AI Data Center Power Supply

Course No: W02-002
Credit: 2 PDH

Erik Norgren, P.E.

CED CANADA
engineering.ca
CPD FOR THE PROFESSIONAL

Continuing Education and Development, Inc.

P: (877) 322-5800
info@cedengineering.ca

Table of Contents

Introduction

The modern Artificial Intelligence (AI) data center, typically implemented as state-of-the-art training clusters, are 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 (GPU), which generate large language models and video/image models. This course provides a general overview of data center infrastructure (Section 2), then provides power supply options including the grid and on-site generation. The grid interconnection process is detailed along with current policy efforts aimed to better accommodate the exploding growth of AI data centers

AI data centers typically cost between $5 billion to $50 billion USD or more and require creative private equity partnerships and financing arrangements to fund. Examples of AI data centers operating or in 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 (US) and Berkeley National Lab conducted research to quantify the rate of US data center growth in terms of annual electricity consumption.

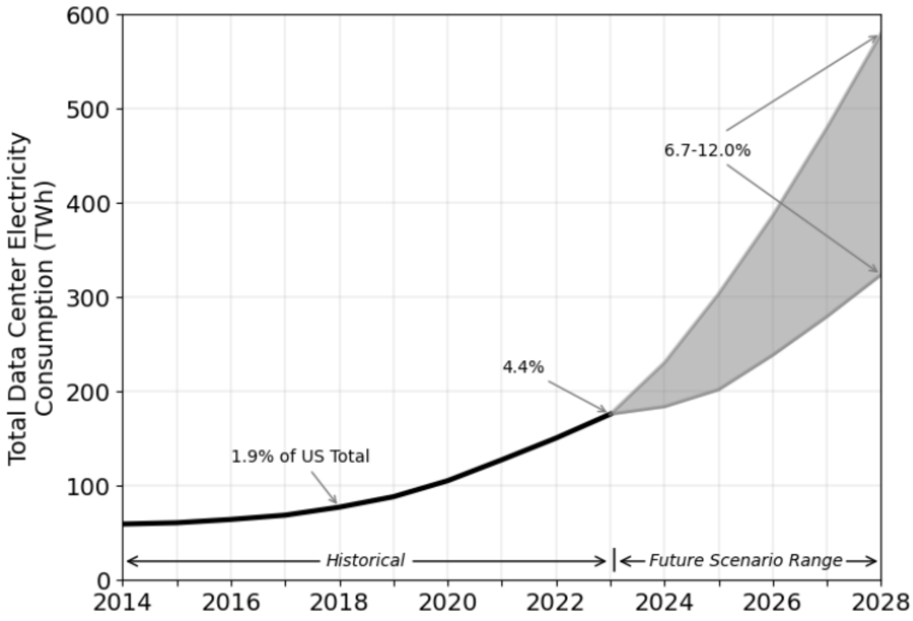

Figure 1: Total Data Center Electricity Consumption (Source: DOE and Berkeley National Lab)

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 to 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.

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

The Physical Building

The data center building typically resembles a large single-story warehouse and 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 is 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 on flood potential. Locating in a flood zone adds considerable risk and presents significant challenges related to insurance.

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 which 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 the provide a higher level of security. The designer may elect seismic racks which 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.

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 (CPU) 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 is 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 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 analytics dashboard to track the health of clusters, be able to predict failures, and plan potential capacity expansions.

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 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 centers 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 handle 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 which uses internal fans and a heat exchanger to cool the water circuit and return the chilled water to the air handling unit inside the data center. Alternatively, the chiller will not have internal fans, and instead will include a secondary water circuit which loops from the chiller and to/from the cooling tower (depicted as #9 below). The cooling tower uses forced air driven by a large fan to cool the air using primarily evaporative heat exchange. This evaporation effect, along with the need to blowdown scaled water to sewer, contributes to a significant amount of water use for the data center.

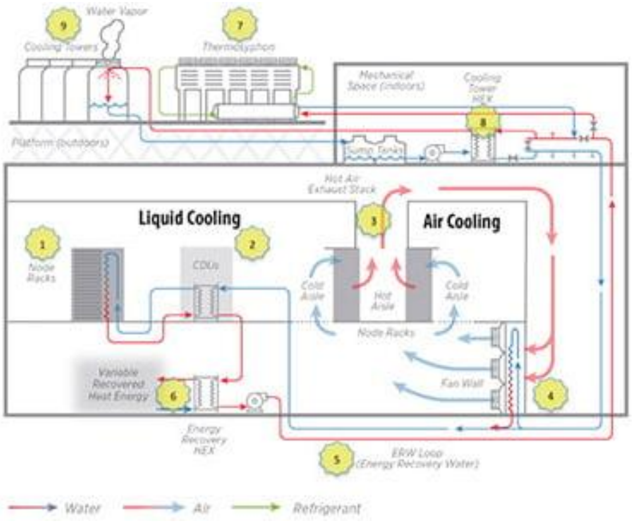

Figure 4: Cooling System Diagram (Source: NREL https://www.nrel.gov/computational-science/data-center-cooling-system)

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 typical approach is to use direct-to-chip cooling which circulates chilled water to heat exchange plates mounted directly to the hottest components.

