

Data Centers and Their Energy Consumption: Frequently Asked Questions

August 26, 2025

Congressional Research Service

<https://crsreports.congress.gov>

R48646



R48646

August 26, 2025

Martin C. Offutt

Analyst in Energy Policy

Ling Zhu

Analyst in
Telecommunications
Policy

Data Centers and Their Energy Consumption: Frequently Asked Questions

In its simplest form, a data center is a physical facility that houses and runs large computer systems. U.S. data center annual energy use in 2023 (not accounting for cryptocurrency) was approximately 176 terawatt-hours (TWh), approximately 4.4% of U.S. annual electricity consumption that year, according to a report by Lawrence Berkeley National Laboratory. A data center typically contains multiple computer servers, data storage devices, and network equipment that can provide information technology (IT) infrastructure service for organizations to store, manage, process, and transmit large amounts of data. Some projections show that data center energy consumption could double or triple by 2028, accounting for up to 12% of U.S. electricity use.

Roughly one-half or greater of the electric power demand of data centers stems directly from the operation of electronic IT equipment. Much of the rest is for cooling. The operation of the IT equipment raises the temperature of the ambient room air, necessitating a cooling strategy. Centralized cooling resources are of two types: (1) those moving chilled air through large duct work; or (2) those moving chilled water in a piped *cooling loop* that exchanges heat with the environment. An alternative to these centralized systems is room-scale air conditioners. One type, called computer room air conditioners (CRACs), is common in smaller data centers. Exchanging the heat with the environment can happen faster with methods that directly consume water. The source of the water can be the local water utility and can also be on-site reservoirs or other colocated water resources. A study by the International Energy Agency estimates for illustration that a 100-megawatt U.S. data center would consume roughly the same amount of water as 2,600 households, accounting only for direct water consumption and averaged across the various cooling strategies.

Currently there are no legally binding energy standards that apply explicitly to operation of data centers in the private sector. For use within the federal government, the U.S. Department of Energy has published guidance on how to optimize energy use in its data centers. Another non-binding program, Energy Star, certifies data centers with a focus on the building and infrastructure. Since 2012 the Department of Energy has regulated the energy efficiency of CRACs, one type of cooling strategy.

The federal government has made some efforts to gather data using information collection methods suitable for later scale-up. A 2021 report by the U.S. Energy Information Administration on a pilot study of energy use in 50 data centers received 9 respondents. Private firms maintain data sets that can provide direct or proxy information on data centers.

In 2023, a letter from five Senators and three Representatives urged EPA to use its authority under Section 114 of the Clean Air Act to implement a “mandatory disclosure regime” on cryptocurrency mining facilities. In the 119th Congress, S. 1475, the Clean Cloud Act of 2025, introduced in the Senate and referred to the Committee on Public Works, would amend the Clean Air Act (42 U.S.C. §§7401 et seq.) to provide the U.S. Environmental Protection Agency and the U.S. Energy Information Administration with authority to collect data and information on annual electricity consumption of data centers and cryptocurrency mining facilities.

Contents

Introduction	1
What is a data center?	1
What do data centers look like?.....	2
How have data centers been defined in federal law and guidance?	2
What are the energy requirements of data centers?	3
Why are data centers energy-intensive?	3
What equipment uses energy in a data center?	5
What contributes to the need for cooling in data centers?	6
How are data centers typically cooled?	7
What are some additional data center cooling strategies?	8
What are some potential concerns about water consumption by data centers?	8
Which, if any, federal regulations cover the energy consumption of data centers?.....	9
Are there reports on the actual energy use of U.S. data centers?	10

Figures

Figure 1. Notional Power Draws of IT Equipment and Infrastructure	6
---	---

Contacts

Author Information.....	11
-------------------------	----

Introduction

U.S. data center annual energy use in 2023 (not accounting for cryptocurrency) was approximately 176 terawatt-hours (TWh), approximately 4.4% of U.S. annual electricity consumption that year.¹ Some projections show that data center energy consumption could double or triple by 2028, accounting for up to 12% of U.S. electricity use.² Data centers provide information technology (IT) infrastructure services for processing large amounts of data, such as for the rapidly growing field of artificial intelligence (AI). Roughly one-half or greater of the electric power demand of data centers stems directly from the operation of IT equipment. Much of the rest is for cooling.

What is a data center?

In its simplest form, a data center is a physical facility that houses and runs large computer systems. A data center typically contains multiple computer servers, data storage devices, and network equipment that can provide IT infrastructure service for organizations to store, manage, process, and transmit large amounts of data. An increasing number of private and public entities have used data centers to support their expanding IT needs, particularly as data volumes continue to grow.

