cold aisle configurations can be served by overhead or under-floor air distribution systems. When an overhead system is used, supply outlets that “dump” the air directly down should be used in place of traditional office diffusers that throw air to the sides, which results in undesirable mixing and recirculation with the hot aisles. The diffusers should be located directly in front of racks, above the cold aisle. In some cases, return grilles or simply open ducts have been used. The temperature monitoring to control the air handlers should be located in areas in front of the computer equipment, not on a wall behind the equipment. Use of overhead variable air volume allows equipment to be sized for excess capacity and yet provides optimized operation at part-load conditions with turn down of variable speed fans. Where a rooftop unit is being used, it should be located centrally over the served area—the required reduction in ductwork will lower cost and slightly improve efficiency. Also keep in mind that overhead delivery tends to reduce temperature stratification in cold aisles as compared to under-floor air delivery.

Under-floor air supply systems have a few unique concerns. The under-floor plenum often serves both as a duct and a wiring chase. Coordination throughout design and into construction and operation throughout the life of the center is necessary since paths for airflow can be blocked by

electrical or data trays and conduits. The location of supply tiles needs to be carefully considered to prevent short circuiting of supply air and checked periodically if users are likely to reconfigure them. Removing or adding tiles to fix hot spots can cause problems throughout the system. Another important concern to be aware of is high air velocity in the under-floor plenum. This can create localized negative static pressure and draw room air back into the under-floor plenum. Equipment closer to downflow CRAC units or computer room air handlers (CRAH) can receive too little cooling air due to this effect. Deeper plenums and careful layout of CRAC/CRAH units allow for a more uniform under-floor air static pressure. For more description on CRAH units as they relate to data center energy efficiency, refer to section 5.2 Air Handlers.

4.4 Other Benefits of Raising Temperature Set Points

Air-side economizer energy savings are realized by utilizing a control algorithm that brings in outside air whenever it is cooler than the return air and when humidity conditions are acceptable (refer to section 5.4.1 Airside Economizer in section 5.4 Free Cooling for further detail on economizer control optimization). In order to save energy, the temperature outside does not need to be below the data center's temperature set point; it only has to be cooler than the return air that is exhausted from the room. As the return air temperature is increased through the use of good air management, as discussed in the preceding sections, the temperature at which an air-side economizer will save energy is correspondingly increased. Designing for a higher return air temperature increases the number of hours that outside air, or a waterside economizer/free cooling, can be used to save energy.

A higher return air temperature also makes better use of the capacity of standard package units, which are designed to condition office loads. This means that a portion of their cooling capacity is configured to serve humidity (latent) loads. Data centers typically have very few occupants and small outside air requirements, and, therefore, have negligible latent loads. While the best course of action is to select a unit designed for sensible cooling loads only or to increase the airflow, an increased return air temperature can convert some of a standard package unit's latent capacity into usable sensible capacity very economically. This may reduce the size and/or number of units required.

A warmer supply air temperature set point on chilled water air handlers allows for higher chilled water supply temperatures which consequently improves the chilled water plant operating efficiency. Operation at warmer chilled water temperatures also increases the potential hours that a water-side economizer can be used (refer to section 5.4.2 Water-Side Economizer in section 5.4 Free Cooling for further detail).

5 Cooling Systems

When beginning the design process and equipment selections for cooling systems in data centers, it is important to always consider initial and future loads, in particular part- and low-load conditions, as the need for digital data is ever-expanding.

5.1 Direct Expansion Systems

Packaged DX air conditioners likely compose the most common type of cooling equipment for smaller data centers. These units are generally available as off-the-shelf equipment from manufacturers (commonly described as CRAC units). There are, however, several options available to improve the energy efficiency of cooling systems employing DX units.

Packaged rooftop units are inexpensive and widely available for commercial use. Several manufacturers offer units with multiple and/or variable speed compressors to improve part-load efficiency. These units reject the heat from the refrigerant to the outside air via an air-cooled condenser. An enhancement to the air-cooled condenser is a device which sprays water over the condenser coils. The evaporative cooling provided by the water spray improves the heat rejection efficiency of the DX unit. Additionally, these units are commonly offered with air-side economizers. Depending on the data center's climate zone and air management, a DX unit with air-side economizer can be a very energy-efficient cooling option for a small data center. (For further discussion, refer to section 4.4 Raising Temperature Set Points and subsection 5.4.1 Air-Side Economizer in section 5.4 Free Cooling.)

Indoor CRAC units are available with a few different heat rejection options. Air-cooled CRAC units include a remote air-cooled condenser. As with the rooftop units, adding an evaporative spray device can improve the air-cooled CRAC unit efficiency. For climate zones with a wide range of ambient dry bulb temperatures, apply parallel VSD control of the condenser fans to lower condenser fan energy compared to the standard staging control of these fans.

CRAC units packaged with water-cooled condensers are often paired with outdoor drycoolers. The heat rejection effectiveness of outdoor drycoolers depends on the ambient dry bulb temperature. A condenser water pump distributes the condenser water from the CRAC units to the drycoolers. Compared to the air-cooled condenser option, this water-cooled system requires an additional pump and an additional heat exchanger between the refrigerant loop and the ambient air. As a result, this type of water-cooled system is generally less efficient than the air-cooled option. A more efficient method for water-cooled CRAC unit heat rejection employs a cooling tower. To maintain a closed condenser water loop to the outside air, a closed loop cooling tower can be selected. A more expensive but more energy-efficient option would be to select an oversized open-loop tower and a separate heat exchanger where the latter can be selected for a very low (less than 3°F) approach. In dry climates, a system composed of water-cooled CRAC units and cooling towers can be designed to be more energy efficient than air-

cooled CRAC unit systems. (Refer to subsection 5.3.1 Efficient Equipment in section 5.3 High-Efficiency Chilled Water Systems for more information on selecting an efficient cooling tower.)