The National Renewable Energy Laboratory (NREL) explains the remainder of the diagram as follows:

All heat energy from the data center is captured to the energy recovery water (ERW) loop (5), which is a closed-loop system. There are three heat rejection options for this IT load, which operate according to the following hierarchy:

- *When possible, heat energy from the energy recovery loop is transferred to the building process hot water (PHW) loop, which provides heat for the office and laboratory spaces within the building. The energy recovery heat exchanger (6) transfers heat from the ERW loop to the PHW loop.*
- *After re-use potential is exhausted, warm ERW water flows to the fourth floor mechanical room. When temperatures permit, heat is dissipated through a thermosyphon (7), which is an advanced dry cooler that uses refrigerant in a passive cycle to dissipate heat.*
- *Remaining heat is transferred from the ERW loop to a tower water (TW) open loop via the cooling tower heat exchanger (8). Cooling towers (9) cool the TW loop by cascading that water across fill material while drawing ambient air across the fill material. This provides a very energy-efficient way to cool water with sensible (heat dissipated to air without evaporation) and latent (heat dissipated with evaporation) heat transfer. However, this evaporative cooling process requires a continuous source of water.*

Source: NREL https://www.nrel.gov/computational-science/data-center-cooling-system

Power Infrastructure

The power infrastructure supported data center is the biggest cost driver and typically the longest constraint in terms of schedule. Nearby water and fiber are critical as well, however available power is first priority in terms of finding the proper location. Data center developers are clamoring to obtain site control at or near electrical infrastructure that can support the anticipated load. Later this course will explore how developers view electrical infrastructure and the process and potential costs to connect.

Listed below are the major power infrastructure elements:

- Existing Substation and Transmission. Data centers are built near existing substations or transmission lines that can support the load

- Transformers to step down transmission voltage to medium voltage. This is typically 35 kV or 13.8 kV

- Medium voltage distribution along with step-downs to 480V or 415 V

- Uninterruptible Power Supply (UPS) and additional battery backup including lead-acid or lithium-ion systems

- Backup generators (diesel or natural gas) to provide power during a utility outage
  Optional:
- Prime power generators, typically natural gas, which can be used to supplement or replace utility grid power

- Utility-scale battery storage systems to assist the grid or prime power generators with fast-response capacity as AI load fluctuates

- Solar PV and/or wind generation, either located nearby or via remote Virtual Power Purchase Agreement (VPPA) to provide renewable energy to the facility. Installing or procuring renewable energy is often a function of the AI owner’s desire to include renewable energy in the overall power supply profile

Power infrastructure is also heavily redundant which means there are backup systems in place to ensure reliable supply of power. The design will typically require a tier-based model of redundancy which calls the base minimum number of components, typically notated as ‘N’. A redundant design will then incorporate added backup components where N+1 means the design allows for a single failure and has a backup component in place. 2N would mean the design has a full redundant set of backup components.

To summarize, AI data centers include advanced hardware, high-performance networking and communications, advanced cooling systems, and redundant power infrastructure. All of these components represent a complex ecosystem optimized for machine learning. AI models will continue to expand in scale and complexity and the supporting infrastructure will need to evolve and grow with it. Also as the IT load density increases, designers will continue to seek to maximize efficiency to create the most computational capability per unit of energy consumption.

Power Usage

Next, we’ll look at what draws power in a data center. The data center has energy intensive IT equipment, HVAC air-handling equipment, cooling systems, and backup power systems. The facility will also include equipment and appliances to accommodate human occupancy with offices to support on-site staff. The power usage of a data center can be characterized by a term called *Power Utilization Effectiveness* (PUE), which is a measure of the total power used by the facility relative to the power consumption by just the IT equipment. As a simple example, a PUE of 2 would mean 50% of the power it used by IT equipment and 50% is the remainder. For large AI data centers, PUEs of 1.1-1.25 are currently being achieved through use of liquid-based cooling.

Here is a visual conveying a data center with a PUE of 2.

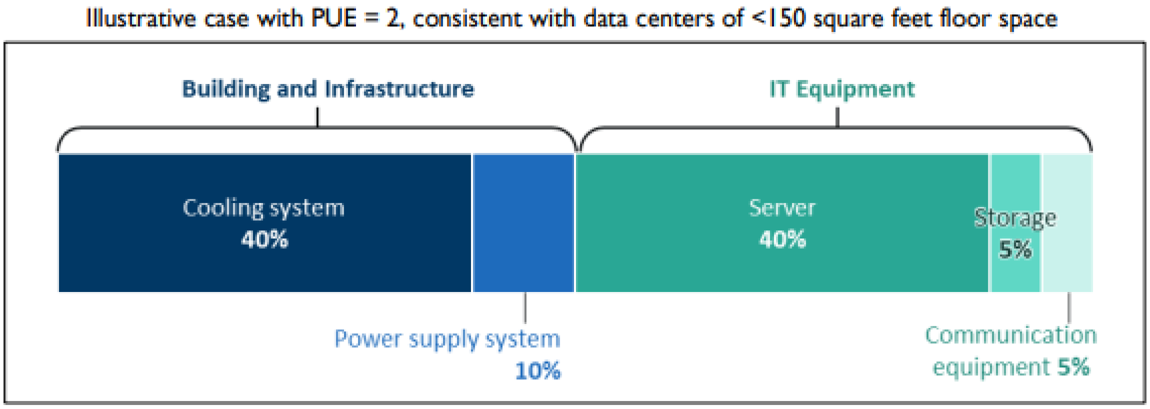
Illustrative case with PUE = 2, consistent with data centers of <150 square feet floor space

Figure 5: Data Center Power Utilization (Source: Congressional Research Service)

Data centers are energy intensive due to the IT load. The core IT equipment is the server racks which contain Central Processing Units (CPUs) and Graphics Processing Units (GPUs). Other elements include data storage drives routers, switches, and memory chips. 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.[1] An advanced data center-level GPU can have a maximum TDP rating between 350W and 700W.[2]

[1] Congressional Research Service: In Focus IF12899, Data Centers and Cloud Computing: Information Technology Infrastructure for Artificial Intelligence, by Ling Zhu. [2] 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/.