Types of data centers vary based on their ownership or intended purposes. For example, a large company may choose to build, own, and operate an on-premises data center (also known as an “enterprise data center”) to house and manage its own IT infrastructure.³ Other organizations, especially those lacking the space, staff, or IT resources, often choose to rent a space, equipment, or services within a colocation data center (also known as a “managed data center”) owned and operated by a third-party company.⁴ Some online service providers operate geographically distributed and interconnected data centers and allow multiple users to remotely access computing resources such as data processing chips, software, data storage, networks, and applications hosted by these data centers, which are called “cloud data centers.”⁵ Cloud computing service providers may also operate and maintain smaller data centers physically located closer to end users (called “edge data centers”) to reduce network latency and data communication delay, speed up content distribution, optimize real-time, data-intensive workloads, and improve application performance and user experience.⁶

The ever-increasing demand for data storage and processing capacities, especially for intensive computational tasks such as AI development and deployment, has led to construction and

¹ A. Shehabi et al., *2024 United States Data Center Energy Usage Report*, Lawrence Berkeley National Laboratory (LBNL), LBNL-2001637, December 2024, p. 5. While cryptocurrency is one type of service supported by data centers, not all studies of energy consumption of data centers touch upon energy usage related to cryptocurrency. The report noted that its calculations assumed that data centers would operate consistently with how they were commissioned and designed but that results may differ.

² Shehabi et al., *2024 United States Data Center Energy Usage Report* (LBNL report), p. 6.

³ Equinix, “What Is a Data Center? What Are Different Types of Data Centers?,” August 1, 2024, <https://blog.equinix.com/blog/2022/10/13/what-is-a-data-center-what-are-different-types-of-data-centers/>.

⁴ Stephanie Susnjara and Ian Smalley, “What Is a Data Center?” IBM, September 4, 2024, <https://www.ibm.com/think/topics/data-centers>.

⁵ Cisco, “What Is a Data Center?,” <https://www.cisco.com/c/en/us/solutions/data-center-virtualization/what-is-a-data-center.html#~:infrastructure-evolution>.

⁶ Susnjara and Smalley, “What Is a Data Center?”

operation of “hyperscale data centers” notable for their sheer size.⁷ According to industry analysts, to be considered a hyperscale data center, a facility should contain at least 5,000 computer servers and a large scale of network equipment and occupy at least 10,000 square feet of physical space, with an electric power rating, sometimes referred to as *power draw*, exceeding 100 megawatts (MW; one megawatt is equal to 1 million watts).⁸ Roughly 100 MW of electric power is sufficient to support the electricity needs of 80,000 U.S. households.⁹

What do data centers look like?

Data centers have architectural configurations ranging from closets, to larger rooms within a single enterprise, to dedicated standalone structures serving the needs of multiple customers or tenants, known as *colocation*. *Colocation centers* are touted for their flexibility in allowing customers to specify exactly the amount of hardware and software resources they require. The largest, or “hyperscale,” facilities occupy whole buildings or groups of buildings. When reporting on the number of data centers, some analysts only count whole-building data centers, which would include the hyperscale and colocation centers noted above. The majority of computer servers are found in these latter two types of centers—74% in 2023, according to a Lawrence Berkeley National Laboratory (LBNL) report.¹⁰

How have data centers been defined in federal law and guidance?

The term *data center* has been defined in federal laws and guidance in the context of energy consumption and federal use of data centers. For instance, Section 453(a)(1) of the Energy Independence and Security Act of 2007 (P.L. 110-140) defines a data center as “any facility that primarily contains electronic equipment used to process, store, and transmit digital information, which may be (A) a free-standing structure; or (B) a facility within a larger structure, that uses environmental control equipment to maintain the proper conditions for the operation of electronic equipment.”¹¹

In its guidance for federal agencies to implement the Federal Data Center Enhancement Act of 2023 (P.L. 118-31), the Office of Management and Budget specified that an agency data center covered in the memorandum (1) is composed of certain types of permanent structures and operates in a fixed location; (2) houses IT equipment, including servers and other high-

⁷ Phill Powell and Ian Smalley, “What Is a Hyperscale Data Center?” IBM, March 21, 2024, <https://www.ibm.com/think/topics/hyperscale-data-center>.

⁸ Powell and Smalley, “What Is a Hyperscale Data Center?” See also VIAVI Solutions, “What Is a Hyperscale Data Center?,” <https://www.viavisolutions.com/en-us/resources/learning-center/what-hyperscale-data-center>.

⁹ This calculation assumes a continuous source of electric power generation equal to 100 MW, operating year-round, and that the average U.S. household consumes 10,566 kWh of electricity per year. The latter figure is sourced from the U.S. Energy Information Administration (EIA), “EIA Releases Consumption and Expenditures Data from the Residential Energy Consumption Survey,” press release, March 29, 2023, <https://www.eia.gov/pressroom/releases/press530.php>.

¹⁰ Shehabi et al., *2024 United States Data Center Energy Usage Report* (LBNL report), p. 37.

¹¹ 42 U.S.C. §17112(a)(1).

performance computing devices, or data storage devices; and (3) hosts information and information systems accessed by other systems or by users on other devices.¹²

What are the energy requirements of data centers?

