SUSTAINABLE DATA CENTERS ROADMAP

CHAPTER 1

Data Center Energy Use

Eric Masanet

October 2025

Innovation for Cool Earth Forum

1 Data Center Energy Use

Eric Masanet

A. Background 2
B. How Much Energy Do Data Centers Currently Use? 15
C. Where Is Data Center Energy Use Headed? 19
D. Recommendations 25
E. References 27

A. Background

i. What is a data center?

Data centers can be thought of as the brains of the internet. These facilities process, store and communicate digital information, providing myriad services including website hosting, email, online shopping, banking, stock trading, maps and navigation, gaming and much more. According to the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE), a data center can be defined as any dedicated space with information technology (IT) equipment loads greater than 10 kW and a floor area with an IT power density greater than 20 W/ft^2.[a,1]

In recent years, the data center sector has expanded rapidly to provide an increasingly salient societal service: the computations that make artificial intelligence (AI) and its many applications possible.[2,3] Data centers range vastly in size from small rooms with a few racks of IT equipment requiring tens of kW to massive dedicated buildings with thousands of racks requiring many hundreds of MW (Figure 1-1).

a While most cryptocurrency mining centers meet this definition, they are outside the scope of this report (see Box 1-2).

While approaches vary, data centers are commonly categorized as follows[2]:

- **Internal data centers:** Run by organizations internally, for their own use. These can range in size from large standalone buildings to small rooms and closets located in existing buildings. Such data centers can also be referred to as enterprise data centers.

- **Retail data centers:** Include both co-location and wholesale data centers built by real estate companies. Co-location data centers host multiple tenants, each of whom leases dedicated space to install and operate their own IT equipment, with cooling, power and other services provided by the data center. A wholesale data center is leased in its entirety to a single tenant.

- **Cloud data centers:** Offer cloud services (i.e., hosted software, hardware, infrastructure and platforms) to consumers who, instead of building and maintaining these systems themselves, pay the cloud provider for remote access and usage.

- **Edge data centers:** Tend to be smaller data centers located close to consumers to reduce latency.

- **High-performance computing (HPC) data centers:** Generally dedicated to hosting supercomputers for scientific computing purposes.

- **Hyperscale data centers:** Built and operated by companies that deploy internet services and platforms at massive scales. Examples include Google, Amazon Web Services (AWS), Microsoft, Oracle, Meta and Apple.[4]

In reality, these categories can have significant overlap. For example, Meta is classified as a hyperscaler but operates its data centers in internal fashion.[5,6] Many cloud providers are also hyperscalers, such as Google Cloud, Microsoft Azure or AWS, and these companies can also operate using space leased at co-location data center facilities.[7] Spaces within co-location data centers can be operated as edge data centers by some clients. Still, such classifications are typically useful for understanding key energy and environmental attributes of data center operators, such as greenhouse gas emissions from power consumption or type(s) of cooling technologies associated with each category[8,9] (see Chapter 2.3 of this Roadmap).

Figure 1-1. *The many shapes and sizes of data centers: (a) Meta's hyperscale data center campus in Odense, Denmark, which contains more than 200,000 m2 of floor area[10]; (b) an example of a tenant-leased space in a co-location data center with secured access[11]; (c) a high-performance computing (HPC) data center occupying around 13,000 m[12,13]; and (d) a small data center room in an office building, occupying less than 100 m[14].*

The total number of data centers globally is unknown. This uncertainty is largely attributable to internal/enterprise data centers that are operated mostly by non-tech companies and have little visibility outside of the organization. For example, recent estimates by Lei et al. (2024)[15] suggest there may be around 1.4 million small- and medium-sized internal data centers operating in the United States alone. However, such data centers are also expected to host a small share of the overall total number of data center IT racks compared to—typically much larger—retail, cloud and hyperscale data centers. For example, Lawrence Berkeley National Laboratory (LBNL) estimated that all types of internal data centers represented only around 15% of servers installed in the United States in 2023, which is currently the world's largest data center market.[2,3] Conversely, LNBL estimated that around 75% of all US servers were located in co-location, cloud and hyperscale data centers that same year. Therefore, when it comes to energy use, while internal data centers are still important (especially from the firm-level perspective), the retail, cloud and hyperscale markets will drive national and global data center energy use moving forward.

While estimates vary, more is known about the number of retail, cloud and hyperscale data centers globally. For example, Synergy Research Group (2025) estimated that the number of large data centers operated by hyperscale providers increased to 1136 at the end of 2024. More than half of these are located in the United States. Online trackers of co-location data centers put their current global numbers around 6600-9000 globally.[16-18] However, these numbers may change quickly, as the sector is expanding rapidly in response to the ongoing AI boom.[19] For example, Synergy Research Group (2025) expects to see 130-140 additional hyperscale data centers coming online each year for the next several years.

ii. How do data centers use energy?

The energy use of data centers can be roughly divided into two major categories: IT equipment and infrastructure equipment (Figure 1-2).[20] IT equipment energy use refers to the electricity consumed by data center IT equipment racks, which contain servers, storage arrays and/or network switches (Figure 1-3).[21] Servers provide the computations and memory necessary for software applications, such as websites, online ordering, and AI model training and inference. Storage arrays are used for storing files, such as documents, databases and streaming videos, and can comprise hard disk drives (HDDs), solid-state drives (SSDs) and even tape storage drives (for archival storage). Network switches enable data communications between racks and to and from the data center via the global internet.

Infrastructure energy use refers to energy consumed by non-IT equipment that provides space conditioning and ensures reliable power within the data center. Space conditioning involves cooling to keep the IT equipment from overheating and humidification or dehumidification to avoid electrostatic discharge, corrosion and condensation that would damage the IT equipment. Data centers use many different types of cooling equipment, the choice of which depends on local factors. Such factors include ambient climates, energy and water costs, building envelope, and rack power densities,[9,22] which are further discussed in Chapter 2.3 of this Roadmap.

Power provision equipment typically consists of transformers and switchgear, uninterruptible power supplies (UPSs), power distribution units (PDUs), and backup generators that provide power during grid outages. All infrastructure equipment runs on electricity except backup power generators, which require fuels such as diesel or natural gas for operation. (See Chapter 3.1 of this Roadmap.)

Figure 1-2. Schematic diagram of data center electricity use (derived from Brown et al., 2007[20])

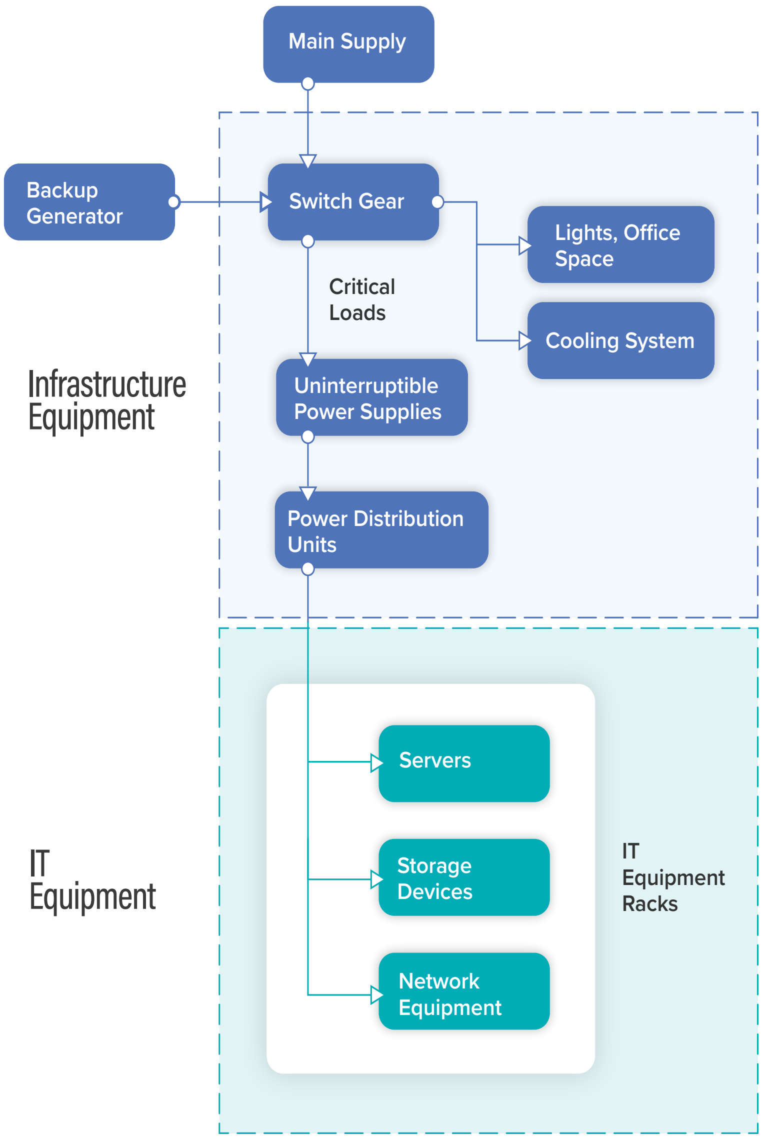

Typical data center electrical components

Ideally, most electricity consumed by a data center would be used by the IT equipment itself, which performs the main revenue-earning function of the data center. When infrastructure uses a high fraction of the total electricity consumed, a data center is considered relatively inefficient. This concept is captured by a key metric used to assess data center efficiency: power usage effectiveness (PUE), which is defined as the ratio of a data center's total energy usage to the energy used by its IT equipment (see Box 1-2).

Figure 1-4 provides a breakdown of the most recent estimates of IT equipment electricity use versus infrastructure electricity use for all data centers in world regions[2,3,23,24] for which such estimates are available. Note that these types of data will typically be estimates due to widespread lack of empirical and reported data on data center energy use, which is a topic addressed later in this chapter. On average for these regions, infrastructure is estimated to account for 30-33% of electricity use, whereas IT equipment is estimated to account for 67-70% electricity use.

