



AI Data Center Power Supply

An Online Continuing Education Course for Engineers

Course Number: E-2139

Credit: 2 Hours / 2 PDH / 2 CPD

AI Data Center Power Supply

Erik B. Norgren, P.E.

1. Introduction

The modern Artificial Intelligence (AI) data center, typically implemented as state-of-the-art training clusters, is at the heart of a new frontier of large language models. A single facility can demand 25 to 250 MW of electrical power, and in some instances, 1,000 MW or more. They can contain up to 100,000 graphics processing units (GPUs), which generate the large language models and video/image models. This course provides a general overview of data center infrastructure (Section 2), then covers power supply options, including the grid and on-site generation. The grid interconnection process is detailed along with current policy efforts aimed at better accommodating the exploding growth of AI data centers.

AI data centers typically cost between \$5 billion and \$50 billion USD or more, and require creative private equity partnerships and financing arrangements to fund. Examples of AI data centers operating or under construction include OpenAI's "Stargate" data center in Abilene, Texas. This massive buildout costs nearly \$500 billion and involves OpenAI, Oracle, Softbank, and Crusoe. This is considered a "Hyper-Scale" application with more than 400,000 GPUs and 1.2 gigawatt power capacity. Another example is xAI's "Colossus 1" supercomputer in Memphis, which began operating in July 2024. It is one of the world's largest AI centers and supports Elon Musk's X platform, Grok, and some SpaceX functions.

To provide an initial perspective on the growth trajectory, the Department of Energy (U.S.) and Berkeley National Lab researched to quantify the rate of U.S. data center growth in terms of annual electricity consumption.




Figure 1: Total Data Center Electricity Consumption (Source: DOE and Berkeley National Lab)

The growth was accelerated in 2017 by the emergency of GPU cluster computing. The usage trend continues to increase, with usage expected to hit somewhere between 350 and 600 TWh in the next few years. This equates to roughly 6% to 12% of overall power consumption in the U.S. This will translate to a peak demand requirement of approximately 75 to 130 GW of power. There are upper and lower bounds in this graph, as it can't be perfectly predicted how GPUs will advance or how liquid-based cooling system efficiencies will continue to improve.

This course will focus on Data centers and their power supply needs. To develop an understanding, the course will first look at data centers and their core components. The reader will begin to understand these core components and why there is such a high power demand. The course will then focus on how data center developers are planning to support their large load either from the grid or by installing on-site generation or storage resources. Grid interconnection processes will be explored by looking at electric utility company requirements that are in place now, as well as developing policy matters. In addition to the grid connection, the course will then provide a detailed view of on-site generation options. Different technology providers will be explored along with supply chain considerations.

2. Data Center Infrastructure Overview

This part of the course provides a general overview of each major component of a data center, which can be broken down into the following key components:

- a. The physical building
- b. Racking, cabling, and communications
- c. Computing hardware, including GPUs, and software
- d. Cooling systems, including piping and equipment
- e. Power Infrastructure

2.1 The Physical Building

The data center building typically resembles a large single-story warehouse. It is designed to securely contain all the IT infrastructure, including servers, networking equipment, office space, and all supporting infrastructure. The design and architecture of the facility are intended to optimize security and power use efficiency.

The structural design of the facility uses reinforced materials like concrete and steel to provide a durable frame and high load-bearing capacity. Fire resistivity is a key consideration for material selection and facility layout. The floors are often raised for underfloor ventilation, and overhead cable trays fill the space above the rack layout. The layout of the facility is typically divided into specific areas, including the "white space" server rooms, office areas, electrical rooms, mechanical rooms, and loading docks.

The image below is xAI's data center, Colossus 1, which took four months to build. It contains 200,000 GPUs and 194 petabytes/s of memory bandwidth. This is a good example of a data center that has utilized a hybrid mix of power supply resources, including on-site natural gas turbines, on-site battery storage, and grid connection. This project is further discussed in Section 4 – Power Supply Options.



Figure 2: Colossus Data Center by xAI (Source: <https://x.ai/colossus>)

The physical security element includes fencing around the perimeter, video surveillance, exterior controlled access points, and interior segmented zones. Safeguarding against unauthorized energy is critical to ensuring data security and facility uptime.