The Department of Energy (DOE) examined the nationwide energy consumption of data centers in response to direction by Congress in the Energy Act of 2020 (Section 1003 of Division Z, P.L. 116-260). The study, performed by Lawrence Berkeley National Laboratory (LBNL), found that U.S. data center annual energy use in 2023 (not accounting for cryptocurrency) was approximately 176 terawatt-hours (TWh), approximately 4.4% of U.S. annual electricity consumption that year.¹³ An analysis by the Electric Power Research Institute (EPRI) similarly estimated that data centers consumed 4% of U.S. electricity in 2023.¹⁴ In a separate analysis, EPRI estimated that AI consumes 10% to 20% of data center energy.¹⁵

Why are data centers energy-intensive?

Roughly one-half or greater of the electric power demand of data centers stems directly from the operation of electronic IT equipment.¹⁶ Much of the rest is for cooling. The core hardware components of a data center include computer servers, which contain computing chips (e.g., central processing unit [CPU] and graphics processing unit [GPU]), memory chips, data storage drives, and network routers and switches.¹⁷ Data from a major CPU chip manufacturer show that its data center-level CPU series in early 2025 had an average thermal design power (TDP) rating between 150 watts (W) and 350W.¹⁸ An advanced data center-level GPU can have a maximum TDP rating between 350W and 700W.¹⁹ The computing chips of CPU and GPU typically consume the most electrical power inside a server, discussed further in “What contributes to the need for cooling in data centers?” According to an industry report published in November 2024,

¹² Office of Management and Budget, *Implementation Guidance for the Federal Data Center Enhancement Act*, M-25-03, January 14, 2025, p. 2, https://bidenwhitehouse.archives.gov/wp-content/uploads/2025/01/M-25-03_Implementation-Guidance-for-the-Federal-Data-Center-Enhancement-Act.pdf.

¹³ Shehabi et al., *2024 United States Data Center Energy Usage Report* (LBNL report), p. 5.

¹⁴ See Figure ES-1 of Electric Power Research Institute (EPRI), *Powering Data Centers: U.S. Energy System and Emissions Impacts of Growing Loads*, October 2024, p. 3, <https://www.epri.com/research/products/000000003002031198>.

¹⁵ EPRI, *Powering Data Centers*, p. 7.

¹⁶ Shehabi et al., *2024 United States Data Center Energy Usage Report* (LBNL report), p. 47.

¹⁷ For more information on the use of CPUs and GPUs in data centers, see CRS In Focus IF12899, *Data Centers and Cloud Computing: Information Technology Infrastructure for Artificial Intelligence*, by Ling Zhu.

¹⁸ Intel, “Intel Xeon 6 Processors,” <https://www.intel.com/content/www/us/en/products/details/processors/xeon.html>. Intel defines *thermal design power* (TDP) as “the average power, in watts, the processor dissipates when operating at Base Frequency with all cores active, under an Intel-defined, high-complexity workload.” Intel, “11th Gen Intel® Core™ Mobile Processor Technical Specifications,” <https://edc.intel.com/preview/content/www/us/en/products/performance/benchmarks/11th-gen-intel-core-mobile-processor-technical-specifications/>.

¹⁹ See, for example, Nvidia, “NVIDIA H100 Tensor Core GPU,” <https://www.nvidia.com/en-us/data-center/h100/>. Nvidia defines the term TDP differently from Intel. The TDP of an Nvidia GPU is “the maximum power that a subsystem is allowed to draw for a ‘real world’ application, and also the maximum amount of heat generated by the component that the cooling system can dissipate under real-world conditions.” Nvidia, *GeForce GPU Power Primer*, https://www.nvidia.com/content/dam/en-zz/Solutions/GeForce/technologies/frameview/Power_Primer.pdf.

computing power and server systems account for roughly 40% of electricity consumption in a data center, while network and data storage equipment use about 10%.²⁰

Each piece of the electronic IT equipment generates heat as it operates. Many chipsets incorporate a safety mechanism called “thermal throttling” that reduces the chip performance to prevent overheating and protect the hardware.²¹ Data centers require cooling systems to help dissipate the heat and maintain optimal performance and overall system stability. The cooling systems could account for another 38% to 40% of electricity consumption in a data center.²² (See **Figure 1.**)

While general-purpose workloads typically require only CPUs, GPUs are considered better than CPUs for handling computation-intensive processes such as AI training.²³ Under full workload conditions, a GPU performing AI training tasks may operate near its maximum capacity and draw power close to its maximum TDP over extended periods of time.