A type of water-side economizer can be integrated with water-cooled CRAC units. A pre-cooling water coil can be added to the CRAC unit upstream of the evaporator coil. When ambient conditions allow the condenser water to be cooled (by either drycooler or cooling tower) to the point that it can provide a direct cooling benefit to the air entering the CRAC unit, condenser water is diverted to the pre-cooling coil. This will reduce, or at times eliminate, the need for compressor-based cooling from the CRAC unit. Some manufacturers offer this pre-cooling coil as a standard option for their water-cooled CRAC units.

5.2 Air Handlers

5.2.1 Central vs. Modular Systems

Better performance has been observed in data center air systems that utilize specifically designed central air handler systems. A centralized system offers many advantages over the traditional multiple distributed unit system that evolved as an easy, drop-in computer room cooling appliance (commonly referred to as a CRAH unit). Centralized systems use larger motors and fans that tend to be more efficient. They are also well suited for variable volume operation through the use of VSDs and maximize efficiency at part-loads.




Figure 5-1. Comparison of distributed air delivery to central air delivery

Source: Rumsey Engineers

In Figure 5-1, the pie charts show the electricity consumption distribution for two data centers. Both are large facilities, with approximately equivalent data center equipment loads, located in adjacent buildings, and operated by the same company. The facility on the left uses a multiple

distributed unit system based on air-cooled CRAC units, while the facility on the right uses a central air handler system. An ideal data center would use 100% of its electricity to operate data center equipment—energy used to operate the fans, compressors and power systems that support the data center is strictly overhead cost. The data center supported by a centralized air system (on the right) uses almost two-thirds of the input power to operate revenue-generating data center equipment, compared to the multiple small unit system that uses just over one-third of its power to operate the actual data center equipment. The trend seen here has been consistently supported by benchmarking data. The two most significant energy saving methods are water-cooled equipment and efficient centralized air handler systems. CRAH units can also be installed in or adjacent to a data center. A CRAH unit that uses 55°F or higher chilled water, and efficient fans with VSD with low pressure drop air design can minimize the energy use of cooling systems.

Most data center loads do not vary appreciably over the course of the day, and the cooling system is typically significantly oversized. A centralized air handling system can improve efficiency by taking advantage of surplus and redundant capacity to actually improve efficiency. The maintenance benefits of a central system are well known, and the reduced footprint and maintenance traffic in the data center are additional benefits. Implementation of an airside economizer system is simplified with a central air handler system. Optimized air management, such as that provided by hot aisle/cold aisle configurations, is also easily implemented with a ducted central system. Modular units are notorious for battling each other to maintain data center humidity set points. That is, one unit can be observed to be dehumidifying while an adjacent unit is humidifying. Instead of modular units independently controlled, a centralized control system using shared sensors and set points ensures proper communication among the data center air handlers. Even with modular units, humidity control over make-up air should be all that is required.

5.2.2 Low Pressure Drop Air Delivery

A low-pressure drop design (“oversized” ductwork or a generous under-floor) is essential to optimizing energy efficiency by reducing fan energy and facilitates long-term buildout flexibility. Ducts should be as short and straight as possible in length, and sized significantly larger than typical office systems, since 24-hour operation of the data center increases the value of energy use over time relative to first cost. Since loads often only change when new servers or racks are added or removed, periodic manual air flow balancing can be more cost-effective than implementing an automated air flow balancing control scheme.

5.3 High-Efficiency Chilled Water Systems

5.3.1 Efficient Equipment

Use efficient water-cooled chillers in a central chilled water plant. A high-efficiency variable frequency drive (VFD)-equipped chiller with an appropriate condenser water reset is typically the most efficient cooling option for large facilities. Chiller part-load efficiency should be considered since data centers often operate at less than peak capacity. Chiller part-load

efficiencies can be optimized with variable frequency driven compressors, high evaporator temperatures, and low entering condenser water temperatures.

Oversized cooling towers with VFD-equipped fans will lower water-cooled chiller plant energy. For a given cooling load, larger towers have a smaller approach to ambient wet bulb temperature, thus allowing for operation at lower cold condenser water temperatures and improving chiller operating efficiency. The larger fans associated with the oversized towers can be operated at lower speeds to lower cooling tower fan energy compared to a smaller tower.

Condenser water and chilled water pumps should be selected for the highest pumping efficiency at typical operating conditions, rather than at full load condition.

5.3.2 Optimize Plant Design and Operation

Data centers offer a number of opportunities in central plant optimization, both in design and operation. A medium-temperature, as opposed to low-temperature, chilled water loop design using a water supply temperature of 55°F or higher improves chiller efficiency and eliminates uncontrolled phantom dehumidification loads (refer to section 5.7 Humidification and section 5.8 Controls). Higher temperature chilled water also allows more water-side economizer hours, in which the cooling towers can serve some or the entire load directly, reducing or eliminating the load on the chillers. The condenser water loop should also be optimized; a 5°F to 7°F approach cooling tower plant with a condenser water temperature reset pairs nicely with a variable speed chiller to offer large energy savings.

5.3.3 Efficient Pumping

A well-thought-out, efficient pumping design is an essential component to a high-efficiency chilled water system. Pumping efficiency can vary widely depending on the configuration of the system, and whether the system is for an existing facility or new construction. Listed below are general guidelines for optimizing pumping efficiency for existing and new facilities of any configuration.

Existing Facilities:

- Reduce the average chilled water flow rate corresponding to the typical load.
- Convert existing primary/secondary chilled water pumping system to primary-only.
- Convert existing system from constant flow to variable flow.
- Eliminate unnecessary bypassed chilled water by replacing 3-way chilled water valves with 2-way valves. Reduce the pressure drop of the chilled water distribution system by opening pump balancing valves and allowing pump VFDs to limit the flow rate.
- Reduce the chilled water supply pressure set point.
- Add a chilled water pumping differential pressure set point reset control sequence.