However, these proportions can differ substantially at individual data centers depending on their purpose, types of IT and cooling equipment, efficiency levels, climate and other factors, as shown in Figure 1-5.[3]

The share of overall data center electricity use attributable to infrastructure has steadily declined over the past two decades. This decline is primarily due to concerted efforts by data center operators to improve the energy efficiency of their cooling systems (as discussed in Chapter 2.3 of this Roadmap) and to adopt more efficient electrical equipment, such as high-efficiency UPS systems. For example, in 2007 LBNL estimated that infrastructure equipment consumed as much power as IT equipment in

Figure 1-3. *A typical data center information technology (IT) equipment layout with individual IT racks arranged in rows. In this depiction, each row contains 13 racks.*

US data centers, resulting in an estimated nationwide average PUE of 2.0. However, by 2023, thanks to the aforementioned efficiency gains, infrastructure equipment consumed less than half as much power as IT equipment, with an estimated nationwide average PUE around 1.43.[2,20]

Servers generally account for the vast majority of energy use on the IT equipment side (Figure 1-5). The energy use characteristics of each major type of IT and infrastructure equipment are discussed briefly below.

Figure 1-4. *Estimated average IT equipment vs. infrastructure equipment energy use for different regions and countries.*

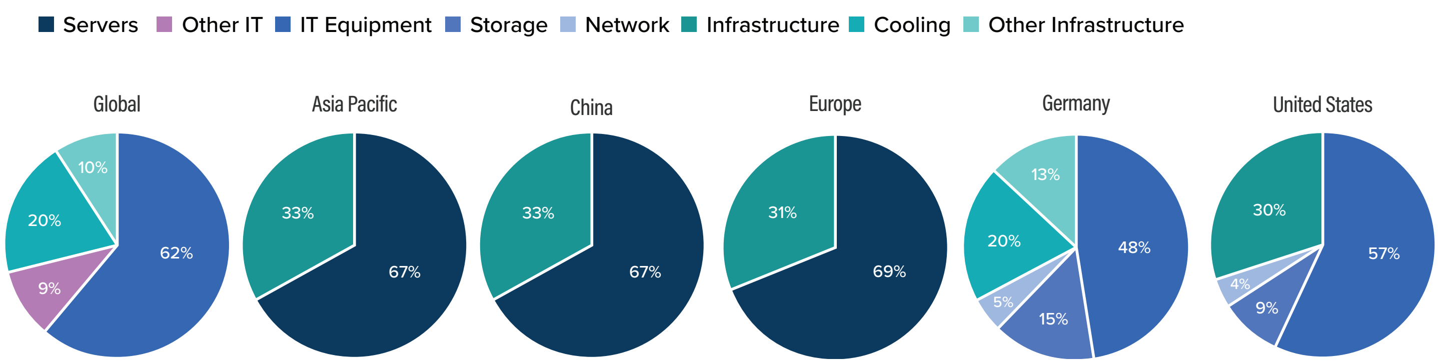

ii(a). Servers

Servers are generally the single largest consumer of electricity in most data centers. While there are several different types of servers used in data centers, over the years, the average power requirements of servers has been growing steadily as has the total number of servers installed in data centers globally. For example, LBNL estimates that the total stock of servers in the United States doubled from 2014 (around 15 million) to 2023 (around 30 million).[2,25] This implies a compound annual stock growth rate (CAGR) of around 8%. While absolute stock numbers are not available at the global level, IEA estimated that the worldwide stock of servers grew at an even faster CAGR of 10% between 2015 and 2024.[3]

Figure 1-5. International Energy Agency (IEA) estimates of shares of electricity consumption by data center and equipment type, 2024[3]

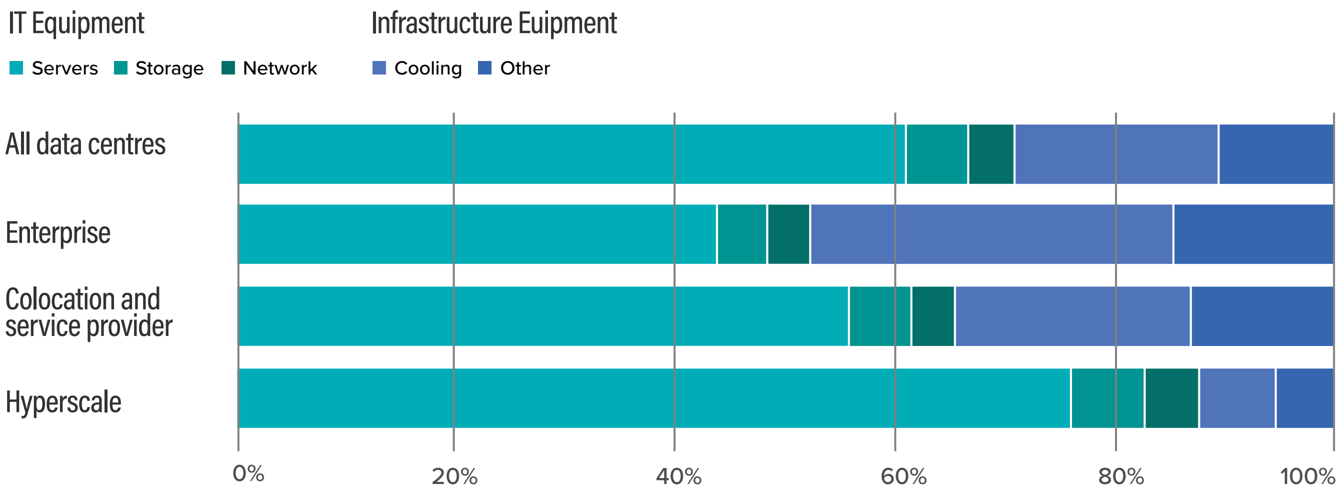

The average power draw has increased over the past decade for all types of data center servers according to stockwide average estimates from LBNL (Figure 1-6).[2] While the data in Figure 1-6 are for the United States, global average power draws are expected to be similar given that servers are commodity products. Conventional servers are single, dual, 4-, and 8-processor servers (Figure 1-6a). These types of servers have historically been associated with traditional data center services, such as cloud computing, database hosting, websites, online commerce and so on. Despite steady energy efficiency gains at the component level, the average power requirements of conventional servers have increased due to a combination of higher wattage central processing units (CPUs), greater onboard memory and storage capacities, and higher utilization levels compared to the past. Accelerated servers (Figure 1-6b) are those that are additionally equipped with accelerated processors, such as graphic processing units (GPUs), which are required for machine learning (ML) and AI model training and inference applications (see Chapter 2.2 of this Roadmap). While accelerated servers are not new, their global stocks have expanded rapidly with the proliferation of AI-focused data centers.

Stocks of accelerated servers can consume far more power per unit compared to stocks of conventional servers (Figure 1-6). As a result, IEA estimates that while accelerated servers accounted for less than 5% of global servers in 2024, their operations accounted for around 20% of worldwide data center electricity demand that same year (accounting for associated cooling, storage and network energy use).[3] The shift to more accelerated servers will have major implications for server energy use in the coming years. For example, IEA estimates that around 8% of the global server stock will be accelerated servers by 2030, whereas LBNL estimates that 21-32% of the US server stock will be accelerated servers by 2028 (8-12 million accelerated servers out of 37 million total servers in 2028).[2]

Figure 1-6. Average per-server operational power draw (kW) of US stocks of different server types for conventional and AI-accelerated servers, historical and projected.. Source: Shehabi et al (2024).[2]

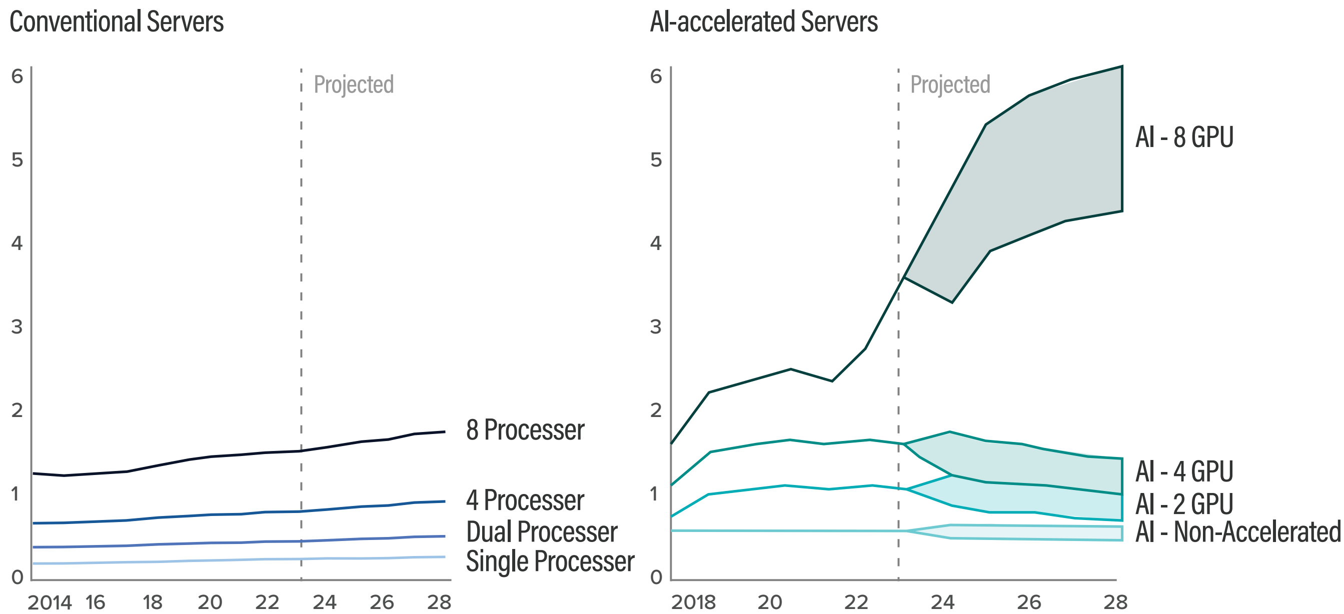

Finally, LBNL's projections for the future power draw of all server types suggest that the average power draws of conventional servers and of the largest AI accelerated servers will continue to rise for the next several years (Figure 1-6). These trends, coupled with expanding server stocks and shifts to AI-accelerated servers, suggest strong growth in global data center energy use in the near term (see Section C).

ii(b). Storage devices

Some servers contain sufficient onboard storage to support their applications. Many others, however, require external storage arrays to store files and data that may be accessed either frequently or infrequently. The two major external storage array types are HDDs and SSDs, although tape storage is still utilized in some data centers for infrequently accessed files, as well as for backup purposes, due to its longevity.[26]

The choice between HDDs and SSDs is typically driven by application requirements, costs and performance. In general, HDDs are cheaper than SSDs and offer higher storage densities, but they have lower data transfer speeds, have higher cooling loads and can be less durable in shock-prone environments. As a very mature technology, HDD arrays can also be expanded and upgraded easily and cost-effectively. In contrast, SSDs are faster and more durable, offering performance advantages in critical applications, while being quieter and running cooler than HDDs. The primary disadvantage of SSDs is their higher cost.[27]

On average, storage is typically the second largest consumer of electricity in the IT rack, albeit at much smaller scales than servers (Figure 1-5).[2] Both LBNL (2024)[2] and IEA (2025)[3] estimate that, as of 2023, HDDs accounted for more than 80% of data center storage capacity with the remaining share attributable to SSDs. The estimated power draws and capacities of both types of storage are summarized in Figure 1-7a and 1-7b, respectively.