The development and deployment of state-of-the-art, large AI models may even require multiple GPUs to work concurrently by distributing large volumes of data and computational tasks across the computing chips (known as “parallel computing”).²⁴ A study released in December 2024 observed that, when training a large AI model using a computer system with eight advanced GPUs for eight hours, the GPUs were near full utilization most of the time (an average of 93%) and the median amount of electrical power consumed by the chips was 7.92 kilowatts (kW; one kilowatt is equal to 1,000 watts), with a total energy consumption of 62 kW-hours (kWh).²⁵

Cutting-edge chip technologies support high-speed GPU-to-GPU data communication among hundreds of GPUs across multiple servers, enabling the creation of a massive data processing cluster to support large-scale AI training but also further increasing these data centers’ energy consumption. A report released in April 2025 estimated that training a specific large AI model required a total power draw of 25.3 MW and that the power required to train these models could double annually.²⁶ The report stated that “[t]he rising power consumption of AI models reflects the trend of training on increasingly larger datasets.”²⁷ Another study released in May 2025 estimated that training another large AI model consumed 50 gigawatt-hours (GWh; one gigawatt is equal to 1 billion watts) of energy, “enough to power San Francisco for three days.”²⁸

Multiple industry reports indicate that data processing demands of AI and related remote computing services have spurred new construction and upgrades of data centers. The computing

²⁰ Karthik Ramachandran et al., “As Generative AI Asks for More Power, Data Centers Seek More Reliable, Cleaner Energy Solutions,” Deloitte Center for Technology Media and Telecommunications, November 19, 2024, <https://www2.deloitte.com/us/en/insights/industry/technology/technology-media-and-telecom-predictions/2025/genai-power-consumption-creates-need-for-more-sustainable-data-centers.html>.

²¹ Intel, “What Is Throttling and How Can It Be Resolved?” May 25, 2023, <https://www.intel.com/content/www/us/en/support/articles/000088048/processors.html>.

²² Ramachandran et al., “As Generative AI Asks for More Power, Data Centers Seek More Reliable, Cleaner Energy Solution.”

²³ Imran Latif et al., “Empirical Measurements of AI Training Power Demand on a GPU-Accelerated Node,” arXiv, December 20, 2024, <https://arxiv.org/abs/2412.08602>.

²⁴ DigitalOcean, “Multi-GPU Computing: What It Is and How It Works,” February 14, 2025, <https://www.digitalocean.com/resources/articles/multi-gpu-computing>.

²⁵ Latif et al., “Empirical Measurements of AI Training Power Demand on a GPU-Accelerated Node,” p. 10.

²⁶ Nestor Maslej et al., *The Artificial Intelligence Index Report 2025*, Institute for Human-Centered Artificial Intelligence, Stanford University, April 2025, p. 72, https://hai-production.s3.amazonaws.com/files/hai_ai_index_report_2025.pdf.

²⁷ Maslej et al., *AI Index Report 2025*, p. 72.

²⁸ James O’Donnell and Casey Crownhart, “We Did the Math on AI’s Energy Footprint. Here’s The Story You Haven’t Heard,” *MIT Technology Review*, May 20, 2025.

resources hosted by these centers would, in turn, lead to increased power demand. For example, one report estimated that the computing capacity (measured by the amount of electrical power consumed by IT equipment) of data centers under construction in North America at the end of 2024 reached a record-high 6,350 MW, more than double the figure from a year earlier.²⁹ Another report indicated that new hyperscale data centers have been built with capacities from 100 MW to 1,000 MW each, “roughly equivalent to the load from 80,000 to 800,000 homes.”³⁰

What equipment uses energy in a data center?

Data centers in whole building structures contain energy-consuming IT equipment, cooling and air handling equipment, and backup power supplies.³¹ The latter may include uninterruptible power supplies and backup diesel generators.³² Backup strategies involving batteries consume electricity while improving the quality of the electricity by evening out the highs and lows of the electric voltage.

For a small data center within an office or research building, the mix of energy consumption could be roughly 50% attributable to IT physical machines and 50% attributable to cooling and power supply, according to a large industrial and energy equipment maker.³³ (See **Figure 1**.) In addition, data centers include appliances and equipment for human occupants, not depicted in the figure. Overall, the energy use within data centers is not well described nationally, as elaborated further in “Are there reports on the actual energy use of U.S. data centers?” An expression or metric sometimes used to describe the energy performance of data centers is *power utilization effectiveness* (PUE), which is the ratio of all the power used by the data center to the power used by just the IT equipment.³⁴ For the example just discussed, the PUE would be two (2).

The LBNL report estimated energy use in data centers by assigning energy consumption to servers and arriving at a total energy consumption value based on an inventory of the number of servers and an assumed activity rate of the computing resources.³⁵ According to this estimate, U.S. data center annual energy use in 2023 (not accounting for cryptocurrency) was approximately 176 terawatt-hours (TWh), approximately 4.4% of U.S. annual electricity consumption that year.³⁶

Discussion of IT equipment rapidly invokes terms descriptive of nonphysical entities such as clouds, virtual machines, architectures, nodes, and protocols not directly associated with a real-

²⁹ CBRE, “North America Data Center Trends H2 2024: Surging Demand Drives Record New Data Center Development,” February 26, 2025, <https://www.cbre.com/insights/reports/north-america-data-center-trends-h2-2024>.

³⁰ EPRI, *Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption*, May 28, 2024, p. 2.