New Construction:

- Reduce the average chilled water flow rate corresponding to the typical load.
- Implement primary-only variable flow chilled water pumping.
- Specify an untrimmed impeller; do not install pump balancing valves and instead use a VFD to limit pump flow rate.
- Design for a low water supply pressure set point.
- Specify a water pumping differential pressure set point reset control sequence.
- Design a low pressure drop pipe layout for pumps.
- Specify 2-way chilled water valves instead of 3-way valves.
- Install VFDs on all pumps and run redundant pumps at lower speeds.

5.4 Free Cooling

5.4.1 Air-Side Economizer

The cooling load for a data center is independent of the outdoor air temperature. The maximum recommended air inlet temperature for most IT equipment is 80°F (per the guidelines in section 3.1), which allows for many more hours of economizer operations than an office building. Most nights and during mild winter conditions, the lowest cost option to cool data centers is an air-side economizer; however, a proper engineering evaluation of the local climate conditions must be completed to evaluate whether this is the case for a specific data center.

Due to the contamination concerns associated with data centers, careful control and design work may be required to ensure that cooling savings are not lost because of excessive filtration requirements. Data center professionals are split in the perception of risk when using this strategy. It is standard practice, however, in the telecommunications industry to equip their facilities with air-side economizers. Some IT-based centers routinely use outside air without apparent complications, but others are concerned about contamination and environmental control for the IT equipment in the room. If minimum efficiency reporting value (MERV) 13 filters are used, the negative effects of outside air pollutants are small. Nevertheless, outside air economizing is implemented in many data center facilities and results in energy-efficient operation. In fact, many data centers in cool climates use only economizer (air and/or water) cooling and no DX cooling.

Control strategies to deal with temperature and humidity fluctuations must be considered along with contamination concerns over particulates or gaseous pollutants. For data centers with active humidity control, a dewpoint temperature lockout scheme should be used as part of the air-side economizer control strategy. This scheme prevents high outside air dehumidification and humidification loads by tracking the moisture content of the outside air and locking out the economizer when the air is either too dry or too moist. Mitigation steps may involve filtration or other measures. Other contamination concerns such as salt or corrosive matter should be

evaluated. Generally, concern over contamination should be limited to unusually harsh environments such as pulp and paper mills or large chemical spills.

Wherever possible, outside air intakes should be located on the north side of buildings in the northern hemisphere where there is significantly less solar heat gain compared to the south side.

Finally, air-side economizers are not easy to add to an existing data center due to its large duct installations. That is where the water-side economizer may come in handy as discussed in the next section.

5.4.2 Water-Side Economizer

Free cooling can be provided via a waterside economizer, which uses the evaporative cooling capacity of a cooling tower to produce chilled water to cool the data center during mild outdoor conditions. Free cooling is usually best suited for climates that have wet bulb temperatures lower than 55°F for 3,000 or more hours per year. It most effectively serves chilled water loops designed for 50°F and above chilled water or lower temperature chilled water loops with significant surplus air handler capacity in normal operation. A heat exchanger is typically installed to transfer heat from the chilled water loop to the cooling tower water loop while isolating these loops from each other. Locating the heat exchanger upstream from the chillers, rather than in parallel to them, allows for integration of the water-side economizer as a first stage of cooling the chilled water before it reaches the chillers. During those hours when the water-side economizer can remove enough heat to reach the chilled water supply set point, the chilled water can be bypassed around the chillers. When the water-side economizer can remove heat from the hot chilled water but not enough to reach set point, the chillers operate at reduced load to meet the chilled water supply set point. Many data centers in cool climates use only water-side economizer cooling and no chillers.

5.5 Thermal Storage

Thermal storage is a method of storing thermal energy in a reservoir for later use, and is particularly useful in facilities with particularly high cooling loads such as data centers. It can result in peak electrical demand savings and improve chilled water system reliability. In climates with cool, dry nighttime conditions, cooling towers can directly charge a chilled water storage tank, using a small fraction of the energy otherwise required by chillers. A thermal storage tank can also be an economical alternative to additional mechanical cooling capacity; for example, water storage provides the additional benefit of backup make-up water for cooling towers.

5.6 Direct Liquid Cooling

Direct liquid cooling refers to a number of different cooling approaches that all share the same characteristic of transferring waste heat to a fluid at or very near the point the heat is generated, rather than transferring it to room air and then conditioning the room air. Liquid cooling can serve higher heat densities and be much more efficient than traditional air cooling, as water flow is a much more efficient method of transporting heat. Energy efficiencies will be realized when

such systems allow the use of a medium temperature chilled water supply and by reducing the size and power consumption of fans serving the data center. These warmer chilled water supply temperatures facilitate the pairing of liquid cooling with a water-side economizer, further increasing potential energy savings. Higher ASHRAE W-classes (see section below) may allow the use of dry coolers to reject heat to atmosphere saving even more energy and water. An ASHRAE TC 9.9 white paper, *Emergence and Expansion of Liquid Cooling in Mainstream Data Centers*, indicates liquid cooling will become more mainstream due to increasing socket power (CPU/GPU/TPU) and included Figure 5-2 that shows when a transition from air to liquid cooling based on socket power is likely.




Figure 5-2. Air cooling versus liquid cooling, transition, and temperatures

Source: [Emergence and Expansion of Liquid Cooling in Mainstream Data Centers](#), ASHRAE

There are a variety of different liquid-cooling technologies that can be used within data centers. Liquid-cooling fluids on the IT loop side can range from treated water to glycol-based solutions to dielectric fluids (nonconductive). Several liquid-cooled solutions are considered hybrid solutions, in that the majority (but not all) of the heat load from IT equipment is captured and removed by the liquid—with the remaining heat load removed by traditional air cooling. However, some liquid-cooled solutions capture practically all the heat load without the use of any fans, which is the most energy-efficient. Most liquid-cooling approaches involve a cooling distribution unit (CDU), which interfaces with the facility cooling loop and provides cooling liquid at the appropriate temperature, pressure, and chemistry for the IT equipment (Figure 5-3).