For both types of drives, capacities have been increasing steadily as storage companies continue to innovate. However, a different story emerges for HDDs compared to SSDs. The former are also decreasing in power draw per drive, even as capacities increase, leading to steady reductions in energy intensity (the average watts of power draw per terabyte of storage). The latter have seen increases in power draw per drive as capacity per drive has increased, indicating lower energy efficiency gains compared to HDDs. On a watts per terabyte basis, by 2027 HDDs are expected to exhibit substantially lower energy intensities than SSDs, which means they may be the more energy efficient choice for applications where speed and performance are not mission critical (Figure 1-7c).

Figure 1-7. *Trends in estimated power draw per drive, capacity per drive and power draw per terabyte for hard disk types in Figure 1-8 are reasonable for other world regions drives (HDDs) and solid-state drives (SSDs), 2010-2027, derived from panels (left) and (center). Source: Shehabi et al (2024).*[2]

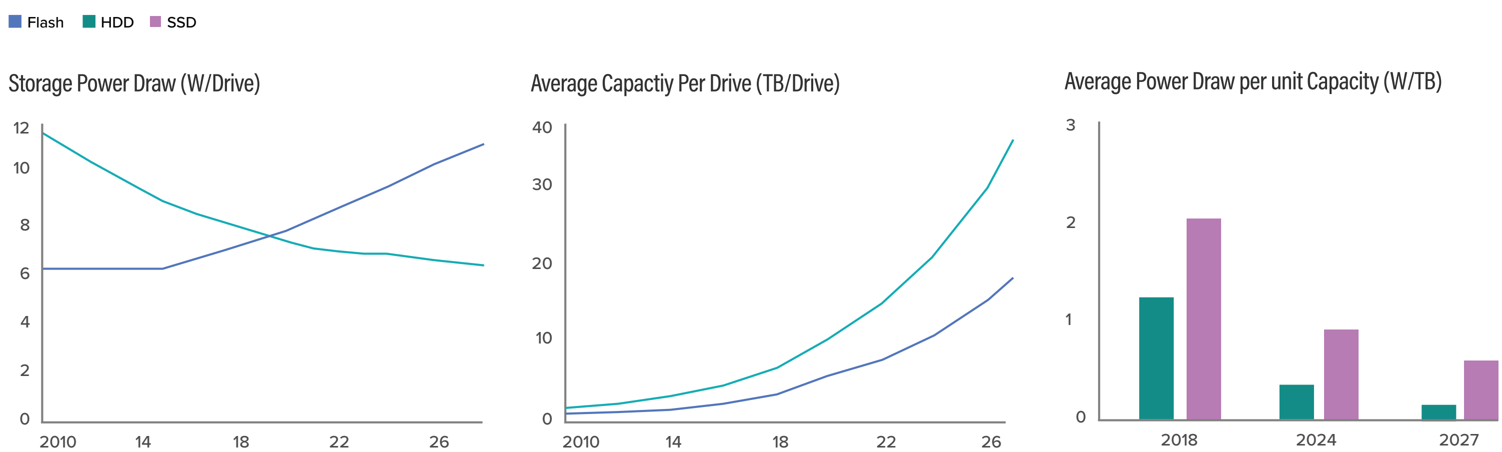

ii(c). Network switches

Network switches generally represent the smallest share of IT equipment energy use by a significant margin. For example, network switches comprised only around 6-7% of estimated national data center IT equipment electricity use in the United States (2023) and Germany (2024).[2,24] Similar to servers and storage, the overall electricity use of network switches is driven by the total stocks of different network switch types and the average per-unit power draws of each switch type. At a high level, network switches fall into two categories: ethernet switches and InfiniBand switches. Ethernet switches have been the dominant class with historical port speeds ranging from 0.1-100 GB/s and are primarily associated with conventional servers. InfiniBand switches are a recent class of ultra-high speed switches (with port speeds up to 400 GB/s presently) that are increasingly used in high-performance and AI computing applications.

Data from market firms such as International Data Corporation (IDC),[28] which are used by some analysts in data center energy analyses, provide network switch shipments in terms of network port numbers and speeds. Figure 1-8 summarizes estimates by LBNL (2024)[2] of the average power draw of different port types over time (left panel)—categorized by port speed—and the corresponding total stock of different ports in the United States (right panel). Similar data are not available for the rest of the world, but IEA (2025)[3] has estimated that both the power draws and the shares of different port types in Figure 1-8 are reasonable for other world regions.

Figure 1-8. Estimated average power draws of different types of network ports and shares of each port type in the United States, 2014-2028. Source: Shehabi et al (2024).[2]

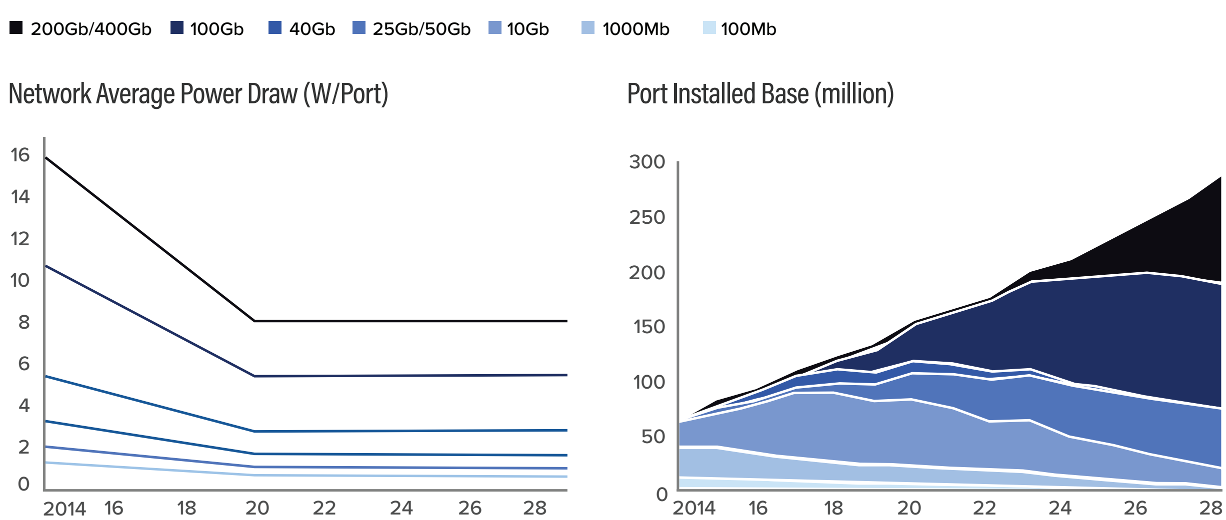

The total electricity use of network switches is a function of the power per port and the total number of ports installed in data centers (as seen in Figure 1-8). While steady per-port power reductions were seen from 2014 to 2020, since that time LBNL's estimates have held constant. Meanwhile, the number of US ports has risen rapidly, with shares steadily shifting to faster port speeds over time as new switch innovations become available. In the United States, the net effect has been around a three-fold increase in network port electricity use from 2014 to 2023[2]; however, overall shares of network switch energy use remained small due to growth in IT electricity use overall.

As more AI-accelerated servers are deployed, the number of InfiniBand ports connecting AI clusters is also expected to increase rapidly. Due this growth, InfiniBand ports are expected to account for 45% of electricity use of all US network switches by 2028[2], underscoring how shifts to AI-accelerated servers will increase not only server electricity use, but network switch electricity use, as well. LBNL (2024)[2] estimates that InfiniBand switches require roughly 11 W/port, which is similar to the high per port power draws of the fastest ethernet switches (Figure 1-8, left panel).

ii(d). Power chain

The electricity use associated with the data center power chain generally arises from conversion and operating losses within transformers, switchgear, UPSs and PDUs. Collectively, these losses can amount to around 10% of total data center energy use on average (Figure 1-4). However, their scales can also vary significantly by data center type (Figure 1-5), depending on the equipment efficiency levels the data center has adopted. For example, Table 1-1 summarizes typical efficiency ranges of the UPS systems different types of US data centers employ, which generally represent the largest losses in the data center power chain. While the combined losses associated with electricity transformation and distribution equipment generally amount to less than 10% of total data center energy use, losses can be much less for the most efficient data centers.[8,20]

Table 1-1. Ranges of uninterruptable power supply (UPS) efficiences by US data center type. Shehabi et al (2024).[2]

Data Center Type	UPS Efficiencies
Commercial edge Enterprise branch Small and medium businesses Telco Edge	77 - 85%
Comms service providers Internal data centers Colocation	80 - 94%
Hyperscale AI data centers	90 - 99%

ii(e). Cooling

Cooling systems generally account for the largest share of infrastructure energy use in most data centers (Figure 1-2). The energy use of cooling systems depends on many factors, including the types of cooling equipment employed, their efficiency levels, data center temperature set points, local climate and more, as discussed in Chapter 2.3 of this Roadmap. A common proxy for the efficiency of cooling systems is PUE (see Box 1-2), given that cooling systems typically consume the largest share of energy for infrastructure at a data center (Figures 1-4 and 1-5). The most efficient cooling systems generally deliver low PUE values, whereas the opposite is true for less efficient cooling systems. Figure 1-9[2] depicts the wide range of PUE values associated with different common cooling systems and approaches (left panel), estimated PUE ranges for different US data center types as of 2023 (upper right panel), and IEA's[3] estimated regional averages for 2024 (lower right panel). Higher PUEs are generally associated with today's small, midsized and co-location data centers, which typically rely on air-cooled racks with chiller-based systems or even direct-expansion (DX) cooling systems in smaller spaces. The listed technologies include liquid cooling systems, which are increasingly deployed to cool high power-density racks at AI data centers and can generally achieve very low PUE values when well designed. However, as AI data center campuses become increasingly large, with some announced data centers set to consume gigawatts of power, even low PUE values can still translate into very large electricity loads.[29]