³¹ These non-IT components may be termed “infrastructure” for purpose of energy tracking. A. Shehabi et al., “Data Center Growth in the United States: Decoupling the Demand for Services from Electricity Use,” *Environmental Research Letters*, vol. 13 (2018), p. 124030. Whole-building data centers are described elsewhere in the present report.

³² UPS strategies range from full standby to active regeneration of the waveform of the power supply. CyberPower, “How Does an Uninterruptible Power Supply (UPS) Work?,” September 10, 2015, <https://www.cyberpowersystems.com/blog/how-does-a-ups-work/>.

³³ ABB, *HVAC Motors: Motors in Data Centers*, <https://new.abb.com/motors-generators/nema-low-voltage-ac-motors/hvac-motors>.

³⁴ The Green Grid, *PUE: A Comprehensive Examination of the Metric*, White Paper #49, 2012; and N. Casale, “Data Centers: Where and How Should PUE Be Improved?” *ASHRAE Journal*, vol. 63, no. 6 (June 2021).


³⁵ Shehabi et al., *2024 United States Data Center Energy Usage Report* (LBNL report), p. 14.

³⁶ Shehabi et al., *2024 United States Data Center Energy Usage Report* (LBNL report), p. 5. The report noted that its calculations assumed that data centers would operate consistently with how they were commissioned and designed and that this is often not the case in hindsight.

world energy-consuming process. Understanding the energy consumption necessitates discussing hardware (i.e., physical machines) to which energy consumption can be assigned, described below in “What contributes to the need for cooling in data centers?”

Figure 1. Notional Power Draws of IT Equipment and Infrastructure

Illustrative case with PUE = 2, consistent with data centers of <150 square feet floor space



Source: CRS, adapted from ABB, *HVAC Motors: Motors in Data Centers*, <https://new.abb.com/motors-generators/nema-low-voltage-ac-motors/hvac-motors>.

Notes: PUE = power usage effectiveness, the ratio of all power used by a data center to that used by the information technology (IT) equipment (i.e., by servers, data storage, and communication). The depicted PUE of 2 is typical of a small data center, i.e., one having less than 150 square feet of floor space, as reported in A. Shehabi et al., *2024 United States Data Center Energy Usage Report*, Lawrence Berkeley National Laboratory (LBNL), LBNL-2001637, December 2024, p. 47.

What contributes to the need for cooling in data centers?

From an electrical engineer’s perspective, as computer chips of various types perform their function, the impedance of the electrical pathways generates heat, both on and between chips.³⁷ The majority of heat arises dynamically while the processors are performing useful functions. Heat-generating activities from computation take place related to, or in support of, CPU/GPU activity. Generally speaking, the energy consumption of a server scales according to the number of CPUs or GPUs.³⁸ On-chip (i.e., on the physical CPU/GPU) energy consumption includes dynamic losses associated with the basic concept of operation of a transistor, known as *switching*.³⁹ Further losses occur associated with computing chips when they access on-chip SRAM (static random access memory) and networking functions. Additional losses on-chip are due to control functions (“clock”), power, and current leakage.⁴⁰ Off-chip energy can be lost when electricity moves between the computing chip (CPU/GPU) and a separate memory chip. The time

³⁷ Servers, named after today’s client-server architecture, may include chips for computing, memory, and networking.

³⁸ Shehabi et al., *2024 United States Data Center Energy Usage Report* (LBNL report), p. 16.

³⁹ Integrated circuits, also known as chips, contain billions of transistors. Heat is the movement of energy between two bodies of different temperatures that are in thermal contact. Colloquially, heat may be referred to as a “loss” if the heat does not provide a useful (valorized) service.

⁴⁰ Leakage current is a static loss not associated with the chip’s performance of computational functions. Not all the energy consumed by the IT equipment performs the useful computing functions of the computer chips. Some energy heats the equipment without performing a function.

needed for very large transmission of data between separate chips, as in the last example, is known as the *memory wall* or *memory bottleneck*.

Unlike in a desktop computer, the activity rates of chips in a data center can be extremely high, and this activity rate increases the cooling needs as the hot equipment raises the temperature of the ambient air. The more common cooling strategies treat and condition this ambient air by lowering its temperature and humidity.

Historically the amount of computing power per watt has improved significantly. The move to multi-core processing at the turn of the millennium lowered the heat produced for the same amount of computing power.⁴¹ According to Nvidia, the last 10 years have seen a 4,000-fold improvement in the GPU's computational performance per watt of power.⁴² More conservatively, the International Energy Agency estimates the change in GPU performance per watt to have improved 100-fold or greater between 2008 and 2023.⁴³

How are data centers typically cooled?