The primary power draw associated with the IT load is from the GPU's. Recent trends in GPU power draw have risen very sharpy, with GPUs drawings 400-500 watts around the 2021/22 timeframe. More recently the load has increased to 1,200 – 1,500 watts. The most notable GPU provider is NVIDIA, which continues to make headlines such as this example of a new tech release below.

NVIDIA® announced NVIDIA Rubin CPX, a new class of GPU purpose-built for massive-context processing. This enables AI systems to handle million-token software coding and generative video with groundbreaking speed and efficiency.

- *The NVIDIA Rubin CPX GPU is purpose-built to handle million-token coding and generative video applications.*
- *The NVIDIA Vera Rubin NVL144 CPX platform packs 8 exaflops of AI performance and 100TB of fast memory in a single rack.*
- *Companies can monetize at an unprecedented scale, with $5B in token revenue for every $100M invested.*
- *AI innovators like Cursor, Runway and Magic are exploring how Rubin CPX can accelerate their applications*[3]

[3] https://nvidianews.nvidia.com/news/nvidia-unveils-rubin-cpx-a-new-class-of-gpu-designed-for-massive-context-inference

Figure 6: NVIDIA Example GPU (Source: NVIDIA at link below)

Increased rack density is also driving the power consumption increase. Data centers are creating higher density layouts and current rack design has a power draw of approximately 30-40 kW per rack. This is expected to increase to 50+ kW by 2027 and beyond.

Figure 7: Data Center Racks (Source: DOE)

Power Supply Options

AI datacenters are being developed by aggressive private entities who are profit driven. They are also becoming more central to the economic competitiveness of nations, and as such ensuring a scalable and resilient power supply has become a major focus. The following section explores the power supply options that data centers have today. Each power supply option is discussed and how they differ economically and operationally. In addition to a grid connection with the respective utility companies in the area, there are several on-site supply options available.

The Grid

Most AI data centers are seeking a nearby connection to the electrical grid. Despite certain issues such as time delays, aging infrastructure and grid congestion, the grid remains as the most reliable and predictable source of power for large loads. This relationship however is becoming much more complex as the number of AI data centers seeking connection skyrockets. Section 6 explores regulatory and policy matters that are evolving in effort to deal with the magnitude of large load interconnection requests.

Today's AI data centers are have greatly increased the required power draw versus previously cloud computing facilities which may have only drawn maybe 10 MW as an example. AI data centers are now seeking grid connections with 250 to 500 MW of load or in some cases 1,000 MW or more. The scale of these loads is similar to steel mills, certain industrial or chemical facilities, or small cities. Data centers can no longer simply connect but must undergo a thorough planning and review process with the electric utility company.

The grid represents a vast network of interconnected utility grids with a myriad of power sources including fossil fuel generation (coal and natural gas), renewables including wind and solar plants, hydroelectric facilities, battery storage, and nuclear plants. The exact mix of generations will depend on the general region. For example, the Pacific Northwest has a large percentage of hydroelectric power. Areas in the Midwest generally have a higher percentage of fossil fuel generation. Texas on the other hand has a very high concentration of wind and solar compared to other parts of the country.

It's not just the magnitude of the required power that makes a grid connection challenging, but it is also the load shape. AI power demands can fluctuate greatly over the course of hours and down to minutes and seconds. Accommodating such a high and fluctuating load profile at a specific point in the utility's transmission system can be very challenging. Grid operators must manage large generation resources while attempting to control power flow congestion to feed power to this large single point load.

Requesting a grid connection is typically referred to as a large interconnect load request and is a formal process with the utility and the regional grid operator often referred to as the Independent System Operator (ISO). This process is explored in the next section.

The studies performed during the large load study process will determine the necessary transmission upgrades that need to be performed to accommodate the project. Local facilities to directly connect the project to the transmission system will also be quantified in terms of scope, schedule and cost. In some cases the transmission network simply cannot deliver the power needed for a large AI data center, and the associated cost is not feasible.

Most AI developers have available transmission capacity at the top of their siting criteria. The type of siting strategy is similar to the site selection process for semiconductor facilities and heavy industrial facilities with large loads. Siting near a substation or lines with sufficient capacity will typically reduce costs and in some cases reduce the total time to build. A reinforced grid backed by substations suitable to support these types of loads will continue to be a key building block. Although securing access to power is a critical step in the process, it may come with years of development and could cost hundreds of millions of dollars to connect to the grid.

On-Site Natural Gas Generation: Reliability and Speed

On-site power generation can be used to either A) supplement the grid or B) eliminate the need for the grid. This is true for several types of on-site generation discussed in the following sections.

Natural gas generation is considered a reliable and quick option. The types of natural gas generators are simple cycle gas turbines, combined cycle gas turbines and reciprocating engine generators. Each provide their own merits in terms of efficiency, cost, space requirement, and operational characteristics.

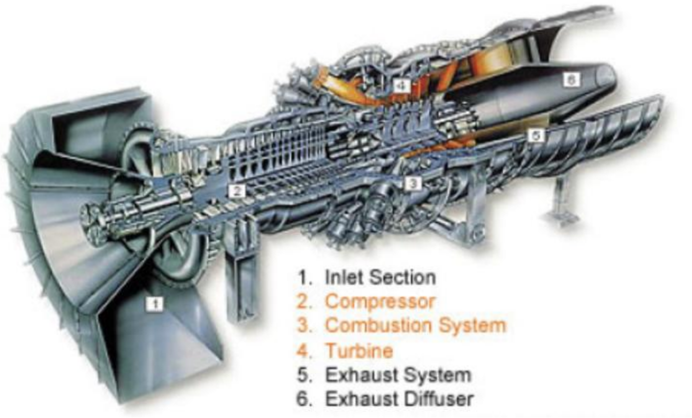

Figure 8: Simple Cycle Natural Gas Turbine (Source: DOE)

Natural gas generators can offer several key advantages:

Fast deployment:

Gas generation can often be deployed faster than the electric utility can upgrade or build new transmission lines and substations. Gas generation is often deployed on a 3 to 5 year timeline while the utility's timeline may be longer.