***Figure 1-9.** Typical power usage effectiveness (PUE) ranges by cooling system equipment type, U.S. data center type in 2023 and world region in 2024.*

Cooling system equipment type

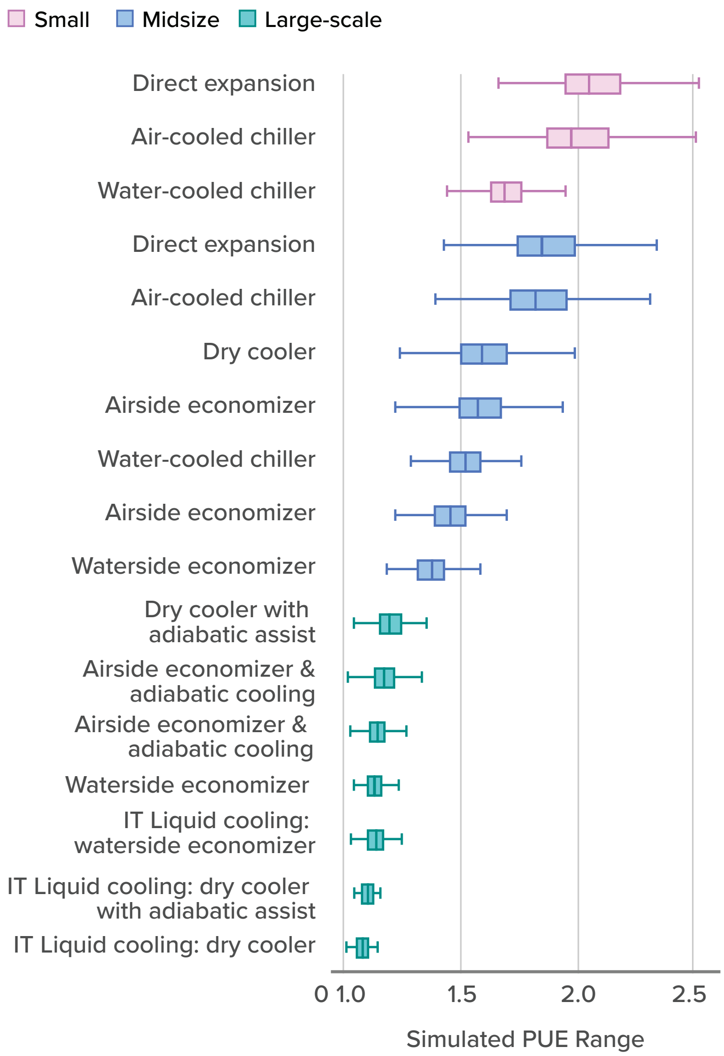

US data center type (2023)

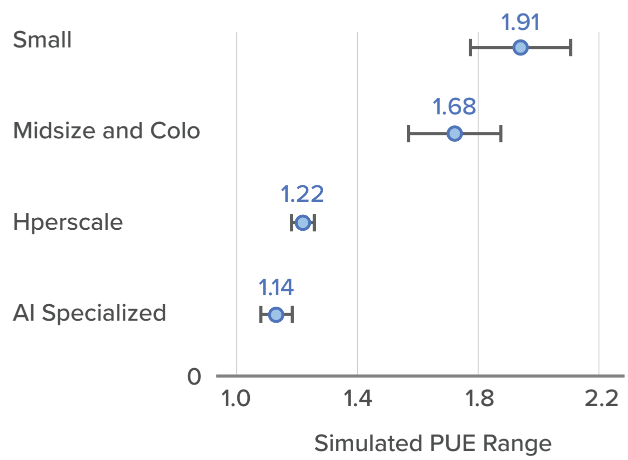

World region (2024)

Power Usage Effectiveness (PUE)	
World	**1.41**
North America | US	**1.32**
Central and South America	**1.70**
Europe	**1.45**
Africa	**1.81**
Middle East	**1.92**
Asia Pacific | China	**1.5**

B. How Much Energy Do Data Centers Currently Use?

i. Estimation approaches

Unlike some other sectors, there are no comprehensive energy statistics available for all data centers at national or global levels. While some large operators report the annual electricity use of their data centers, others only report electricity use for their entire organizations (which can include office buildings, research labs, retail spaces and more), and some operators report no data at all.[29] Even when such data are collected from major operators by governments, as in Ireland,[30] the many small internal data centers operated by non-tech companies are often excluded. While the share of global electricity use attributable to internal data centers is small compared to hyperscale, co-location and AI data centers, as discussed earlier in this chapter, their electricity use is still substantial in an absolute sense. While the reporting situation may change with emerging policies that will require more disclosures (see Chapter 6 of this Roadmap), historically, the only way to quantify the energy use of all data centers at different spatial scales has been to use estimation methods.

In general, estimation methods fall into the following broad categories[2,31,32]:

- **Bottom-up calculations:** Estimate total stocks of different types of IT equipment by region and/or data center type, assume average power draws for each, and apply PUE assumptions (e.g., as discussed in Section A.2) to arrive at grand totals of energy use. Typically considered the "gold standard" of estimation methods due to its rigor and explanatory power, [29,33] its main drawback is that many (and often commercial) data are required for comprehensive estimates.

- **Top-down calculations:** Rely on energy consumption data that are collected by governments, reported by companies and/or are associated with data center populations (e.g., advertised critical IT capacities of co-location data centers) within a given region. While such aggregated data are partially available, this method risks omitting segments of the market that do not report (e.g., internal data centers and non-reporting operators).

- **Extrapolation-based calculations:** Develop estimates by applying growth rates to the estimates of previous years (which can be bottom-up, top-down or even extrapolative in nature), wherein proxies, such as internet traffic, investments or past energy growth, are typically used for growth rates. This method can introduce large uncertainties for multi-year extrapolations and lacks explanatory depth.

Each method has advantages and disadvantages, and each method is prone to uncertainties. While examples of all methods exist in the literature, the most rigorous and transparent analyses generally use bottom-up calculation approaches, such as recent US national and global analyses published by LBNL (2024)[2] and IEA (2025),[3] respectively. Kamiya and Coroama (2025)[31] offer a detailed discussion of estimation methods.

ii. Estimated current energy demand

Summaries of the most recent and authoritative estimates of data center energy demand are provided in Figures 1-10 and 1-11.[a]

At the global level, IEA (2025)[3] recently estimated that data center energy demand increased roughly 3-fold over the past two decades, rising to 416 TWh in 2024, or around 1.5% of global electricity use that same year (Figure 1-10, left panel). However, much of that growth has occurred since around 2017, after a long period of relatively modest electricity demand growth.[33] The post-2017 rise was primarily due to growth in cloud computing, increasing online media consumption and, importantly, increasing deployments of AI computing globally. These trends are reflected by shifts in installed IT power capacity away from internal/enterprise data centers and toward hyperscale, cloud and co-location data centers (Figure 1-10, center). Over the same time period, substantial growth occurred in the world's three largest data center markets: the United States, Europe and China—with the latter experiencing the greatest relative growth due to large investments in AI data centers[34] (Figure 1-10, right panel).

Figure 1-11 summarizes best-available country- and region-level results by market segment and/or equipment type (when available) and for the most recent historical years reported. In terms of individual countries, the United States and China represent by far the greatest estimated shares of recent data center energy demand, at 8 and 5 times the scale, respectively, of Germany, which is the next largest country at around 20 TWh in 2024 (Figure 1-11). As a region, however, Europe is estimated by IEA (2025) to be the third largest data center energy user, at 68 TWh in 2024, followed by the rest of Asia Pacific (i.e., excluding China) at 48 TWh in 2024. Within Europe, data centers in Germany consume 2-4 times the energy of the next largest countries, which include France, the Netherlands and Ireland. Also notable in Figure 1-11 are the very small current scales of data center energy use in the large world regions of Africa, the Middle East, and Central and South America, whose collective energy use (around 5 TWh in 2024) is comparable to that of Ireland alone.

a Global estimates and those for all other countries besides the United States and Europe are from IEA (2025),[3] which relies primarily on bottom-up methods. Estimates for the United States are from LBNL (2024),[2] which also relies on bottom-up methods. Estimates for Europe (except Germany) are from Kamiya and Bertoldi (2024),[23] which aggregates data from different bottom-up, top-down and other types of studies. Estimates for Germany are from Murzakulova et al. (2025),[24] which uses bottom-up methods.

Expressed as a share of total national or regional energy use, a different perspective on data center energy use arises (Figure 1-11). Namely, while the United States and China account for the largest absolute amounts of total data center electricity use globally, data centers comprised only around 4% and 1% of their national electricity consumption in 2024, respectively.[3] Conversely, while Ireland's total data center electricity use is small from global perspectives, at the national level, data centers are a major source of electricity use, comprising 22% of Ireland's national electricity use in 2024.[35] Therefore, the importance of data center electricity use can also vary by country, even if country-level electricity use is a small contributor to global totals.