The operation of the IT equipment raises the temperature of the ambient room air, necessitating a cooling strategy. Generally speaking, the cooling strategies for IT equipment differ from cooling needed for personal comfort, with computer servers tolerant of higher temperatures but requiring lower humidity.⁴⁴ Federal guidelines and research advise using large centralized cooling resources for data centers.⁴⁵ There are two types of centralized cooling resources: (1) those moving air through large ductwork to deliver chilled air and remove warm air; or (2) those moving water or other heat transfer fluid through a piped *cooling loop* that exchanges heat with the environment. The centralized cooling resources achieve higher efficiency at the larger scale. Those of the first type, which moves air, can improve their energy efficiency by utilizing variable speed fans.⁴⁶ Though less common, these methods can, if required, be used to heat the interior of the data center using air-ducted or looped strategies (i.e., that use heat transfer fluids).

An alternative to these centralized systems is room-scale air conditioners. One type, called computer room air conditioners (CRACs), is common in smaller data centers.⁴⁷ With CRACs, the air is looped and filtered within the room but the heat is sent outside the building using refrigerant or other fluid.⁴⁸

⁴¹ K. Bourzac, "Fixing AI's Energy Crisis," *Nature*, October 17, 2024.

⁴² Bourzac, "Fixing AI's Energy Crisis."

⁴³ International Energy Agency (IEA), "Efficiency Improvement of AI Related Computer Chips, 2008-2023," October 17, 2024, <https://www.iea.org/data-and-statistics/charts/efficiency-improvement-of-ai-related-computer-chips-2008-2023>.

⁴⁴ 65 *Federal Register* 48830 (August 9, 2000).

⁴⁵ DOE Federal Energy Management Program (FEMP) and National Renewable Energy Laboratory (NREL), *Best Practices Guide for Energy-Efficient Data Center Design*, July 2024, p. 14, https://www.energy.gov/sites/default/files/2024-07/best-practice-guide-data-center-design_0.pdf.

⁴⁶ DOE FEMP, *Best Practices*, p. 16.

⁴⁷ In DOE's definition, CRACs are "[u]sed in computer rooms, data processing rooms, or other information technology." DOE FEMP, *Best Practices*, p. 14; and DOE, *Technical Support Document: Energy Efficiency Program for Commercial and Industrial Equipment: Certain Categories of Commercial Air Conditioning and Heating Equipment*, August 2019, p. 2-1.

⁴⁸ Trane, *Engineers Newsletter: Understanding the Selection of Direct Expansion (DX)*, vol. 52, no. 1 (March 2023), <https://www.trane.com/content/dam/Trane/Commercial/global/learning-center/engineers-newsletters/ADM-APN086-EN.pdf>; R. Waldron, "CRAC Units: Computer Room AC Basics," Rasmussen Mechanical Services, January 25, 2023, (continued...)

The above methods rely on a sequence in which computing equipment first heats the room air, after which the cooling system accepts heat from the room air, either moving a heated fluid outside, where the heat is returned to the environment, or removing the air itself and replacing it with cooled return air. Exchanging the heat with the environment can happen faster with methods that directly consume water.

What are some additional data center cooling strategies?

High-performance computing (HPC)⁴⁹ equipment has necessitated cooling methods that are thermodynamically closer to the chips and intercept the energy before it has substantially raised the room air temperature. These *direct liquid cooling* technologies can address the higher power of HPC.

Another option that may be used during shoulder season (i.e., spring and fall in temperate climates) or winter, known as *free cooling*, imports water chilled passively by environmental conditions. Compared with methods involving conventional *mechanical cooling* that use powered compressors, free cooling is relatively inexpensive. Optimizing the cooling of the data center will generally involve some combination of methods.

Facilities with substantial cooling needs sometimes employ thermal storage. The National Institutes of Health, for example, has a large chilled water storage facility at its Bethesda, MD, campus.⁵⁰ The thermal storage system operates by drawing down the chilled water reservoir during periods of intense cooling requirements and replenishing the reservoir during off-peak hours.

What are some potential concerns about water consumption by data centers?

A study by the International Energy Agency estimates for illustration that a 100 MW U.S. data center may consume roughly the same amount of water as 2,600 households, accounting only for direct water consumption and averaged across the various cooling strategies.⁵¹ One study

https://www.rasmech.com/blog/crac-units-computer-room-ac-basics/?srsltid=AfmBOooCLecm_NVPaRw7UJw8Qby5OKw2mKN73EF9CyG-a2spbHTuyBz-; and R. Schmidt and M. Iyengar, “Thermodynamics of Information Technology Data Centers,” *IBM Journal of Research and Development*, vol. 53, no. 3 (August 2009).

⁴⁹ The definition of high-performance computing is not rigorous but is associated with more complex tasks. See, for example, Nvidia, “What Is High-Performance Computing?,” <https://www.nvidia.com/en-us/glossary/high-performance-computing/>.

⁵⁰ National Academies of Sciences, Engineering, and Medicine, *Managing the NIH Bethesda Campus Capital Assets for Success in a Highly Competitive Global Biomedical Research Environment* (National Academies Press, 2019), p. 42.