Figure 5-3. CDU liquid-cooling system within a data center

Source: Emergence and Expansion of Liquid Cooling in Mainstream Data Centers, ASHRAE

Liquid-cooling technologies can be categorized into three technology classes:

- Localized air-to-liquid heat exchanger (rear-door heat exchangers): Involves air-to-liquid heat exchanger (coil) placed at the back of a server-rack that captures server heat, resulting in energy efficiency gains by transferring heat to liquid closer to the server-rack source (versus computer room air-handler or air-handler units).
 - o Variations include rear-door heat exchanger (also referred to as cooling doors), enclosed cabinet, in-row coolers, and overhead cooling, all utilizing some type of air-to-liquid heat exchanger close-coupled with the IT source.
 - o Options for cooling doors include active system where additional fans draw air through coils, or passive system where server fans move air through coils.
- Cold plates: Standard fin-based heat sinks on chips are replaced with cold plates that have liquid flowing through channels to remove heat (solution generally involves a CDU).
 - o Variety of technology solutions exist, and this approach can also be applied to memory and other heat-producing components.
- Immersion: Electronics submerged in a dielectric fluid (nonconductive); approaches include:
 - o Single-phase immersion cooling, in which heat load is transferred to dielectric oil/fluid that is pumped around electronics by a CDU.
 - o Two-phase immersion, in which an engineered dielectric fluid with a boiling point below IT components' maximum operating temperatures removes heat load by undergoing a liquid-to-gas phase change, with that vapor transferring heat to a vapor-to-liquid heat exchanger, which then condenses back to a liquid in a passive cycle.

The cold plate and immersion classes are considered direct-liquid-cooling technologies. High-performance computing data centers have been early adopters of direct liquid cooling due to rack power densities (where densities of 60 kW per compute rack were observed in 2013, and recently surpassing 125+ kW per compute rack).

ASHRAE has also recently renamed the ASHRAE water classifications to clearly define the server supply temperature range. The upper temperature limits are now included in the W number (in degrees Celsius). The classes are newly named as follows: W17 (previously W1), W27 (W2), W32 (W3), W40 (new), W45 (W4), and W+ (W5). This change to ASHRAE W classes has been included in the fifth edition of Thermal Guidelines for Data Processing Environments (released in 2021) and will be included in the upcoming third edition of the Liquid Cooling Guidelines for Datacom Equipment Centers.

5.7 Humidification

Most data centers do not require humidification to maintain the recommended minimum humidity of 15.8°F dew point (~10% relative humidity). Low-energy humidification techniques can replace traditional electric resistance humidifiers with an adiabatic approach that uses the heat present in the air or recovered from the computer heat load for humidification. Ultrasonic humidifiers, evaporative wetted media and micro droplet spray are some examples of adiabatic humidifiers. An electric resistance humidifier requires about 430 Watts to boil one pound of 60°F water, while a typical ultrasonic humidifier only requires 30 Watts to atomize the same pound of water (although ultrasonic humidifiers may require deionized water). These passive humidification approaches also provide evaporative cooling of the air, in contrast to an electric resistance humidifier heating the air, which further saves energy by reducing the load on the cooling system.

5.8 Controls

More options are now available for dynamically allocating IT resources as computing or storage demands vary. Within the framework of ensuring continuous availability, a control system should be programmed to maximize the energy efficiency of the cooling systems under variable ambient conditions as well as variable IT loads.

Variable speed drives on CRAH and CRAC units allow for varying the airflow as the cooling load fluctuates. For raised-floor installations, the fan speed should be controlled to maintain an under-floor pressure set point. However, cooling air delivery via conventional raised floor tiles can be ill-suited for responding to the resulting dynamic heat load without either over-cooling the space or starving some areas of sufficient cooling. Variable air volume air delivery systems are a much better solution for consistently providing cooling when and where it is needed. Supply air and supply chilled water temperatures should be set as high as possible while maintaining the necessary cooling capacity. All CRAH/CRAC units should communicate and operate at the same temperature and humidity.

Data centers often over-control humidity, which results in no real operational benefits and increases energy use. Tight humidity control is a carryover from legacy data centers and generally can be relaxed or eliminated for most locations.

Humidity controls are frequently not centralized. This can result in adjacent units serving the same space fighting to meet the humidity set point, with one humidifying while the other is dehumidifying. Humidity sensor drift can also contribute to control problems if sensors are not regularly recalibrated. One very important consideration to reducing unnecessary humidification is to operate the cooling coils of the air-handling equipment above the dew point (usually by running chilled water temperatures above 50°F), thus eliminating unnecessary dehumidification.

On the chilled water plant side, variable flow pumping and chillers equipped with variable speed driven compressors should be installed to provide energy-efficient operation during low load conditions. Another option to consider for increasing chiller plant efficiency is to actively reset the chilled water supply temperature higher during low load conditions. In data centers located in relatively dry climates and which experience relatively low partial loads, implementing a water-side economizer can provide tremendous savings over the course of the year (see earlier discussion in section 5.4.2 Water-Side Economizers).

6 Electrical Systems

Similar to cooling systems, it is important to always consider initial and future loads, in particular part- and low-load conditions, when designing and selecting equipment for a data center's electrical system.

6.1 Power Distribution

Data centers typically have an electrical power distribution path consisting of the utility service, switchboard, switchgear, alternate power sources (i.e. backup generator), paralleling equipment for redundancy (i.e., multiple UPSs and PDUs), and auxiliary conditioning equipment (i.e. line filters, capacitor bank). These components each have a heat output that is tied directly to the load in the data center. Efficiencies can range widely between manufacturers and variations in how the equipment is designed. However, operating efficiencies can be controlled and optimized through thoughtful selection of these components.

6.1.1 Uninterruptible Power Supplies (UPS)

UPS systems provide backup power to data centers, and can be based on battery banks, rotary machines, fuel cells, or other technologies. A portion of all the power supplied to the UPS to operate the data center equipment is lost to inefficiencies in the system which can lead to significant energy losses. The first step to minimize these losses is to evaluate which equipment, if not the entire data center, requires a UPS system. For instance, the percentage of IT power requiring UPS by a scientific computing facility (high-performance computing) can be significantly lower than the percentage required for a financial institution.