Figure 1-10. *Estimates of global historical (left) energy use by data center components, (middle) installed power capacity by data center type and (right) data center energy use by major world region, 2005-2024. Source: IEA (2025).*[3]

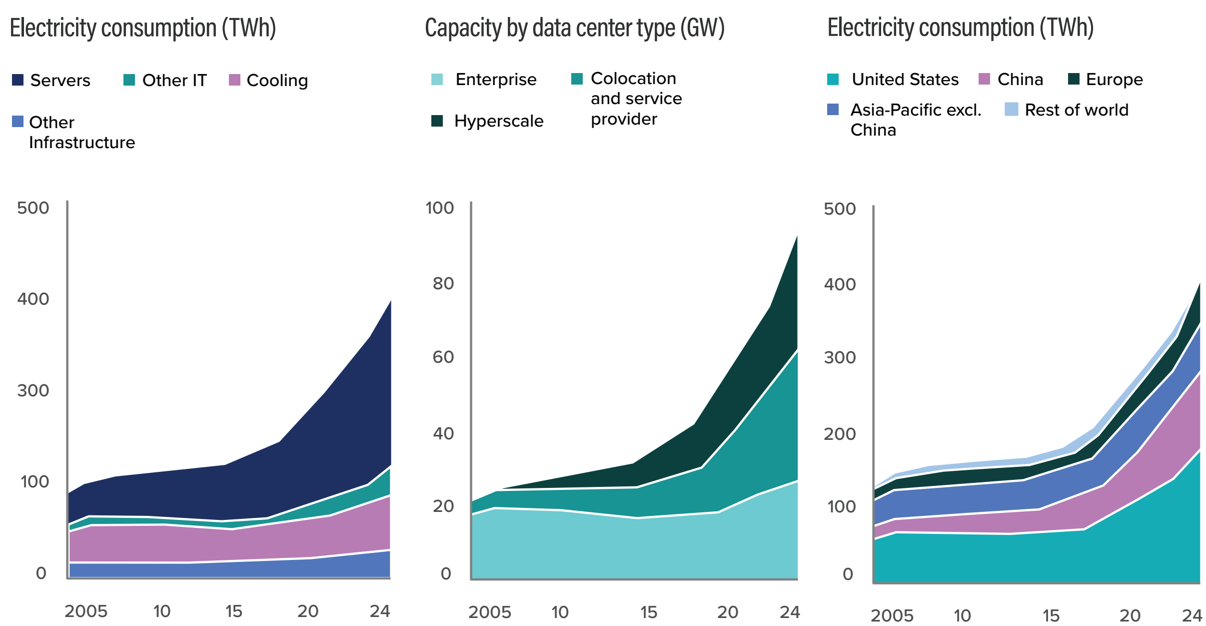

At subnational scales, data centers can affect electricity grids even more acutely. For example, in the City of Santa Clara, California-in the heart of the U.S. Silicon Valley-data centers account for around 60% of total electricity use, leading to electricity rate increases for all inhabitants to support grid infrastructure projects.[36] In the U.S. state of Virginia, which hosts an estimated 13% of reported data center capacity globally,[37] data center electricity use comprised around 25% of total electricity demand,[38] making data centers a critical focus of utility planning, rate structuring, energy policy, and community activism activities in recent years.

Figure 1-11. *Most recent country- and region-level estimates of data center energy demand grouped by greater and less than 50TWh/year. Sources: IEA (2025),[3] Shehabi et al. (2024),[2] Kamiya and Bertoldi (2024),[23] and Murzakulova et al.*

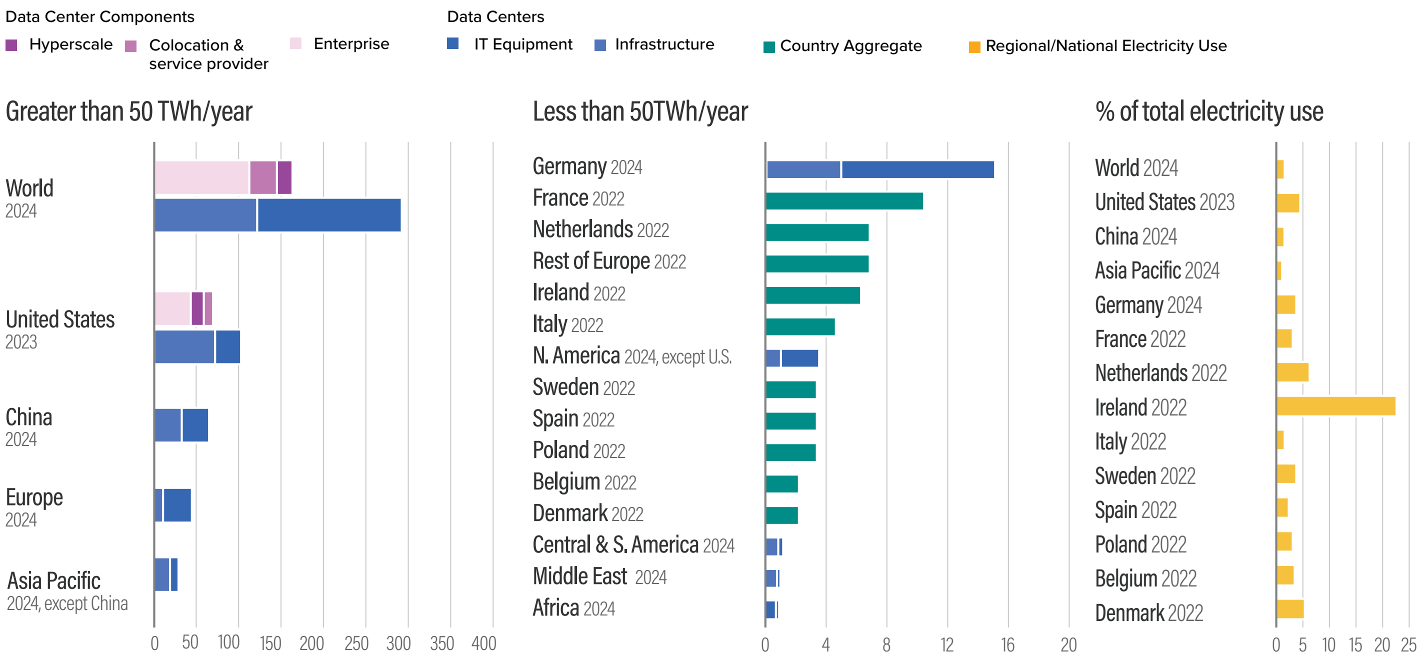

iii. What about carbon emissions?

Accurate estimates of the operational carbon emissions associated with the world's data centers are typically precluded by several factors. First, Scope 1 emissions data[3] that are specific to the onsite and/or backup generators used at data centers are rarely reported. Second, the specific locations of many of the world's data centers, especially internal/enterprise data centers, are unknown. Third, even when the specific locations of data centers are known, analysts must often use region- or country-level emissions factors for Scope 2 emissions.[4] Fourth, when some companies report Scope 2 emissions, they report only "market-based" emissions that reflect renewable energy credits and virtual power plant agreements rather than "location-based" emissions that reflect actual emissions from their local grids. Market-based emissions tend to be smaller than location-based emissions. (Scope 1, 2 and 3 emissions accounting is discussed in Chapter 3 of this Roadmap.)

However, recent estimates can offer a rough sense of scale on regional and global data center carbon emissions.

In 2023 in the United States, according to LBNL, Scope 2 emissions from data centers were 61 billion kg (61 Mt) of carbon dioxide equivalents (CO_2e). This figure is based on an estimate of 176 TWh of electricity (4.4% of total US electricity consumption) and assumed spatial concentrations of data centers within the country.[2] For context, total 2023 emissions from the US power sector amounted to 1454 billion kg (1454 Mt) CO_2e,[40] meaning data centers accounted for around 4.2%. These results suggest that US data centers are located in power grids that deliver, on average, slightly lower CO_2 emissions intensities than the national average.

At the global scale, IEA (2025)[b] estimated that data centers accounted for around 180 Mt of CO_2 emissions in 2024 from the consumption of electricity, which excluded emissions from backup power generation. As a share of total combustion-related emissions globally, this value translates to around 0.5%.

These results suggest that, while data center electricity demand is poised to grow rapidly, thus far it has made relatively small contributions to global CO_2 emissions compared to other energy-intensive sectors. However, this may change in the near future depending on the electric power sources that are employed moving forward to meet increasing data center electricity demand.

C. Where Is Data Center Energy Use Headed?

The rapid increase in data center electricity demand in the past several years (Figure 1-10), coupled with many recent announcements of future AI data center expansions, has led to the appearance of numerous recent studies that project future data center energy use at national and global levels. These studies have employed a wide range of different methods, data sources, analysis periods and assumptions in their projections, with no two studies being directly comparable. Due to these different approaches, there is wide variance in their results, as reported by Kamiya and Coroama (2025).[31]

However, the recent IEA (2025)[3] report provides robust insights on possible future trajectories of data center energy demand, as well as their drivers, the latter of which are summarized in Figure 1-12. Between 2024 and 2030, IEA estimates that the major drivers of data center energy demand growth will be major expansions of both conventional server stocks and energy-intensive AI-accelerated server stocks (see Figure 1-6), increased wattage of AI-accelerated servers, and expanded stocks of storage and network switch arrays to accompany expansions in compute capabilities. However, IEA also estimates that the energy demand growth attributable to the above

b Scope 1 emissions for a data center come from equipment on site, such as backup generators and refrigeration equipment.[39]

c Scope 2 emissions for a data center come from the generation of electricity purchased by the data center.[39]

drivers may be partially offset through reduction in data center PUEs associated with the shift to more hyperscale and AI data centers, which are increasingly adopting low-PUE liquid cooling methods (see Chapter 2.3 of this Roadmap).

As a result of these drivers, in its base case scenario, IEA estimates that global data center electricity demand could increase by a factor of 2.3 between 2024 and 2030, from 416 TWh to 946 TWh. The majority of this estimated growth is associated with large increases in the electricity demand of conventional and AI-accelerated servers (Figure 1-13; left panel) and plays out predominantly in the United States and China data center markets (center panel). In the base case, Scope 2 carbon emissions (right panel) are projected to rise rapidly alongside electricity demand growth, increasing from 180 Mt of CO_2 emissions in 2024 to around 320 MtCO_2 by 2030. Emissions in IEA's base case stop growing and instead begin to decline slightly after 2030, as lower-carbon electricity sources begin to penetrate. Future power grids are discussed in Chapter 4 of this Roadmap.