In terms of physical location of the facility, power supply and network connectivity are the key drivers. However, environmental aspects are also considered, such as flood potential and seismic considerations. Large data centers are typically not located in a floodplain, and FEMA provides floodplain data and maps for an initial assessment of flood potential. Locating in a flood zone adds considerable risk and presents significant challenges related to insurance.

2.2 Racking, Cabling, and Communications

Data center racks are fairly standardized and provide the storage framework for the servers, networking, and storage hardware, cabling, and power distribution. Their purpose is to provide an organized and efficient layout that can optimize IT equipment density while providing the necessary supporting infrastructure, mainly power and cooling.

Power supply to the hardware is provided by a power distribution unit (PDU). These can take the form of simple power outlet strips to units capable of energy metering and remote monitoring. Intelligent PDUs allow operators to monitor power usage at a granular rack level and can play a key role in capacity planning.



Figure 3: Rack Layout (Source: DOE – Federal Energy Management Program)

Racks can also include sensors that monitor environmental conditions such as temperature, vibration levels, and humidity. The racks can also include certain security features, like electronic locks for access.

There are distinctions between rack types. Open-frame racks are typically used where security is less of a concern. They offer ease of access for maintenance or hardware changes. These types of racks are more common where air distribution-based cooling systems are the primary means of cooling. Enclosed racks are the most common in today's data centers as they provide a higher level of security. The designer may elect seismic racks that are more rugged and can be placed in earthquake zones.

Data center racks are the building blocks of a data center. The rack design impacts efficiency, total energy consumption, required space, maintenance protocols, and overall facility longevity. Rack system design will continue to evolve as power requirements rise. This drives a need for rack design to continue to become more dense and more efficient.

2.3 Computing Hardware and Software

The foundation of an AI data center is the computing hardware. To support the demands for increased computing capacity, the industry has made a shift from general-purpose central processing units (CPUs) to graphical processing units. GPUs are the most commonly used hardware due to their flexibility in programming. GPUs were originally designed for graphics processing, but are extremely effective at performing parallel computations that are needed for AI machine learning models. GPUs are typically arranged into clusters containing 10 or more, and are able to perform as a single unit with large memory capability supporting "tensor operations."

Tensor operations are mathematical computations structured across numbered arrays. They power the neural network and are able to perform simple operations like addition and subtraction, and more complex operations like dot products and convolution. Tensor operation is used to process text, audio, and image data. These computations require high-bandwidth memory subsystems, which are now becoming the standard in AI data centers. Memory capacity continues to be a challenge as the computational speed now requires multi-terabytes per second to support very large matrix computations.

The interconnectivity between GPU nodes, often referred to as "accelerators," is just as critical. A hyper-scale data center may require thousands of GPU nodes to operate synchronized and thus require low-latency and high-bandwidth communications. The network topology refers to the logical and physical arrangement of switches, routers, GPUs, and storage systems. As the computations increase in size, the network can often become a bottleneck, requiring developers to apply custom network switches and algorithms that relieve data traffic congestion.

For large-scale machine learning, scheduling software is a critical component. There are various providers, including Ray and Slurm, and these proprietary scheduling applications coordinate the GPU nodes, storage devices, and other network components. Scheduling programs are designed to support long-running AI jobs that are resource-intensive and work to minimize data fragmentation.

Lastly, there is an overarching layer of management software that provides monitoring and ties the overall data ecosystem together. GPU performance and utilization, memory consumption and status, thermal management, and data congestion all need to be monitored. Operators are able to utilize the analytics dashboard to track the health of clusters, be able to predict failures, and plan potential capacity expansions.

2.4 Cooling Systems

Tight control of the environmental conditions within a data center is key to ensuring IT equipment performance and safety. High precision ventilation and cooling systems operate to provide a stable temperature and a narrow range of humidity. Maintaining internal temperatures within the space and more specifically within the IT racks, prevents overheating and hardware degradation.

Cooling systems for data centers have needed to be advanced dramatically to accommodate the increasing power load. Rack design is now drawing much higher power requirements per rack compared to traditional server arrangements, with loads exceeding 100 kW. The industry has seen a transition towards liquid cooling systems, where cooling water is delivered to chip-level heat exchange plates. In some cases, designers are pursuing full immersion cooling. The end-goal is to provide the necessary heat removal while also aiming to maximize overall facility power efficiency.