High reliability:

Gas turbines and engines can run for months, assuming no unexpected failures. They are generally reliable and have high uptimes exceeding 95%. The design will often include extra redundant generators to backup other generators when down for maintenance.

Predictable fuel supply:

Natural gas remains a relatively inexpensive fuel and in high supply. Developers can secure long-term contracts from commodity gas suppliers. They also have the option to pay for "firm" gas supply, meaning the supplier has an obligation to provide a consistent and non-curtailable supply.

Grid independence:

Given its speed of deployment, natural gas generators can offer the ability to operate the data center without grid supply. In this case, the facility would be "islanded" from the utility grid. This is also commonly referred to as a "microgrid". In this scenario, the fleet of generators, along with controls and potentially other on-site resources, need to operate in unison and handle the fluctuating AI load. Maintaining consistent frequency and voltage in an islanded grid can be very challenging and requires precision engineering.

Potential drawbacks include environmental concerns and emissions regulations. Those items need to be navigated carefully and addressed properly in the overall project schedule. Also, this option may not agree with certain corporations' sustainability goals.

Fuel Cells

Fuel cells are unique in that they are considered non-renewable, given the fuel is often non-renewable, yet they do not produce harmful emissions. The fuel can be natural gas, biogas, or hydrogen, with the most common being natural gas. Fuel cells produce power via an electrochemical reaction rather than combustion. They are generally known for high efficiency, little noise and extremely low air emissions.

For an AI data center application, fuel cells have the following advantages:

- **High power density** suitable for GPU clusters of 10 to 100 MW modules.
- **Modular deployment** allowing designers to easily match system size to the power demand
- **Continuous operation**, a key requirement for high-availability AI computing

Figure 9: Fuel Cell System (Source: DOE)

The market offers hydrogen based fuel cells as the fuel source. However hydrogen as a feedstock presents significant economic challenges related to the production and storage of the fuel. Also, the vast majority of hydrogen available in the market is "grey" which means it was produced from natural gas without carbon capture. To produce truly "green" hydrogen, one must use a renewable powered electrolysis process. An electrolyzer uses power to split water into hydrogen and oxygen, which is an expensive and inefficient process.

Only a small percentage of AI data center developers are currently installing fuel cells given economic challenges. Some are installing fuel cells using natural gas as the fuel, with hopes of hydrogen fuel becoming more economically feasible. Although not widely adopted, fuel cells represent a promising pathway towards a potentially more environmentally friendly on-site power solution.

Renewable Power

Many data center developers look to procure renewable power to feed their facilities. A Power Purchase Agreement (PPA) is often pursued, which is typically a long term power offtake contract between the data center and a 3rd party wind or solar plant. These types of agreements allow the data center entity to claim renewable energy power usage either with the wind/solar farm directly connected to the data center or alternatively the wind/solar facility is remote. PPAs will typically be for 10 to 20 years contract length.

PPAs with wind and solar plants do present a challenge in that they have intermittent production. The generation output often times does not line up with the facility demand. Attempting to line up real-time energy use with renewable facility output is often referred to as real-time energy matching. Separately, a data center developer may also contract with an energy storage facility which can allow the data center to claim improved real-time energy matching to better meet their sustainability targets and claim Renewable Energy Credits (RECs).

Many data center developers attempt to co-locate their facility with renewable sources. This may include:

- Installing a solar facility adjacent to the data center
- Locating wind turbines within the same regional grid
- On-site battery energy storage used to firm intermittent renewable facility output

In general, renewable wind and solar facilities are usually supplemental and contribute to a data center's sustainability goals. However these sources will rarely replace the need for either grid power or natural gas-fired power sources. AI facilities require 24/7 power needs with unpredictable and fluctuating load profiles. With solar output peaking at mid-day and wind output varying by hour and seasonally, the intermittency often does not match up with the data center load. Energy storage can help assist with addressing the intermittency, but often the energy storage solutions are need big enough to completely solve this issue.

Renewable energy will remain valuable to data centers, but overall insufficient as a standalone solution. If a data center seeks to provide their own on-site power resources, it will rely on a core natural gas solution for the moment, along with optional supplemental resources like wind, solar and/or battery storage. Incorporating multiple forms on generation and storage options is typically referred to as a "hybrid" solution.

Battery Storage

Utility-scale battery storage is playing an increasing role when it comes to powering IA data centers. They are becoming more popular and valuable, however remain as a supplemental resource rather than a primary resource. The most common technology is lithium-ion based storage solutions, which typically come in outdoor containers. They can provide rapid response to the data center, can smooth out the intermittency of renewables, provide grid frequency regulation and may be an alternative backup power solution to diesel generators.

There is a practical limitation to how much and how long a battery can serve the AI data center load. For example, an AI facility may be near 400 MW fairly consistently across an 8 hour span. Serving this load entirely would require a 400 MW x 8 hour storage duration = 3,200 megawatt-hour storage solution. This would be one of the largest battery storage applications in the world. The most common utility-scale energy storage solution is usually smaller around 1,000 megawatts-hours or less. Battery storage system costs are improving; however the economics still do not support increasing a battery solution to this extreme size.