Increasing the UPS system efficiency offers direct, 24-hour-a-day energy savings, both within the UPS itself and indirectly through lower heat loads and even reduced building transformer losses. Among double conversion systems (the most commonly used data center system), UPS efficiency has improved from 85% to 90% in the 1990s, to 95% or higher in 2023. When a full data center equipment load is served through a UPS system, even a small improvement in the efficiency of the system can yield a large annual cost savings. For example, a 15,000-square-foot data center with IT equipment operating at 100 Watts/square foot requires 13,140 megawatt-hours (MWh) of energy annually for the IT equipment. If the UPS system supplying that power has its efficiency improved from 90% to 95%, the annual energy bill will be reduced by 768,421 kWh, or about \$90,000 at \$0.12/kWh, plus significant additional cooling system energy savings from the reduced cooling load. For battery-based UPS systems, use a design approach that keeps the UPS load factor as high as possible. This usually requires using multiple smaller units.

Redundancy in particular requires design attention; operating a single large UPS in parallel with a 100% capacity identical redundant UPS unit (n+1 design redundancy) results in very low load factor operation, at best no more than 50% at full design buildout. Consider a UPS system sized for two UPS units with n+1 redundancy, with both units operating at 30% load factor. If the same load is served by three smaller units (also sized for n+1 redundancy), then these units will operate at 40% load factor. This 10% increase in load factor can result in a 1.2% efficiency increase (see

Figure 6-1). For a 100-kW load, this efficiency increase can result in savings of approximately 13,000 kWh annually.




Figure 6-1. Typical UPS efficiency curve for 100 kVA capacity and greater

Source: Rumsey Engineers

Some double conversion UPS systems (which offer the highest degree of power conditioning) have the ability to operate in the more efficient line conditioning mode, usually advertised as “economy” or “eco” mode or “energy saver” mode, which can bring efficiency up to 99%.

6.1.2 Power Distribution Units (PDU)

A PDU passes conditioned power that is sourced from a UPS or generator to provide reliable power distribution to multiple pieces of equipment. It provides many outlets to power servers, networking equipment and other electronic devices that require conditioned and/or continuous power. Maintaining a higher voltage in the source power lines fed from a UPS or generator allows for a PDU to be located more centrally within a data center. As a result, the conductor lengths from the PDU to the equipment are reduced and less power is lost in the form of heat.

Specialty PDUs that convert higher voltage (208V AC or 480V AC) into lower voltage (120V AC) via a built-in step-down transformer for low voltage equipment are commonly used as well. Transformers lose power in the form of heat when voltage is being converted. The parameters of the transformer in this type of PDU can be specified such that the energy efficiency is optimized. A dry-type transformer with a 176°F temperature rise uses 13% to 21% less energy than a 302°F rise unit. The higher-efficiency 176°F temperature rise unit has a first-cost premium; however, the cost is usually recovered in the energy cost savings. In addition, many transformers tend to operate most efficiently when they are loaded in the 20% to 50% range. Selecting a PDU with a

transformer at an optimized load factor will reduce the loss of power through the transformer. Energy can also be saved by reducing the number of installed PDUs with built-in transformers.

6.1.3 Distribution Voltage Options

Another source of electrical power loss for both AC and DC distribution is that of the conversions required from going from the original voltage supplied by the utility (usually a medium voltage of around 12 kV AC or more) to that of the voltage at each individual device within the data center (usually a low voltage around 120V AC to 240V AC). Designing a power distribution network that delivers all of the required voltages while minimizing power losses is often a challenging task. Following are general guidelines for delivering electrical power in the most energy-efficient manner possible:

- Minimize the resistance by increasing the cross-sectional area of the distribution path and making it as short as possible.
- Maintain a higher voltage for as long as possible to minimize the current.
- Use switch-mode transistors for power conditioning.
- Locate all voltage regulators close to the load to minimize distribution losses at lower voltages.

6.2 Demand Response

Demand response refers to the process by which facility operators voluntarily curb energy use during times of peak demand. Many utility programs offer incentives to business owners that implement this practice on hot summer days or other times when energy demand is high and supply is short. Demand response programs can be executed by reducing loads through a building management system or switching to backup power generation.

For reducing loads when a demand response event is announced, data center operators can take certain reduction measures such as dimming a third of their lighting or powering-off idle office equipment. Automated network building system solutions can make this a simple, efficient, and inexpensive process.

6.3 DC Power

In a conventional data center, power is supplied from the grid as AC power and distributed throughout the data center infrastructure as AC power. However, most of the electrical components within the data center, as well as the batteries storing the backup power in the UPS system, require DC power. As a result, the power must go through multiple conversions resulting in power loss and wasted energy.

One way to reduce the number of times power needs to be converted is by utilizing a DC power distribution. This has not yet become a common practice in data centers (but used widely in telecom facilities) and, therefore, could carry significantly higher first costs, but it has been tested at several facilities. A study done by Lawrence Berkeley National Laboratory in 2007

compared the benefits of adopting a 380V DC power distribution for a datacom facility to a traditional 480V AC power distribution system. The results showed that the facility using the DC power had a 7% reduction in energy consumption compared to the typical facility with AC power distribution. Other DC distribution systems are available including 575V DC and 48V DC. These systems offer energy savings as well.

6.4 Lighting

Data center spaces are not uniformly occupied and, therefore, do not require full illumination during all hours of the year. UPS, battery, and switch gear rooms are examples of spaces that are infrequently occupied. Therefore, zone-based occupancy sensors throughout a data center can have a significant impact on reducing the lighting electrical use. Careful selection of an efficient lighting layout (e.g., above aisles and not above the server racks), and type (e.g., LED) will also reduce not only the lighting electrical usage but also the load on the cooling system. The latter leads to secondary energy savings.