Air-Cooling:

A traditional air-cooling design incorporates airflow paths that work front to back through the rack. The design includes perforated doors and cable management to promote efficient cooling. A key design

consideration is to prevent thermal hotspots across the rack/server layout. A large majority of early data center designs have utilized hot aisle and cold aisle design, along with raised floors for cold air distribution.

The right side of the figure below shows the air-cooling-based system. A fanwall, CRAC unit, or air handler received chilled water from a chiller machine. The air handling unit sends cold air to an underfloor air plenum, which then distributes the air up through floor registers to the hot racks. The cold air is drawn across and through the data racks, heating the air stream. The air is then returned usually through ceiling return inlets and back to the air handler to be cooled again.

The water chiller depicted as #7 in the diagram below can either be air-cooled utilizing internal fans or water-cooled. The diagram reflects an air-cooled chiller that uses a condenser fan and a heat exchanger to cool the water circuit and return the chilled water to the data center. Alternatively, the chiller can be water-cooled, which loops from a secondary water circuit (cooling tower) which uses forced air and evaporation effects to cool the water. This contributes to a significant amount of water

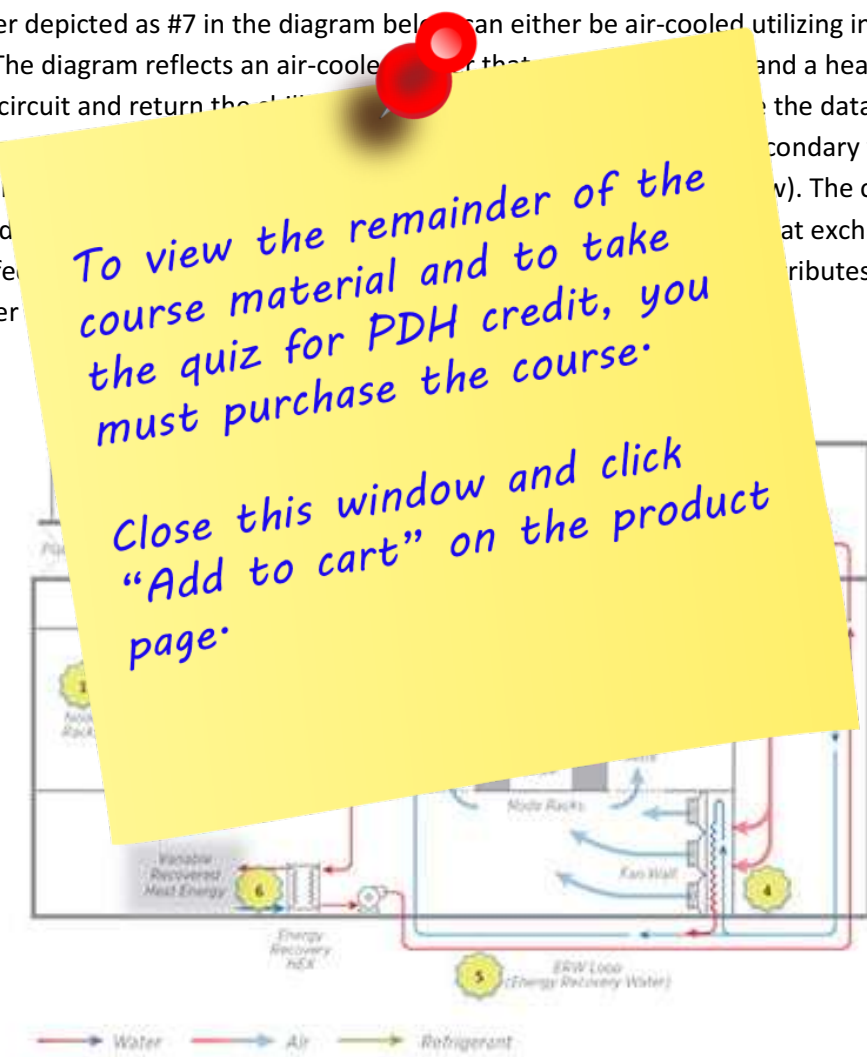


Figure 4: Cooling System Diagram (Source: NREL <https://www.nrel.gov/computational-science/data-center-cooling-system>)

Liquid Cooling:

Nearly all large AI data centers have transitioned to liquid-based cooling due to intense heat load and superior efficiency. Air-based cooling is no longer a sufficient method for modern AI workloads. The