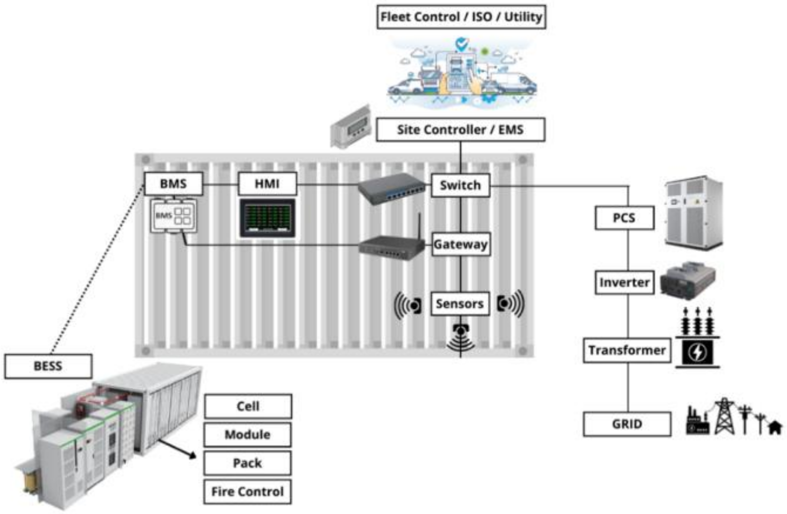

Figure 10: BESS Components (Source: DOE)

Battery storage is common for the following use cases:

- Providing bridge power during the period of increase AI load and natural gas generator startup. This usually requires several minutes
- Stabilizing voltage and frequency also known as improving power quality
- Supplement grid power when the AI load peaks, also known as “peak shaving”
- Provide temporary backup power during a grid outage event

For a hybrid-type system consisting of other on-site generation resources, a battery can pair with these resources to provide a more stable and flexible overall solution. In summary, battery storage can provide supplemental services that may be valuable. The data center developer will need to weigh these benefits against the added cost to determine if battery storage makes sense for the application.

Microgrids

A microgrid is an electrical system with its own generation sources which can operate independent of the utility grid. The system may or may not include the option to operate connected with the grid. AI data centers increasingly are exploring options to establish a microgrid and combine multiple power sources into a single coordinated system.

A typical AI microgrid might include:

- Grid interconnection
- On-site gas turbines or fuel cells

- Battery storage
- On-site solar arrays
- Control systems capable of autonomous islanding

Microgrids offer an added level of resilience to a data center. If the utility grid fails, the microgrid can switch to its on-site generation resources. If the design includes a grid connection, the grid can be used to supplement the on-site generation resources and pick up the supply if on-site generation is down for maintenance. Also if the facility has connected solar or wind power and the output spikes, the microgrid can store surplus energy in its batteries.

AI data centers cannot tolerate downtime, so this added level of resilience is invaluable. Microgrids can also help operators optimize costs by allowing the operators to leverage different power supply options and fuel which can fluctuate in real-time. Although a microgrid will require considerable added capital investment they are becoming more and more popular in modern AI data center design.

Small Modular Reactors

Small modular reactors (SMR) are becoming more popular in the data center discussion for long-term power supply. These small nuclear reactors can be designed to produce hundreds of megawatts and can offer several key advantages. These advantages include continuous "baseload" power, small footprints, very low emissions and minimal fuel requirements. Baseload power refers to providing a steady and continuous supply of power where other resources such as solar may provide intermittent power.

SMRs can be an ideal match for AI data centers:

- SRMs provide stable baseload power suitable for nonstop GPU power supply
- SRMs reduce or eliminate the need for an alternative baseload power source such as the grid
- SRMs provide reliability and predictability for the data center

Despite its advantages, SMRs face substantial hurdles including regulatory, economic and commercial hurdles. There are very few designs have received full regulatory approval and development and construction timelines remain long. There are also some unknowns related to cost which can challenge the formation of an economic thesis for these projects. The technology also faces public perception and local acceptance obstacles given it involves nuclear fuel and accident and disposal concerns.

Many companies are considering SMRs despite and some are willing to make the bet that SMR technology will provide to be a great choice for baseload power supply. Over time, SMRs may become the most transformative power solutions available to data centers, especially as global computation demands continue to accelerate.

Backup Power: Diesel, Gas, and Battery Hybrids

Nearly all data centers are installing backup power systems. Diesel generators can be a common choice for backup power as they have been a long-standing backup power choice for other types of facilities. They generally have high reliability, long shelf life of diesel fuel and are relatively easy to operate. The main concern with diesel generators is their environmental impact on greenhouse gas emissions. Given these concerns, AI data centers may consider alternate options including:

- **Gas-fired backup generators** which emit less pollutants to the air compared to diesel generators
- **Battery or flywheel backup** providing instant response until generators start, get up to speed and finally drop in to provide power. This can take a few minutes
- **Fuel cell backup systems** which can offer cleaner greatly reduced emissions

When setup as a backup power system, these sources are rarely operated to provide continuous power when the grid is available. Their purpose is to maintain reliable data center computing by backing up the power grid or backing up other generators which are designated as primary power sources. Losing data center operations has an enormous financial impact, so these backup power systems are critical.