7 Use of Waste Heat for Energy-Efficient Design

7.1 Use of Waste Heat

The higher the cooling air or water temperature leaving the server, the greater the opportunity for using waste heat. The direct use of waste heat for low temperature heating applications such as preheating ventilation air for buildings or heating water will provide the greatest energy savings. Heat recovery chillers may also provide an efficient means to recover and reuse heat from data center equipment environments for comfort heating of typical office environments.

The following conditions are recommended for a heat reuse project:

1. A heat host (or heat consumer) nearby the data center area, possibly adjacent to it, or tied into a district loop (refer to A Guide to Energy Master Planning of High-Performance Districts and Communities).
2. A temperature level that works for the heat host; optimally the heat will be directly used without heat pumps.
3. The same owner for the heat host and data center, ideally.
4. A champion, incentives, and policies that support heat reuse.

In most cases, data centers will have a redundant cooling system able to remove the heat if the heat host is unavailable. Deploying heat reuse allows a reduction or elimination of the use of chillers and in some cases of cooling towers that evaporate water (Pecchioli, Comella, Sickinger, and VanGeet 2023). Heat reuse enables water savings and offsetting fossil fuel used for heating, and this is why it appears in the Key Steps to Sustainable Data Centers after energy efficiency.

The Green Grid has proposed and defined a metric for Measuring the Benefit of Reuse Energy from a Data Center; the Energy Reuse Effectiveness, or ERE. For more information see <http://www.thegreengrid.org/en/Global/Content/white-papers/ERE>.

8 Benchmarking

Energy efficiency metrics/key performance indicators (KPIs) and benchmarks can be used to track the performance of—and identify potential opportunities to reduce energy use in—data centers.

An effective organization will consider the total cost of ownership for operational efficiency and cost, utilizing different energy-based metrics and sustainability metrics (e.g., water and carbon) to capture a view of the efficiencies at which a data center performs. The family of PUE-type of metrics (PUE, ERE, WUE, CUE) along with a utilization metric based on industry such as Gflops or Transactions per Watt are recommended.

8.1 Power Usage Effectiveness (PUE)

PUE is defined as the ratio of the total annual energy to run the data center facility to the total annual energy drawn by all IT equipment:

$$PUE = \frac{\text{Total Facility Energy}}{\text{IT Equipment Energy}}$$

Standard	Good	Better
1.6	1.4	1.1

An average data center has a PUE of 1.6; however, several recent super-efficient data centers have been known to achieve a PUE below 1.1. (According to Uptime Institute’s 2022 Global Data Center Survey, in 2022 the annual PUE average for large data centers was 1.55. See <https://uptimeinstitute.com/about-ui/press-releases/2022-global-data-center-survey-reveals-strong-industry-growth>).

PUE was originally defined by The Green Grid, but this organization passed the development, standardization and dissemination of this metric to the International Standards Organization/International Electrotechnical Commission.
<https://www.iso.org/obp/ui/#iso:std:iso-iec:30134:-2:ed-1:v1:en>.

It is important to note that PUE does not define the overall efficiency of an entire data center, but only the efficiency of the supporting equipment (the infrastructure) within a data center. This metric is defined using units of average annual energy (kWh). Using the annual measurements provides the advantage of accounting for variable free-cooling energy savings as well as the trend for dynamic IT loads due to practices such as IT power management.

PUE is defined with respect to site power draw. Another alternative definition could use a source power measurement to account for different fuel source uses.

ENERGY STAR defines a similar metric, defined with respect to source energy, Source PUE as:

$$\text{Source PUE} = \frac{\text{Total Facility Energy}}{\text{UPS Energy}}$$

Link to PUE paper (registration is required to download, but it's free to register):
<https://www.thegreengrid.org/en/resources/library-and-tools/20-PUE%3A-A-Comprehensive-Examination-of-the-Metric>.

8.2 Energy Reuse Effectiveness (ERE)

ERE is defined as the ratio of the total annual energy to run the data center facility minus the reuse energy to the total energy drawn by all IT equipment:

$$ERE = \frac{\text{Total Facility Energy} - \text{Reuse Energy}}{\text{IT Equipment Energy}}$$

Further examination of the properties of PUE and ERE brings out another important result. The range of values for PUE is mathematically bounded from 1.0 to infinity. A PUE of 1.0 means 100% of the power brought to the data center goes to IT equipment and none to cooling, lighting, or other non-IT loads. For ERE, the range is 0 to infinity. ERE does allow values less than 1.0. An ERE of 0 means that 100% of the energy brought into the data center is reused elsewhere, outside of the data center control volume.

Link to ERE paper (registration is required to download, but it's free to register):
<https://www.thegreengrid.org/en/resources/library-and-tools/242-ERE%3A-A-Metric-for-Measuring-the-Benefit-of-Reuse-Energy-From-a-Data-Center>.

8.3 Water Use Effectiveness (WUE)

The site-based WUE metric is defined as the ratio of the annual site water usage to the total energy drawn by all IT equipment:

$$WUE = \frac{\text{Annual site water usage}}{\text{IT Equipment Energy}}$$

This metric has units of liters/kilowatt-hour (L/kWh). There is a source-based definition of the WUE metric that takes into account water used on-site and water used off-site in the production of energy used on-site. See the paper for more details.

Link to WUE paper (registration is required to download, but it's free to register):
[https://www.thegreengrid.org/en/resources/library-and-tools/238-WP#35---Water-Usage-Effectiveness-\(WUE\):-A-Green-Grid-Data-Center-Sustainability-Metric](https://www.thegreengrid.org/en/resources/library-and-tools/238-WP#35---Water-Usage-Effectiveness-(WUE):-A-Green-Grid-Data-Center-Sustainability-Metric).