⁵¹ The IEA estimates water use of data centers to include 60% indirect (at power plants) and 40% direct water use, with the sum equal to the water consumption of 6,500 households. The direct water use would thus be equivalent to the water consumption of 2,600 households. IEA, *Energy and AI*, World Energy Outlook Special Report, April 2025, p. 242, <https://iea.blob.core.windows.net/assets/34eac603-ecf1-464f-b813-2ecceb8f81c2/EnergyandAI.pdf>.

estimated that data centers use roughly seven cubic meters of water per megawatt-hour (MWh) of energy.⁵²

Of the various cooling strategies to return heat to the environment, those that include cooling towers accelerate the rate of heat transfer by spraying water onto surfaces in the cooling tower. The water evaporates to provide cooling and must constantly be replenished. According to one vendor, water-cooled systems may use as little as one-half the electric power of air-cooled systems.⁵³

The source of the cooling water can be the local water utility. A city in Oregon found that nearly 30% of the city's water consumption was attributable to Google data centers, which had tripled water consumption over a five-year period.⁵⁴ The source of water can also be on-site reservoirs or other colocated water resources.⁵⁵ Data centers of smaller size within an office building might only incrementally increase direct water use.

There are two principal sources of water demand from cooling towers. One of the sources of demand is the sprayed-on cooling water just noted. Another water-consuming operation, known as *blowdown*, uses water to flush out hard scale that can build up on surfaces in the cooling tower when sprayed-on water evaporates and leaves behind solid material.⁵⁶ By one estimate, most of the wastewater produced at a data center is from blowdown.⁵⁷

Which, if any, federal regulations cover the energy consumption of data centers?

DOE has published guidance on how to optimize energy use in data centers used by the federal government.⁵⁸ Currently there are no legally binding energy standards that apply explicitly to operation of data centers in the private sector.⁵⁹ The Energy Star program, a voluntary labeling program, certifies data centers with a focus on the energy efficiency of data centers' buildings and infrastructure.⁶⁰ Energy Star uses a calculation method that effectively removes the consumption of the IT equipment, attributing energy use only to that of a *powered shell* (the

⁵² M. A. B. Siddik et al., "The Environmental Footprint of Data Centers in the United States," *Environmental Research Letters*, vol. 16, no. 64017 (2021), p. 4.

⁵³ ABB, *HVAC Motors: Motors in Data Centers*, <https://new.abb.com/motors-generators/nema-low-voltage-ac-motors/hvac-motors>.

⁵⁴ M. Rogoway, "Google's Water Use Is Soaring in the Dalles, Records Show, with Two More Data Centers to Come," *Oregonian*, February 22, 2023.

⁵⁵ R. Miller, "Alligator Patrols Google's Data Center," *DataCenter Knowledge*, December 13, 2012, <https://www.datacenterknowledge.com/hyperscalers/alligator-patrols-google-s-data-center>.

⁵⁶ DOE FEMP, "Best Management Practice #10: Cooling Tower Management," <https://www.energy.gov/femp/best-management-practice-10-cooling-tower-management>.

⁵⁷ DOE FEMP, "Best Management Practice #10: Cooling Tower Management," <https://www.energy.gov/femp/best-management-practice-10-cooling-tower-management>.

⁵⁸ DOE FEMP, *Best Practices Guide*.

⁵⁹ Building codes and standards vary by jurisdiction. Generally, the federal government has jurisdiction over building codes and standards for federal and military buildings; establishes national manufactured housing construction standards; and requires buildings to comply with the Americans with Disabilities Act (ADA; P.L. 101-336) of 1990 and the Fair Housing Act of 1968 (P.L. 90-284). For more information, see CRS Report R47665, *Building Codes, Standards, and Regulations: Frequently Asked Questions*, coordinated by Linda R. Rowan.

⁶⁰ For further information on Energy Star, see CRS In Focus IF10753, *ENERGY STAR Program*, by Corrie E. Clark.

building and its electricity).⁶¹ Energy Star gives a rating based on the performance of the powered shell of the building and rates this against that of similar buildings. The program has certified nearly 300 data centers, with all but one having greater than 10,000 square feet total floor area.⁶²

Energy Star also certifies products found in data centers such as enterprise scale servers, uninterruptible power supplies, memory storage, and networking equipment.⁶³

Since 2012, DOE has regulated the energy efficiency of CRACs.⁶⁴ The regulatory program is authorized by the Energy Policy and Conservation Act, as amended (EPCA; 42 U.S.C. §§6291 et seq.). The standards apply to units at the time they are shipped by manufacturers; the manufacturer is responsible for compliance.⁶⁵

Are there reports on the actual energy use of U.S. data centers?