Integrated Power Strategies

Data centers do not have to choose a single option for power supply. Instead they can choose really any combination of the options presented previously. No single option is typically considered sufficient, so data center developers must carefully analyze their options, the benefits, and the economic impact of each. These integrated strategies can include:

- The grid which supplies predictable baseload power supply
- On-site gas turbine generators or fuel cells which provide resiliency and can replace or be a supplement to the grid
- Renewables (wind and solar) which can help the data center meet sustainability goals. Renewable sources are a supplemental source
- Battery energy storage which can aid in stabilizing the system and smooth out intermittent supply from renewables. Also can provide quick power during on-site generator startup
- Microgrids which provide the ability to operate independently from the grid. A microgrid control system orchestrates the interaction between on-site power sources and the data center load
- SMRs are a future option to provide on-site baseload power

Data center developers will continue to evaluate these hybrid power supply options to ensure reliability and keep their solutions economically competitive. As each of these technologies continues to improve and better craft their offerings to data centers, we will continue to see refinements in how data center developers implement creative hybrid strategies.

Real Project Example

The xAI Colossus data center located in Memphis is an excellent example of a hybrid solution for power supply. This data center is primarily powered by the utility grid which in this case is provided by Memphis Light, Gas and Water (MLGW) and the Tennessee Valley Authority (TVA). The site has natural gas turbine generators and Tesla Megapack battery energy storage systems which are deployed to provide supplemental power and ensure continuous operation.

As of early 2025 xAI Colossus had a load of approximately 150 MW. They have requested expansion plans to the utility companies to increase the load to over 1,000 MW. They plan to do this in phases with the first phase of expansion roughly doubling the load from 150 MW to 300 MW.

The facility will consider a continued hybrid approach to meeting the increasing power demand. Depending on the final timelines for the utility to accommodate the increasing load, more natural gas turbines may be added to accommodate the demand increases. These turbines are considered portable type, as they are fully containerized solutions which can be quickly added or removed. The facility may also add more Tesla megapacks to supplement the power needs. The facility reportedly has more than 150 megapack units installed currently.

Grid Interconnection Process

This section provides an overview of the grid interconnection process for a data center. This is typically referred to in the industry as a large load interconnection and is treated by utility companies in a specialized process to study and approve the proposed connection. The Electric Reliability Company of Texas (ERCOT) oversees the operations of the high voltage transmission system across Texas and provides a good example of a formalized process. Although ERCOT operates an energy-only wholesale market that is largely independent from the Federal Energy Regulatory Commission (FERC), it still provides a good example model that other Independent System Operators (ISO) and utilities are implementing.

In ERCOT specifically, a data center will be treated as a "Large Load Addition" when it triggers one or more of the following criteria:

- The load is greater than 20 MW
- The load increases by 10 MW or more at an existing substation
- A load change that will trigger significant transmission upgrades or will have significant reliability impacts on the local transmission system

Large Load Interconnection Request

When a data center developer is ready to submit their facility load request to the local utility company, they will need to assist in preparing a Large Load Interconnection Request (LLIR). The LLIR serves as the formal starting point for the interconnection process. In most areas of the country, the utility must be the entity that submits the LLIR to the respective Independent System Operator, in this case ERCOT. Other ISOs across the country include CAISO, SPP, MISO, PJM, and ISO-NE. ISO's are responsible for overseeing a regional transmission network by managing power flow and managing the electric wholesale market. The data center customer will assist the utility in preparing the LLIR by providing key data.

A typical LLIR package will include the following items:

- Estimated peak load requirement
- Expected shape of the load, both hourly and seasonal
- Reliability requirements
- Proposed customer-owned substation equipment
- Single-line electrical diagrams
- Power quality considerations such as harmonics

- On-site generation resources (natural gas generators, solar, wind, fuel cells, etc). Typically need to submit equipment data/cut sheets, single line diagrams that include these generators and their connection to the facility, and power flow models

Data center designs often include UPS systems, power electronics, and fast load ramping characteristics. For this reason, the utility and the ISO will want to see detailed modeling of harmonics, voltage behaviors, and frequency response implications.

Reliability Study:

After the LLIR is submitted along with typically a $100,000 or more study fee, ERCOT will review the LLIR to determine if additional reliability studies will be required. Large load interconnections will typically require a comprehensive reliability review to evaluate the impact on nearby transmission system elements including transformers, conductors and protection settings. These evaluations will include N-1 and other contingency scenarios. The analysis for large loads is iterative and collaborative with the utility. Data center developers typically will not directly interact with ERCOT through the study process.

ERCOT will evaluate whether existing transmission lines, transformers, breakers, and substations can support the requested data center load. If the added load triggers overloads, voltage violations or short-circuit issues, ERCOT identifies system upgrades that will be required.

ERCOT and the other ISOs operate their systems under an N-1 standard. This means they predict and assume next single element outage and operate the system accordingly. An example would be assuming the next transmission line to have a thermal overload event is down and out of service. As ERCOT studies the new data center load, they determine whether the system can safely handle the new demand profile under multiple credible contingencies.

Voltage Performance and Stability:

The magnitude of the data center load and its typical location near the distribution-transmission system interface require that voltage performance and stability are closely evaluated. ERCOT and the utility may model voltage recovery, voltage drop during fast data center load ramp, impacts of large UPS batteries, and potential harmonics issues from power electronics. In many cases, ERCOT will require additional modeling of UPS systems or require that specific harmonics filtering is installed.

Short-Circuit Duty and Protection

Large loads connected to substations will increase the fault current contribution. The utility and ERCOT will evaluate whether relays, breakers and overall protection design is capable of handling new fault levels. If any components fall short, replacement of breakers or relays may be required. This can have an impact on timeline or cost as in many cases these upgrades will need to be performed by the utility company. In general, utility upgrades take longer than the speed at which the data center developers want to build and commission their facility.