8.4 Carbon Use Effectiveness (CUE)

The CUE metric is defined as the ratio of the total annual CO₂ emissions caused by the data center to the total energy drawn by all IT equipment:

$$CUE = \frac{\text{Total CO}_2 \text{ emissions caused by the Total Facility Energy}}{\text{IT Equipment Energy}}$$

This metric has units of kilograms of CO₂ equivalent/kilowatt-hour (kg CO₂ eq/kWh). CO₂ emission rate can be obtained from the EPA Power Profile website by entering a U.S. ZIP code. Adjustments can be made to account for energy from onsite renewable systems.

Link to CUE paper (registration is required to download, but it's free to register):

[https://www.thegreengrid.org/en/resources/library-and-tools/241-WP#32---Carbon-Usage-Effectiveness-\(CUE\):-A-Green-Grid-Data-Center-Sustainability-Metric](https://www.thegreengrid.org/en/resources/library-and-tools/241-WP#32---Carbon-Usage-Effectiveness-(CUE):-A-Green-Grid-Data-Center-Sustainability-Metric).

8.5 Related Metrics

The following metrics provide important additional information about the data center.

8.5.1 Rack Cooling Index (RCI)

As discussed in this report, raising the data center temperature needs to be carefully executed to avoid IT equipment overheating. Effective air management needs to be in place and the IT equipment intake air temperatures should be monitored. The RCI metric can be used to help in this process.

RCI is a measure of compliance with ASHRAE/NEBS temperature specifications (recommended and allowable temperatures). This metric is unique in that it not only takes into account the number of equipment temperatures that are outside the recommended range but also with how much. Although, RCI is not a simple ratio between two variables (as with the previous metrics), the interpretation of the metric is straight forward:

- RCI is always $\leq 100\%$
- RCI = 100%: No equipment intake temperature above the recommended range
- RCI < 90%: Signifies poor thermal conditions, possibly with temperatures above the allowable range (risk for IT equipment overheating)

This metric has been incorporated in and is automatically calculated in the U.S.DOE Air Management Tool: <http://datacenters.lbl.gov/tools>.

8.6 Airflow Efficiency

This metric characterizes overall airflow efficiency in terms of the total fan power required per unit of airflow. This metric provides an overall measure of how efficiently air is moved through the data center, from the supply to the return, and takes into account low pressure drop design as well as fan system efficiency.

$\frac{\text{Total Fan Power (Watts)}}{\text{Total Fan Airflow (cubic feet per minute)}}$	Standard	Good	Better
	1.25 W/cfm	0.75 W/cfm	0.5 kW/cfm

8.7 Cooling System Efficiency

There are several metrics that measure the efficiency of an HVAC system. The most common metric used to measure the efficiency of an HVAC system is the ratio of average cooling system power usage (kW) to the average data center cooling load (tons). A cooling system efficiency of 0.8 kW/ton is considered good practice while an efficiency of 0.6 kW/ton is considered a better benchmark value.

$$\frac{\text{Average Cooling System Power (kW)}}{\text{Average Cooling Load (ton)}}$$

Standard	Good	Better
1.1 kW/ton	0.8 kW/ton	0.6 kW/ton

8.8 ISO/IEC 30134

The data center sector uses globally standardized KPIs, which have been defined and published by the International Standards Organization (ISO).

The International Standard ISO/IEC 30134 Information Technology — Data Centers — Key performance Indicators has a series of metrics: Part 1 Overview and general requirements, Part 2: Power usage effectiveness (PUE), Part 3: Renewable energy factor (REF), Part 4: IT Equipment Energy Efficiency for servers (ITEEsv), Part 5: IT Equipment Utilization for servers (ITEUsv), Part 6: Energy Reuse Factor (ERF), Part 7: Cooling Efficiency Ratio (CER), Part 8: Carbon Usage Effectiveness (CUE), Part 9: Water Usage Effectiveness (WUE).

8.9 Data Center Infrastructure Management and Monitoring

Ongoing energy-usage management can only be effective if sufficient metering is in place. There are many aspects to monitoring the energy performance of a data center that are necessary to ensure that the facility maintains the high efficiency that was carefully sought out in the design process. Below is a brief treatment of best practices for data center energy monitoring. For more detail, refer to the Self Benchmarking Guide for Data Center Energy Performance in section 10 Bibliography and Resources.

Energy-efficiency benchmarking goals, based on appropriate metrics, first need to be established to determine which measured values need to be obtained for measuring the data center's efficiency. The metrics listed above provide a good starting point for high-level energy-efficiency assessment. A more detailed assessment could include monitoring to measure losses along the electrical power chain equipment such as transformers, UPS, and PDUs with transformers. (For a list of possible measured values, refer to the Self-Benchmarking Guide for High-Tech Buildings: Data Centers at Lawrence Berkeley National Laboratory's website: <http://hightech.lbl.gov/benchmarking-guides/data.html>).

The accuracy of the monitoring equipment should be specified, including calibration status, to support the level of desired accuracy expected from the monitoring. The measurement range should be carefully considered when determining the minimum sensor accuracy. For example, a

pair of +/- 1.5°F temperature sensors provides no value for determining the chilled water ΔT if the operating ΔT can be as low as 5°F. Electromagnetic flow meters and ultrasonic flow meters are among the most accurate water flow meters available. Three-phase power meters should be selected to measure true root mean square (RMS) power.

Ideally, the Energy Monitoring and Control System (EMCS) and Supervisory Control and Data Acquisition (SCADA) systems provide all of the sensors and calculations required to determine real-time efficiency measurements. All measured values should be continuously trended and data archived for a minimum of one year to obtain annual energy totals. An open protocol control system allows for adding more sensors after initial installation. IT equipment often includes on-board temperature sensors. A developing technology includes a communications interface which allows the integration of the on-board IT sensors with an EMCS.

Monitoring for performance measurement should include temperature and humidity sensors at the air inlet of IT equipment and at heights prescribed by ASHRAE's Thermal Guidelines for Data Processing Environments, 2021. New technologies are becoming more prevalent to allow a wireless network of sensors to be deployed throughout the IT equipment rack inlets.