Figure 1-12. *International Energy Agency (IEA) estimates of the major drivers of global data center energy demand growth from 2024-2030. Source: IEA (2025).*[3]

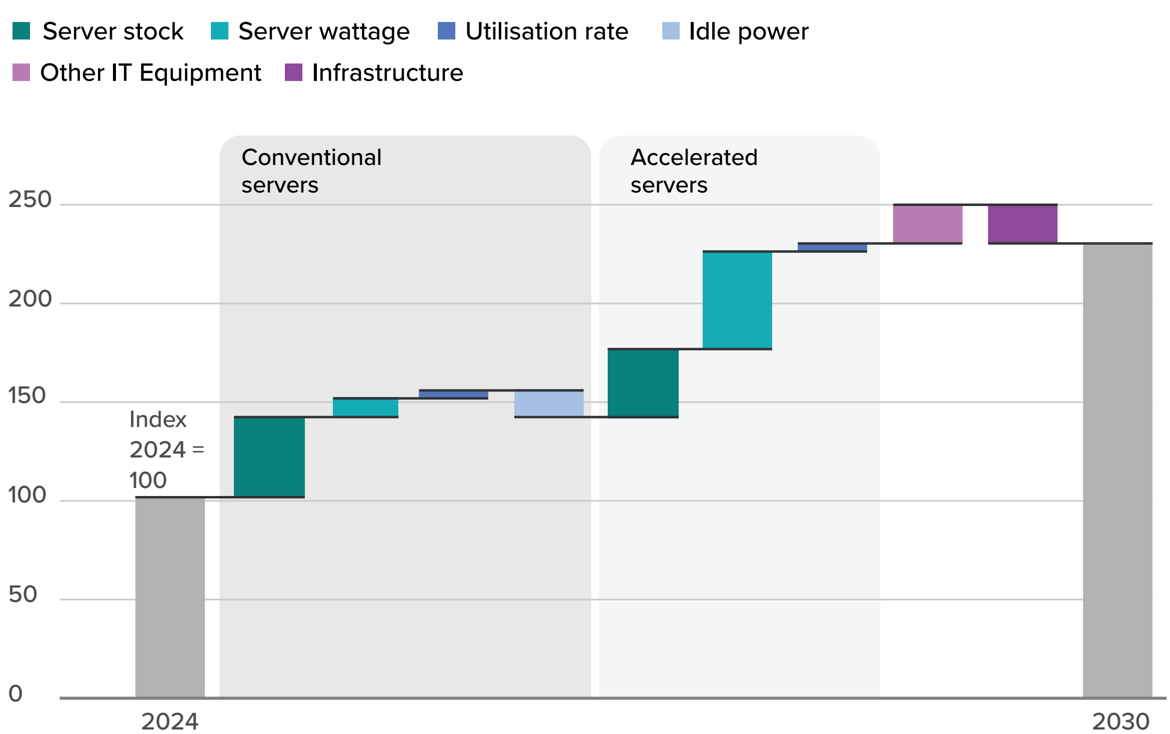

Figure 1-13. *International Energy Agency (IEA) (2025)3 Base Case scenario results for (left) global data center electricty demand by data center component, (center) global data center electricty demand by world region and (right) Associated Scope 2 (location-based) carbon emissions.*

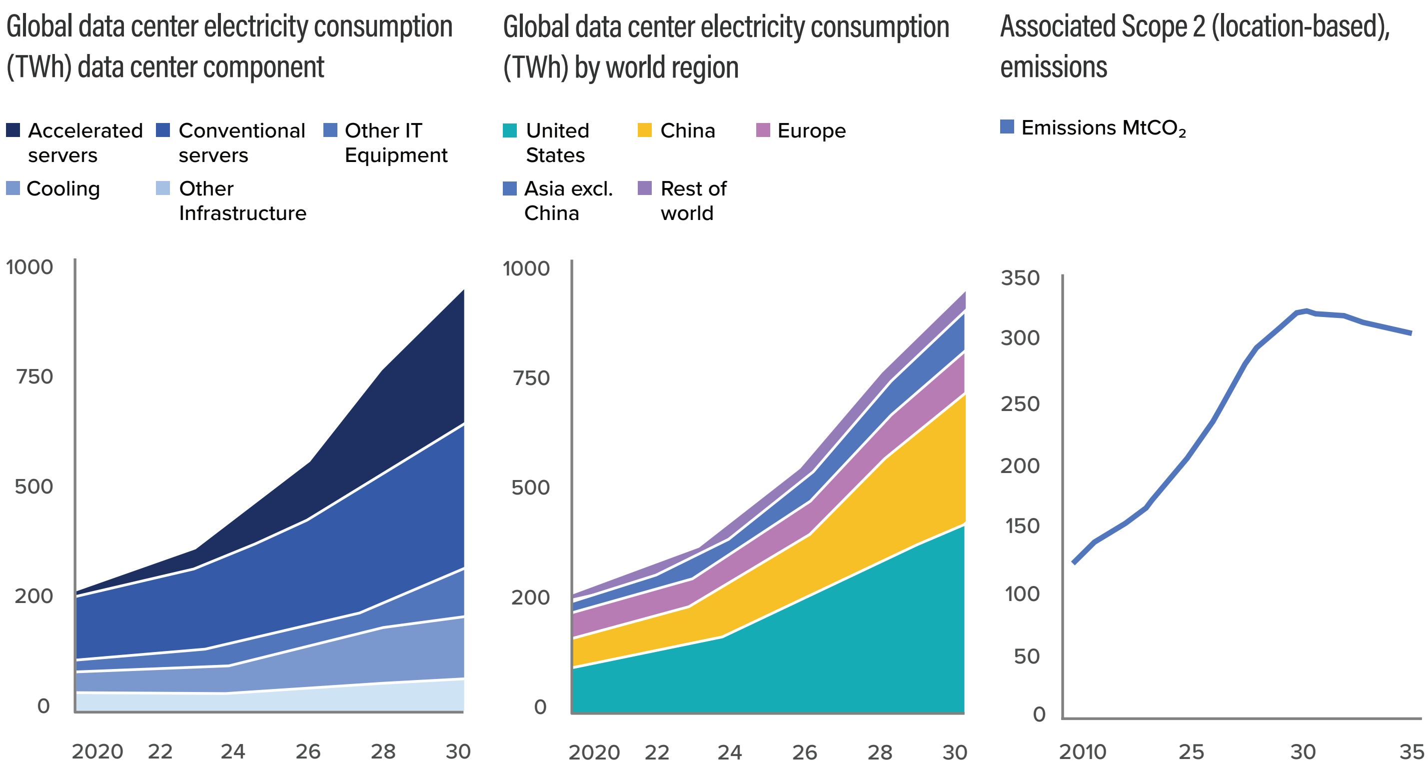

In light of substantial estimation uncertainties, IEA (2025)[3] also considered bounding scenarios that correspond to "bullish" and "bearish" outlooks for growth of AI computing applications. In its "lift-off" scenario, IEA considered stronger AI adoption and increased global demand for digital services, leading to greater deployment of data center facilities compared to the base case. In its "headwinds" scenario, IEA considered plausible constraints to AI's growth, including difficulties financing further investments, strong local constraints on data center development and delayed data center construction due to electricity supply chain constraints. These two extremes result in vastly different outcomes, with data center electricity demand reaching around 1700 TWh by 2035 in the "lift off" scenario and leveling off at around 700 TWh by 2035 in the constrained "headwinds" scenario. These scenarios are summarized in Figure 1-14.

Scenario exercises such as these can shed light on possible future trajectories and what drives them, which can help inform early decisions that may help steer AI and data center energy demand toward the most sustainable outcomes. Such interventions include more efficient hardware and software, more efficient cooling systems, greater use of renewable electricity, reuse of waste heat, and more as discussed in later chapters of this Roadmap. As such, future analyses, especially those meant to guide policy decisions, should develop and implement scenario exercises.[47]

Figure 1-14. *International Energy Agency (IEA) global scenarios for the future of global data center electricity use (TWh), which depict a broad range of potential outcomes depending on market growth, efficiency gains and potential bottlenecks. Source: IEA (2025)*[2,3]

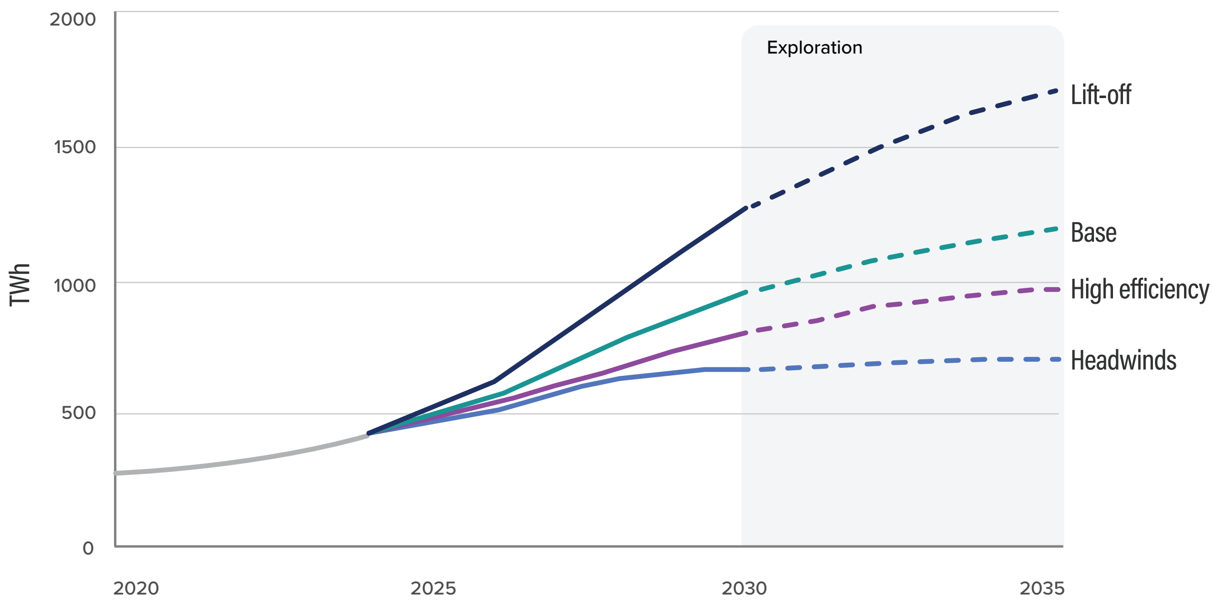

Box 1-1

Jevon's Paradox?

The recent rise in global data center electricity use has been predominantly driven by demand for AI computing. This trend is expected to continue through at least the end of the decade.[2,3] This rising electricity use has occurred despite steady and well-documented gains in the energy efficiency of AI hardware and software that have occurred in parallel (see, e.g., Patterson et al. 2023[41]) and despite growing adoption of the most energy-efficient cooling systems in many AI data centers.