On-Site Generation:

If an on-site generation source plans to provide power to the facility simultaneously with utility power, this is commonly referred to as "paralleling" with the grid. The on-site generation sources need to match the voltage and frequency of the utility grid to ensure safe and reliable operations. In order to gain the utility's approval to operate in parallel with the grid, the utility must study the proposed on-site generation and run power models to assess the impact on the local utility grid. This study, review and approval process is referred to as the grid interconnection process.

Transmission Upgrades

When upgrades are determined to be required, ERCOT and the utility will classify them as either "directly assignable" or "network upgrades". Directly assignable are typically funded by the data center developer. The data center will need to pay for these upgrades according to a proposed payment schedule by the utility. Network upgrades are usually upgrades that the utility funds and recovers through electric rates. Recovery through rates is subject to the Public Utility Commission (PUC) oversight. Data center developers have sought to find sites with limited upgrade requirements. Working with the utility as early as possible in the development process is highly recommended.

Utility Responsibilities and Engineering Design

Once ERCOT completes its reliability and power flow studies, they will move into a phase of determining the required "interconnection facilities" that are required to physically connect the data center to the grid. These new facilities will often include a new utility-owned substation or substation expansion. Depending on where the site is located the utility may need to build a dedicated high-voltage with a typical voltage of 138 kV or 345 kV.

The utility's scope may also include new runs of transmission lines or significant rebuilds of existing lines. The utility and its real estate team will spend time evaluating routing, environmental considerations, right-of-way acquisition, and constructability. Also included in their assessment scope definition are elements at the point of interconnection such as metering, protection relays and communication systems. Data centers will be required to comply with ERCOT telemetry and real-time monitoring requirements.

Interconnection Agreement

After the utility defines its required engineering and construction, it will present the data center with an Interconnection Agreement (IA). The purpose of the IA is to establish the terms and conditions for connecting the data center to the grid including engineering requirements, construction coordination, curtailment rights, and timeline obligations. The data center developer and utility will both sign this agreement which essentially commits both parties to construct the project under the agreed-upon terms.

At the time of signing the IA, the data center developer will also be required to provide financial security usually in the form of cash or letter of credit. Over the course of construction, the utility may draw on this security to recover funds should the data center develop terminate the project or fail to make their required payments to the utility.

ERCOT Approval and Authorization to Energize

The last step in the interconnection process is the Authorization to Energize (ATE). The utility will request ATE from ERCOT once certain conditions are satisfied:

- The construction of the required interconnection facilities are complete
- Protection systems are in place and tested
- Telemetry is in place and operational
- All metering and ERCOT reporting requirements have been met by the developer

After these items are confirmed and the installed system is consistent with the studied configurations, ERCOT will grant the data center and utility approval to energize the substation(s) and facility. The facility can begin to bring their facility online and this may be done in sequential stages through a commissioning process. Commissioning may take 30-60 days.

Regulatory and Policy

This section focuses on policy matters, with a particular focus on the U.S. where policy development is most active. Regulators are extremely active in shaping policy as AI data centers are not “normal” commercial loads and have a considerable impact on the U.S. grid and the economy. Data centers often request tens to hundreds of megawatts and want to develop at speed to bring their facilities online as quickly as possible. The high concentration of requests and the speed create three regulatory challenges:

- Interconnection queues have been traditionally built for generator projects and smaller loads. An interconnection process within each ISO for large loads is still in early emergence
- Cost allocations can become very challenging to allocate when the grid and the ISO are attempting to accommodate data center requests and large generator requests simultaneously
- Rules/regulations related to system operation and reliability need to adapt given data centers can ramp quickly and in some cases are paired with on-site generation. Forming new regulations need to include federal agencies, ISOs and states

Federal policy:

The Federal Energy Regulatory Commission (FERC) governs transmission interconnection processes and tariffs across all interstate transmission systems. They are that most critical entity creating Federal policy for both generator and large load interconnections. In 2023, FERC introduced new rules related to generator interconnection. These new rules were aimed to reshape the generator interconnection process by creating cluster studies and requiring more significant financial investment. The goal was to reduce queue backlogs and speed up the process timing.

The Department of Energy (DOE) has urged FERC to create much more explicit and standardized procedures specifically acknowledging the importance of better accommodating AI data center demands. FERC has also opened consideration for a targeted policy that fast-track interconnection for large loads that are co-located with generators. The thought there is to avoid a portion of capacity upgrades that would be otherwise needed if the generation component were not installed at the data center facility.

NERC:

The North American Electric Reliability Corporation (NERC) will play a key role in further defining reliability expectations for large load interconnections. NERC is focused on solving the gaps where traditional reliability standards did not expect such large loads to come online with very dynamic ramping nature. NERC has made recommendations to address the situation that include more stringent modeling requirements and study processes when a large load seeks interconnection. NERC is also providing clarifications on how system operators should treat large loads during contingency events. They have published draft reliability guidance addressing risk mitigation for large loads that recommends a mix of contractual requirements and interconnection terms and conditions.

Independent System Operators:

Interconnection procedures are implemented by the utilities (transmission owners) and by the respective Independent System Operators (ISO) for the region. At this point in time, each are taking somewhat unique approaches to accommodating large load requests.

- PJM is accommodating large load requests by adding them to long-term planning for transmission. They are also developing means for large loads to participate in demand response programs for financial incentives. PJM is also developing new financial security policies that require data center developers to increase their upfront financial commitments
- MISO is developing improved long-term load forecasting to better predict future load growth. The forecasts will consider future scenarios that account for economic, technology and policy changes. MISO also now requires that large loads make an initial financial commitment of 25% of the construction costs when an interconnection agreement is executed between the utility and the data center developer.
- ERCOT/Texas has deployed an interim large load interconnection process in effort to manage the initial rush of large load requests. Senate Bill 6 addresses large loads greater than 75 MW, and is aimed to increase oversight and increase financial security requirements through the

interconnection process. It will require that the PUCT create new standards for large loads and also require that large load developers share the interconnection costs along with rate-payers.