Owing to recent interest in the energy use of data centers, the federal government has begun data collection to assess their scope and scale and, ultimately, their impacts on energy demand and natural resources. The Energy Information Administration (EIA) has conducted two data collection activities but these were limited by sample size or were curtailed. Congress has shown interest in the results of data collection and in expanding the federal government's authority to collect the data. As noted, the federal government has made some efforts to gather data using information collection methods suitable for later scale-up.⁶⁶ A 2021 EIA report on a pilot study of energy use in 50 data centers received 9 respondents; other responses were either not received, were incomplete, or were not facilities matching the criteria for data centers. EIA noted significant obstacles to collecting the data.⁶⁷ EIA focused on facilities of greater than 50,000 square feet as most likely to be whole building data centers.⁶⁸

In 2024, EIA attempted to collect information focused on cryptocurrency mining data centers. EIA began its survey with emergency approval under the Paperwork Reduction Act of 1980 (PRA; P.L. 96-511) but ceased data collection following an agreement filed with a federal court in Texas in which EIA agreed to destroy any data collected up to that point.⁶⁹

⁶¹ Energy Star, *Data Center Estimates in the United States and Canada*, August 2023, p. 1, https://www.energystar.gov/sites/default/files/tools/Data_Center_Estimates_August_2018_EN%20-%20508%20Blue.pdf. The whole-building certification by Energy Star's computation tool, Portfolio Manager, takes account of electricity and natural gas used on site and calculates an energy use intensity (energy per floor area) that includes the energy needed to generate and deliver the electricity and to deliver the natural gas.

⁶² Energy Star, *Energy Star Certified Data Centers*, <https://www.energystar.gov/buildings/certified-data-centers>.

⁶³ Energy Star, "Data Center Equipment," https://www.energystar.gov/products/data_center_equipment.

⁶⁴ The Energy Policy Act of 1992 (P.L. 102-486) amended the Energy Policy and Conservation Act (42 U.S.C. §6291 et seq.) and gave DOE authority to set energy conservation standards for "commercial package air conditioning and heating equipment." 42 U.S.C. §§6311(8)(A) and 6313(a)(6)(A).

⁶⁵ See CRS Report R47038, *The Department of Energy's Appliance and Equipment Standards Program*, by Martin C. Offutt.

⁶⁶ For further discussion of federal policies on information collection, see CRS Report R48546, *The Office of Information and Regulatory Affairs (OIRA): Overview and Major Responsibilities*, coordinated by Meghan M. Stuessy and Taylor N. Riccard.

⁶⁷ EIA, *2018 CBECS Data Center Pilot*, p. 2.

⁶⁸ *Ibid.*, p. 3.

⁶⁹ Notice of Agreement, *Tex. Blockchain Council v. Dep't of Energy*, No. 6:24-cv-99 (W.D. Tex. Mar. 1, 2024), ECF No. 24.

Private firms maintain data sets that can provide direct or proxy information on data centers. Cushman & Wakefield, a real estate firm with a specialty in data centers, and Baxtel, a real estate and marketing firm that provides specialized services for data center owners, both maintain data sets.⁷⁰

S. 1475, the Clean Cloud Act of 2025, introduced in the Senate and referred to the Committee on Environment and Public Works, would amend the Clean Air Act (42 U.S.C. §§7401 et seq.) to provide the U.S. Environmental Protection Agency (EPA) and EIA with authority to collect data and information on annual energy consumption of the data center or cryptocurrency mining facility, the provider of the electricity, any power purchase agreements, and related topics.

In 2023, a letter from five Senators and three Representatives urged EPA to use its authority under Section 114 of the Clean Air Act to implement a “mandatory disclosure regime” on cryptocurrency mining facilities.⁷¹ The letter further noted that some such facilities would emit greater than the 25,000-ton threshold of carbon dioxide equivalent of greenhouse gases per year necessary to invoke Section 114, an authority that allows EPA to request information from emission source categories for regulatory development or enforcement and for other purposes.⁷²

Author Information

Martin C. Offutt
Analyst in Energy Policy

Ling Zhu
Analyst in Telecommunications Policy

Acknowledgments

Mari Lee prepared the graphic.

Disclaimer

This document was prepared by the Congressional Research Service (CRS). CRS serves as nonpartisan shared staff to congressional committees and Members of Congress. It operates solely at the behest of and under the direction of Congress. Information in a CRS Report should not be relied upon for purposes other than public understanding of information that has been provided by CRS to Members of Congress in connection with CRS’s institutional role. CRS Reports, as a work of the United States Government, are not subject to copyright protection in the United States. Any CRS Report may be reproduced and distributed in its entirety without permission from CRS. However, as a CRS Report may include copyrighted images or material from a third party, you may need to obtain the permission of the copyright holder if you wish to copy or otherwise use copyrighted material.

⁷⁰ EPRI, *Powering Data Centers*, p. 9; and Cushman & Wakefield, “Global Data Center Market Comparison,” 2024, <https://cushwake.cld.bz/2024-Global-Data-Center-Market-Comparison/20/>.

⁷¹ Letter from Sens. Warren, Whitehouse, Markey, Merkley, and Durbin and Reps. Huffman, Tlaib, and Porter to Michael Regan, Administrator, U.S. Environmental Protection Agency, and Jennifer Granholm, Secretary, Department of Energy, February 6, 2023.

⁷² 42 U.S.C. §7414.