The fact that total electricity use has risen despite large energy efficiency gains has led some analysts to invoke the concept of "Jevon's Paradox," which is attributed to economist William Stanley Jevons. Jevon's Paradox proposes that increasing the efficiency of a resource's use can lead to greater consumption of that resource over time, not less. The theory is that, as efficiency increases, the costs of using a resource are lowered, leading to more demand and expanded use of the resource. Scholars are beginning to research Jevon's Paradox in the context of AI computing. This assessment will require careful analysis of technology trends, potential rebound effects, and the extent to which AI substitutes for or complements demand for societal services.[42,43]

However, Jevon's Paradox has also been hotly debated in the energy efficiency community for years. Some leading scholars suggest that efficiency is a critical feature of sustainable growth, leading to enormous societal energy and cost savings over time while expanding access to energy services. They also argue that Jevon's Paradox ignores other growth factors, such as economic development, and that it paints energy efficiency as a negative force rather than a positive one.[44,45]

With respect to current increases in electricity demand for AI, some analysts have noted that AI technologies are in an early growth phase and many AI businesses are expanding principally due to competition for market leadership, not due to energy efficiency improvements.[46] Finally, there are many examples of energy efficiency improvements that have led to overall energy demand reductions, and despite rebound effects, such as the case of energy-efficient lighting technologies in the United States. Therefore, Jevon's Paradox is not the universal axiom that many view it to be.

While debates about Jevon's Paradox are not likely to be settled soon, most analysts agree that energy efficient hardware, software, cooling systems and data center operations will be critical for ensuring that AI data centers develop in the most sustainable ways possible. For example, in its "high efficiency" scenario, IEA estimates that energy efficiency could reduce the world's data centers' electricity consumption by nearly 20% by 2035 compared to its base case. Indeed, many gigawatts of future data center capacity builds have already been announced, not in response to historical energy efficiency gains, but rather for business, economic growth and national security reasons. Ensuring that these planned builds are as energy-efficient as possible will be critical for minimizing their environmental and social impacts.

Box 1-2

Cryptocurrency mining

Cryptocurrency mining refers to the use of specialized computing equipment that performs calculations—most commonly “proof of work (POW)”—to validate blocks of transactions in a blockchain ledger.[48] Operators of this equipment can earn rewards for each successful validation, which is typically paid in cryptocurrency, hence the term “cryptocurrency mining.”

This specialized equipment, referred to as mining “rigs,” can consume substantial amounts of electricity. For example, the global network of mining rigs for Bitcoin, the largest POW cryptocurrency in the world, was estimated to require around 200 TWh per year (as of August 28, 2025).[49] This is an amount that exceeds the total estimated electricity use of all US data centers in 2024 (183 TWh).[3]

Due to these high power demands, the majority of cryptocurrency mining occurs in dedicated facilities that can each contain many thousands of mining rigs. Just like the IT racks in conventional data centers, these mining rigs also generate large amounts of heat that must be removed. Cooling techniques range from simple facility ventilation fans to cutting-edge liquid cooling systems, depending on the operator.[2,50] Although data on the geographical distribution of cryptocurrency mining facilities are scarce, as of 2022, the majority of Bitcoin mining operations were estimated to be located in the United States, China, Kazakhstan, Canada, Russia and Germany.[51]

The highly specialized and physically distinct nature of cryptocurrency mining means that it is often regarded as its own sector, distinct from either conventional or AI data centers. For further information on the energy impacts of cryptocurrency mining, see Shehabi et al. (2024), IEA (2025), and CBECI (2025).

Besides POW, there are several other methods for validating cryptocurrency transactions, such as proof of authority (POA), proof of burn (POB), proof of capacity (POC) and proof of stake (POS).[52] Of these, POS—used by Ethereum and some other cryptocurrencies—is the most common. The energy requirements of POS cryptocurrencies are much lower than those of POW cryptocurrencies,[48] due to more efficient validation algorithms that can be run on standard computing equipment. CBECI (2025) estimates that the global Ethereum network consumes around 4.4 GWh of electricity per year, which is 45,000 times lower than the global Bitcoin network.

Cryptocurrency mining is outside the scope of this Roadmap.

D. Recommendations

Data center energy use has risen rapidly, and according to many projections, may continue to do so in the years ahead. However, our ability to understand and manage this potential growth depends on having models, datasets and scenarios that can accurately estimate where that growth may be headed, which factors may drive that growth and what interventions and improvements can bend the curve toward more sustainable outcomes.

To support such a vision, the following recommendations could improve the depth and quality of the knowledge base on data center energy use:

1. *Governments and regulatory agencies should* ***develop public data repositories on the energy use and characteristics of data centers*** *at national and sub-national scales, including via mandatory data collection initiatives from data center operators with strict data quality and measurement and verification protocols.*

2. *Data center operators should* ***improve reporting and transparency on the energy use, peak power demand, operating characteristics and other environmental attributes (e.g., water consumption) of data centers*** *to improve the empirical knowledge base for data center energy analysts.*

3. *Governments, philanthropies and research institutions should* ***organize and convene forums to establish best practice analysis methods and data sharing initiatives*** *to rapidly improve the state of science for estimating and projecting data center energy use.*

4. *Researchers should* ***conduct regular inter-model comparisons of data center energy models and scenarios*** *to understand model differences, identify potential improvements, and establish and coalesce on best practices.*

5. *Governments, research institutions and international organizations should* ***convene forums specifically aimed at developing and disseminating scenario narratives for exploring future growth pathways for AI and data centers****, including associated frameworks and analysis approaches; similar to the Shared Socioeconomic Pathways,[47] such scenarios would enable inclusion of data center futures in climate change mitigation scenarios.*

6. *Governments and research institutions should **develop and disseminate models and datasets of data center energy use in developing and emerging economies**, which have historically been overlooked by the research community but whose data center electricity use may grow in the future.*

7. *Research institutions should **reinforce the need for best practices in analyses adopted by policymakers**, given wide variance in results associated with low-quality studies.*[31]

8. *Governments and research institutions should **improve approaches to identifying existing and planned data centers**, such as through in-country sources and open shared location datasets such as the IEA AI Observatory. Improved understanding of the spatial patterns of data centers is important for assessing potential local impacts and proactively designing policies that avoid or minimize those impacts.*

9. *Governments and companies should **support research to better understand the CO_2 emissions of AI data center power sources**, focusing on more granular, grid-scale modeling of emissions. This includes closely tracking announced investments in cleaner power technologies, as well as tracking those investments when they ultimately come online, for more accurate forward-looking scenarios.*

E. References

1. American Society of Heating Refrigerating and Air-Conditioning Engineers (ASHRAE) Standing Standard Project Committee (SSPC). ANSI/ASHRAE 90.4-2019: Energy Standard for Data Centers (ISSN 1041-2336); ASHRAE, Atlanta, Georgia, https://webstore.ansi.org/standards/ashrae/ansiashrae902019 (2019).
2. Arman Shehabi, Alex Newkirk, Sarah J Smith, Alex Hubbard, Nuoa Lei, Md Abu Bakar Siddik, Billie Holecek, Jonathan Koomey, Eric Masanet & Dale Sartor. 2024 United States Data Center Energy Usage Report (Report #: LBNL-2001637); Lawrence Berkeley National Laboratory (LBNL), Berkeley, California, https://escholarship.org/uc/item/32d6m0d1 (2024).
3. International Energy Agency (IEA). Energy and AI; IEA, Paris, France, https://www.iea.org/reports/energy-and-ai (2025).
4. Synergy Research Group. Hyperscale Data Center Count Hits 1,136; Average Size Increases; US Accounts for 54% of Total Capacity; Synergy Research Group, Reno, Nevada, https://www.srgresearch.com/articles/hyperscale-data-center-count-hits-1136-average-size-increases-us-accounts-for-54-of-total-capacity (2025).
5. Mary Zhang. Top 250 Data Center Companies in the World as of 2024; Dgtl Infra LLC, New York, New York, https://dgtlinfra.com/top-data-center-companies/ (2024).
6. Synergy Research Group. Synergy Updates Its Ranking of the World's Top 20 Hyperscale Data Center Locations; Reno, Nevada, https://www.srgresearch.com/articles/synergy-updates-its-ranking-of-the-worlds-top-20-hyperscale-data-center-locations (2025).
7. Apple. 2024 Environmental Progress Report; Apple Inc., Cupertino, California, https://www.apple.com/environment/pdf/Apple_Environmental_Progress_Report_2024.pdf (2024).
8. Nuoa Lei & Eric Masanet. Statistical analysis for predicting location-specific data center PUE and its improvement potential. Energy 201, 117556 (2020). https://doi.org/10.1016/j.energy.2020.117556.
9. Nuoa Lei & Eric Masanet. Climate- and technology-specific PUE and WUE estimations for U.S. data centers using a hybrid statistical and thermodynamics-based approach. Resources, Conservation and Recycling 182, 106323 (2022). https://doi.org/10.1016/j.resconrec.2022.106323.
10. Dan Swinhoe. Meta/Facebook to expand Odense data center campus in Denmark; Data Center Dynamics (DCD), London, United Kingdom, https://www.datacenterdynamics.com/en/news/metafacebook-to-expand-odense-data-center-campus-in-denmark/ (2021).
11. Wikimedia Commons. Photos taken of the SDTPA deployment (File:Sdtpa wmf-8.jpg; Creative Commons Attribution-Share Alike 3.0 Unported license); San Francisco, California, https://commons.wikimedia.org/wiki/File:Sdtpa_wmf-8.jpg (2010).
12. Jon Weiner. Energy Department to Invest $16 Million in Computer Design of Materials; Lawrence Berkeley National Laboratory (LBNL), Berkeley, California, https://newscenter.lbl.gov/2016/08/16/energy-department-invest-16-million-computer-design-materials/ (2016).
13. Better Buildings Initiative & Lawrence Berkeley National Laboraotry. Retro-commissioning at LBNL's National Energy Research Scientific Computing Center; US Department of Energy (DOE), Berkeley, California, https://betterbuildingssolutioncenter.energy.gov/showcase-projects/retro-commissioning-lbnls-national-energy-research-scientific-computing-center (Accessed August 2025).

14. Bill Tschudi. The Other Half of the Problem - Server Closets and Small Server Rooms; Lawrence Berkeley National Laboratory (LBNL), Berkeley, California, https://datacenters.lbl.gov/sites/default/files/ServerclosetLBNL11-12-13.pdf (2013).

15. Nuoa Lei, Mohan Ganeshalingam, Eric Masanet, Sarah Smith & Arman Shehabi. Shedding light on U.S. small and midsize data centers: Exploring insights from the CBECS survey. Energy and Buildings339, 115734 (2025). https://doi.org/10.1016/j.enbuild.2025.115734.

16. Datacenters.com. Homepage: Find a Data Center: Search 7,074 Data Centers; Englewood, Colorado, https://www.datacenters.com/ (Accessed August 2025).

17. Data Center Map (DCM). Homepage: Explore The Leading Global Data Center Database!; Egtved, Denmark and Valrico, Florida, https://www.datacentermap.com/ (Accessed August 2025).