Supply air temperature and humidity should be monitored for each CRAC or CRAH unit as well as the dehumidification/humidification status to ensure that integrated control of these units is successful.

9 Conclusion

This guide provided an overview of best practices for energy-efficient data center design spanning the categories of IT systems and their environmental conditions, data center air management, cooling and electrical systems, and heat recovery. An effective organization will consider the total cost of ownership for operational efficiency and cost, utilizing different energy-based metrics, sustainability metrics (water and carbon), and related metrics (thermal, airflow, and cooling) to capture a view of the efficiencies at which a data center performs. Utilizing the following approach will lead to an energy-efficient data center and sustainable operations.

9.1 Key Steps to Sustainable Data Centers

FEMP and NREL developed the following approach for optimizing data center sustainability, listed in order of importance:

1. Reduce energy use by making systems as efficient as possible – the associated data center metric is Power Usage Effectiveness (PUE).
 - o **Maximize the IT equipment intake temperature to maximize cooling energy efficiency while ensuring IT equipment thermal guidelines are met to avoid overheating or compromising reliability.**
 - o Use “free” cooling to reduce or eliminate compressor (chiller, direct expansion [DX])-based cooling.
 - o Optimize fan/pump speeds and uninterruptible power supplies.
2. Reuse heat to achieve the lowest Energy Reuse Effectiveness (ERE) metric possible.
 - o **Maximize compute leaving temperature to maximize energy reuse.**
3. Reject as much remaining heat to dry coolers as possible – water savings reflected in data center metric Water Usage Effectiveness (WUE).
 - o **Maximize compute leaving temperature to maximize heat rejected dry to air.**
4. Maximize energy from renewable systems on-site or within grid region – tracked through data center metric Carbon Usage Effectiveness (CUE).
 - o **Work toward the goal of 100% renewable energy 100% of the time.**

9.2 Federal Energy Management Program

FEMP’s mission is to facilitate the federal government’s implementation of sound, cost-effective energy management and investment practices to enhance the nation’s energy security and environmental stewardship: <https://www.energy.gov/femp/federal-energy-management-program>.

FEMP encourages federal agencies and organizations to improve data center energy efficiency, which can offer tremendous opportunities for energy and cost savings. For questions about

FEMP's data center technical assistance, contact Kendall Kam at Kendall.Kam@hq.doe.gov: <https://www.energy.gov/femp/energy-efficiency-data-centers>.

9.2.1 Training and Tools

FEMP has developed a comprehensive training program, Data Center Energy Practitioner (DCEP), which takes approximately 1 to 4 days depending on the training modules selected. The training program includes hands-on information on energy assessments in data centers. It covers all major energy consuming systems (IT equipment, Air Management, Cooling, and Electrical).

DCEP webpage: <http://datacenters.lbl.gov/DCEP>

Center of Expertise (CoE) for Energy Efficiency in Data Centers: <https://datacenters.lbl.gov>

COE Data Center Energy Efficiency Toolkit: <https://datacenters.lbl.gov/tools>

1. What does the Power Usage Effectiveness (PUE) metric evaluate?

- Server speed efficiency
- Energy efficiency of data center infrastructure
- Cost of power supply units
- Carbon emissions from equipment

2. Which equipment is the largest energy consumer in a typical data center?

- Lighting systems
- Cooling towers
- IT equipment
- Water chillers

3. What is the optimal power supply load level for efficiency?

- 10% to 20%
- 40% to 60%
- 80% to 100%
- 90% to 110%

4. What design configuration separates hot and cold air aisles?

- Open-loop layout
- Rear-door cooling
- Convection-based setup
- Hot aisle/cold aisle containment

5. Which factor is most crucial in achieving energy savings in server fans?

- Use of variable speed fans
- Fixed-speed fans
- Operating at low humidity
- Server stacking orientation

6. What storage technique maximizes storage capacity utilization?

- Data mirroring
- RAID-5 allocation
- Synchronous backup
- Thin provisioning

7. What is the benefit of higher chilled water temperatures?

- Increased cooling efficiency
- Higher humidity levels
- Reduced air pressure
- Faster data transfer

8. What metric assesses how much heat is reused from a data center?

- PUE
- ERE
- CUE
- RCI

9. Which ASHRAE class supports operation up to 113°F?

- Class A1
- Class A2
- Class A3
- Class A4

10. What is the goal of air management in data centers?

- Reduce fan sizes
- Prevent mixing of hot and cold air
- Maximize cable lengths
- Enable open-loop cooling

11. What does virtualization help achieve?

- Higher cooling requirements
- Server consolidation and energy savings
- Increase in power usage
- Redundant data duplication

12. Which cooling method uses outside air for free cooling?

- Air-side economizer
- Chilled beam cooling
- CRAC unit cooling
- Liquid immersion cooling

13.What is the function of a CDU in direct liquid cooling?

- Air filtration
- Distributes and conditions cooling fluid
- Manages dry coolers
- Recycles server data

14.Which fan design reduces static pressure in airflow systems?

- Low-pressure drop ducts
- Overhead return diffusers
- Miniature propeller fans
- Cross-ventilation blowers

15.How can cable congestion impact cooling?

- Prevents data loss
- Reduces airflow efficiency
- Speeds up processors
- Increases server utilization

16.What is the benefit of under-floor clearance in raised installations?

- Easier cable installation
- More uniform pressure distribution
- Higher cooling loads
- Lower plenum height

17.What is the typical range of server utilization in enterprise settings?

- 10% to 20%
- 20% to 40%
- 50% to 70%
- 80% to 100%

18.What type of lighting improves energy efficiency in data centers?

- LED with occupancy sensors
- Fluorescent above racks
- Halogen with dimmers
- Incandescent strip lights

19.What improves UPS efficiency most?

- Larger UPS units
- Resistive load testing
- Operating at high load factors
- Constant voltage regulation

20.Which metric compares total CO₂ emissions to IT energy usage?

- CUE
- PUE
- WUE
- ERE