18. Baxtel. About Us Page: We are a Market Intelligence and Advisory Firm; Denver, Colorado, https://baxtel.com/about-us (Accessed August 2025).

19. Theara Coleman. 2023: the year of the AI boom; The Week US (Future US Inc), New York, New York, https://theweek.com/tech/2023-ai-boom (2023).

20. Richard Brown, Eric Masanet, Bruce Nordman, Bill Tschudi, Arman Shehabi, John Stanley, Jonathan Koomey, Dale Sartor, Peter Chan, Joe Loper, Steve Capana, Bruce Hedman, Rebecca Duff et al.Report to Congress on Server and Data Center Energy Efficiency: Public Law 109-431. (Ernest Orlando Lawrence Berkeley National Laboratory, Berkeley, CA (US), United States, 2007), https://www.osti.gov/biblio/929723; https://doi.org/10.2172/929723.

21. Nuoa Lei & Eric R. Masanet. Global Data Center Energy Demand and Strategies to Conserve Energy in Data Center Handbook: : Plan, Design, Build, and Operations of a Smart Data Center, Second Edition (ed Hwaiyu Geng P.E.) pgs 15-26 (2021, https://doi.org/10.1002/9781119597537.ch2).

22. Nuoa Lei, Jun Lu, Arman Shehabi & Eric Masanet. The water use of data center workloads: A review and assessment of key determinants. Resources, Conservation and Recycling 219, 108310 (2025). https://doi.org/10.1016/j.resconrec.2025.108310.

23. George Kamiya & Paolo Bertoldi. Energy Consumption in Data Centres and Broadband Communication Networks in the EU (ISSN 1831-9424); Publications Office of the European Union, Luxembourg, Luxembourg, https://data.europa.eu/doi/10.2760/706491 (2024).

24. Zhanat Murzakulova, Deniz Öztürk, Rafael Noster, Carla Groß, Lars Sorge, Svitlana Chebotarova, Simon Kieweg, Max Schulze, Jessica Stüer, Ralph Hintemann, Simon Hinterholzer, Severin Beucker, Peter Radgen et al. Stand und Entwicklung des Rechenzentrumsstandorts Deutschland; German Federal Ministry for Economy and Climate Protection, Berlin, Germany, https://www.bundeswirtschaftsministerium.de/Redaktion/DE/Publikationen/Technologie/stand-und-entwicklung-des-rechenzentrumsstandorts-deutschland.pdf (2025).

25. Arman Shehabi, Sarah Smith, Dale Sartor, Richard Brown, Magnus Herrlin, Jonathan Koomey, Eric Masanet, Nathaniel Horner, Inês Azevedo & William Lintner. United States Data Center Energy Usage Report (LBNL-1005775); Lawrence Berkeley National Laboratory, Berkeley, California, https://escholarship.org/uc/item/84p772fc (2016).

26. Tim Danton. Return of the tape: the unlikely renaissance of tape storage and how IBM Research is leading the way; TechFinitive,

London, United Kingdom, https://www.techfinitive.com/features/return-of-the-tape-the-unlikely-renaissance-of-tape-storage-and-how-ibm-research-is-leading-the-way/ (2024).

27. Amazon Web Services Inc. (AWS). What's the Difference Between an SSD and a Hard Drive?;Seattle, Washington, https://aws.amazon.com/compare/the-difference-between-ssd-hard-drive/(Accessed August 2025).

28. International Data Corporation (IDC). Worldwide Quarterly Network Infrastructure Tracker;Needham, Massachusetts, https://my.idc.com/getdoc.jsp?containerId=IDC_P23118 (Accessed August 2025).

29. Eric Masanet, Nuoa Lei & Jonathan Koomey. To better understand AI's growing energy use, analysts need a data revolution. Joule 8, 2427-2436 (2024). https://doi.org/10.1016/j.joule.2024.07.018.

30. Central Statistics Office (CSO) Ireland. Data Centres Metered Electricity Consumption 2023 (Online ISSN: 2811-5422); CSO Ireland, Cork, Ireland, https://www.cso.ie/en/releasesandpublications/ep/p-dcmec/datacentresmeteredelectricityconsumption2023/keyfindings/ (2024).

31. George Kamiya & Vlad C. Coroamă. Data Centre Energy Use: Critical Review of Models and Results; International Energy Agency (IEA) Technology Collaboration Programme on Energy Efficient End-Use Equipment (4E TCP) Efficient, Demand Flexible Networked Appliances (EDNA) Platform (EDNA-IEA 4E TCP), Paris, France, https://www.iea-4e.org/wp-content/uploads/2025/05/Data-Centre-Energy-Use-Critical-Review-of-Models-and-Results.pdf (2025).

32. David Mytton & Masaō Ashtine. Sources of data center energy estimates: A comprehensive review. Joule 6, 2032-2056 (2022). https://doi.org/10.1016/j.joule.2022.07.011.

33. Eric Masanet, Arman Shehabi, Nuoa Lei, Sarah Smith & Jonathan Koomey. Recalibrating global data center energy-use estimates. Science 367, 984-986 (2020). https://doi.org/10.1126/science.aba3758.

34. Meaghan Tobin. China Is Spending Billions to Become an A.I. Superpower; The New York Times, Taipei, Taiwan, https://www.nytimes.com/2025/07/16/technology/china-ai.html (2025).

35. Central Statistics Office (CSO) Ireland. Data Centres Metered Electricity Consumption 2024 (On-line ISSN: 2811-5422); CSO Ireland, Cork, Ireland, https://www.cso.ie/en/releasesandpublications/ep/p-dcmec/datacentresmeteredelectricityconsumption2024/ (2025).

36. Khari Johnson. Crackdown on power-guzzling data centers may soon come online in California;CalMatters, Sacramento, California, https://calmatters.org/economy/technology/2025/02/data-center-crackdown-to-protect-california-electricity-rates/#:~:text=Last%20month%20the%20municipal%20utility%20company%20in,up%2060%25%20of%20energy%20from%20the%20utility (2025).

37. Joint Legislative Audit & Review Commission (JLARC). Data Centers in Virginia; JLARD, Richmond, Virginia, https://jlarc.virginia.gov/landing-2024-data-centers-in-virginia.asp (2024).

38. Lisa Martine Jenkins & Phoebe Skok. Mapping the data center power demand problem, in three charts; Latitude Media, Somerville, Massachusetts, https://www.latitudemedia.com/news/mapping-the-data-center-power-demand-problem-in-three-charts/ (2024).

39. World Resources Institute (WRI) and World Business Council for Sustainable Development (WBCSD). The Greenhouse Gas Protocol: A Corporate Accounting and Reporting Standard (Revised Edition); WRI and WBCSD, Washington, DC and Geneva, Switzerland, https://ghgprotocol.org/sites/default/files/standards/ghg-protocol-revised.pdf (2004).

40. US Environmental Protection Agency (EPA). Inventory of U.S. Greenhouse Gas Emissions and Sinks: 1990-2023 (EPA 430-R-25-003); US EPA, Washington, DC, https://library.edf.org/AssetLink/145ky510ew61fk1tq5c2klp5kq5yp33j.pdf (2025).

41. David Patterson, Joseph Gonzalez, Urs Hölzle, Quoc Le, Chen Liang, Lluis-Miquel Munguia, Daniel Rothchild, David R. So, Maud Texier & Jeff Dean. The Carbon Footprint of Machine Learning Training Will Plateau, Then Shrink. Computer 55, 18-28 (2022). https://doi.org/10.1109/MC.2022.3148714.

42. Alexandra Sasha Luccioni, Emma Strubell & Kate Crawford. From Efficiency Gains to Rebound Effects: The Problem of Jevons' Paradox in AI's Polarized Environmental Debate. arXiv:2501.16548 (2025). https://doi.org/10.48550/arXiv.2501.16548.

43. Cyrus Moulton. What is Jevons Paradox? And why it may — or may not — predict AI's future;Northeastern Global News, Northeastern University, Boston, Massachusetts, https://news.northeastern.edu/2025/02/07/jevons-paradox-ai-future/ (2025).

44. Kelly Vaughn. Jevons Paradox: The Debate That Just Won't Die; Rocky Mountain Institute (RMI), Basalt, Colorado, https://rmi.org/blog_jevons_paradox/ (2012).

45. David Owen. The Efficiency Dilemma; The New Yorker, New York, New York, https://www.newyorker.com/magazine/2010/12/20/the-efficiency-dilemma (2010).

46. Katie Fehrenbacher. Data center electricity demand is not a crisis: report; Axios Pro, Arlington, Virginia, https://www.axios.com/pro/climate-deals/2025/02/05/data-center-electricity-demand-crisis-report (2025).

47. Amy Luers, Jonathan Koomey, Eric Masanet, Owen Gaffney, Felix Creutzig, Juan Lavista Ferres & Eric Horvitz. Will AI accelerate or delay the race to net-zero emissions? Nature 628, 718-720 (2024). https://doi.org/10.1038/d41586-024-01137-x.

48. Nuoa Lei, Eric Masanet & Jonathan Koomey. Best practices for analyzing the direct energy use of blockchain technology systems: Review and policy recommendations. Energy Policy 156, 112422 (2021). https://doi.org/10.1016/j.enpol.2021.112422.

49. Cambrige Bitcoin Electricity Consumption Index (CBECI). Bitcoin network power demand; University of Cambridge Judge Business School, Cambridge Centre for Alternative Finance, Cambridge, United Kingdom, https://ccaf.io/cbnsi/cbeci (Accessed August 2025).

50. The White House. Climate and Energy Implicatons of Crypto-Assets in the United States;Washington, DC, https://bidenwhitehouse.archives.gov/wp-content/uploads/2022/09/09-2022-Crypto-Assets-and-Climate-Report.pdf (2022).

51. Cambrige Bitcoin Electricity Consumption Index (CBECI). Bitcoin Mining Map; University of Cambridge Judge Business School, Cambridge Centre for Alternative Finance, Cambridge, United Kingdom, https://ccaf.io/cbnsi/cbeci/mining_map (Accessed August 2025).

52. Cryptopedia Staff. Blockchain Consensus Mechanisms Beyond PoW and PoS; Cryptopedia, Gemini Trust Company, LLC, New York, New York, https://www.gemini.com/cryptopedia/blockchain-consensus-mechanism-types-of-algorithm (2023).