State Policies:

Several states have begun to create obligations specifically related to data centers and large loads. For example, in many states, developers are required to reveal if they have duplicative interconnection requests. States are also supportive of implementing minimum billing and demand charge protections for utilities, which will help prevent stranded utility investments. Some states are also supportive of specialized interconnection pathways where a data center can seek an expedited interconnection process if they co-locate with generation. The aforementioned Senate Bill 6 in Texas is a good example. This reflects a broader trend where states and their legislators are responding to local pressures and reliability concerns related to the grid. On the other hand, there are some states that are offering tax exemptions and incentives for data centers to attract investment. Minnesota is a good example which provides a 20 year tax exemption on certain equipment and energy.

Cost Allocation:

A critical focus of current policy development is to determine how network upgrades, both at transmission and distribution levels, should be paid for and who is responsible for those payments. Federal and state policies are leaning towards a cost causation principle that essentially requires the new large load developer to be responsible for the costs. This is to include network upgrades and local interconnection facilities. Reform proposals from the DOE and other organizations have explicitly supported this principle while also calling for a clear methodology that also considers pending generation interconnection requests. The idea is to be able to properly delineate what projects, either large load or generator interconnections, are triggering the need for the upgrades and allocating the cost accurately.

Regulators need to balance the cost sharing principle so it is done in a fair manner and does not discriminate against certain types of projects and also considers what, if any, costs are fair to be placed on the general electricity rate-payers in the state or region. Tools within these policies include strong financial assurances, milestones and deadlines within the interconnection queue process, minimum electric billing demands, and finally the contractual requirements on cost responsibility for the needed upgrades. These policies are intended to reduce queue clogging and process time, enhance grid reliability, and share cost responsibilities in a fair manner.

Generator Co-Location:

Many data center developers are seeking to co-locate generation or storage resources at their facility to reduce or eliminate the need for grid power. This can speed up their implementation timeline and give them more control over their power supply resources. Regulators are considering what impacts these co-location arrangements might have on the broader wholesale power markets, as these resources are reducing grid reliance and diverting resources away from arrangements where they might otherwise be connected to the transmission side of the system and participating in the wholesale market. Regulators also have concerns about behind-the-meter generation paralleling with the grid, which could create local reliability issues if these generators are not properly studied or not properly operated.

Reliability Studies:

There are increasing requirements for large loads to be studied for dynamic and transient behavior as well as voltage ride-through considerations. ISO and NERC guidance are pushing for more stringent modeling of large load behavior, including the potential for automatic disconnects from the system under certain conditions. Regulators are also contemplating if data center loads should have mandatory communication telemetry, strict load ramping requirements or curtailment obligations. The goal is to be able to safely add and manage these large loads on the grid and not impact reliable power delivered to the local community.

Cybersecurity:

AI data centers can be considered critical assets and often will be placed on federal lands. For these reasons and the sensitivity of data involved, regulators and agencies are attentive to cybersecurity. The U.S. government is prioritizing secure buildout of AI datacenters and are attempting to strike a balance between access and security/resilience. The National Institute of Standards and Technology (NIST) has developed an AI Risk Management Framework (RMF) and is the principal standard in the U.S. It addresses IT resilience across the full lifecycle of a project. AI operators are being asked to follow the RMF framework alongside more traditional NIST security design measures. In the EU, the AI Act has been created which placed specific obligations on developing AI systems to promote strong security standards and required documentation through the lifecycle. The requirements also allow for facilities to be audited for compliance which will help ensure

Overall to summarize the above policies, they are aimed to promote the following:

- Standardize large load requests and study processes to eliminate duplicative interconnection filings and allow flexibility for unique geographic constraints
- Provide clear cost causation rules so cost responsibility can be accurately and fairly allocated between the data center developer, any generator interconnections in the region, and the rate-payers
- Increase overall transparency of data center developments and their large load connection requests
- Improve and standardize reliability requirements. Interconnection agreements should clearly define telemetry requirements, ride-through specifications, and curtailment clauses
- Standardize generator co-location frameworks that allow for on-site generation without harming the integrity of the overall wholesale market

Summary

The data center landscape is in rapid expansion with developers like Google and Oracle racing to deploy computing capability to support the growing need for AI technologies. It is expected that the total investment to AI data centers will ultimately require trillions of dollars over the next decade to meet the demand. Virginia has long stood as the primary hub for data centers, however that is changing quickly. New hubs include Phoenix, Dallas/Forth Worth and Silicon Valley, however the emergence is coming more and more widespread.

Developers are seeking to locate near existing electrical infrastructure with at least an existing capacity that can serve a large portion of the anticipated load. Their goal is to minimize the required network upgrades and interconnection facilities that the respective utility companies will need to perform. There are many other considerations including the availability of high speed fiber in the area and local water supply to support the IT communications and the water use for cooling needs. To meet the cooling needs, most AI data center developers are pursuing water-based cooling systems, where chilled water is delivered directly to the GPU racks and heat is exchanged via small plate heat exchangers. Water consumption associated with a cooling tower to reject the heat to the atmosphere can be significant.

Developers will consider several options for their power supply including the grid and various on-site resources. On-site resources can include natural gas generation (turbine or engines), fuel cells, battery energy storage, fuel cells and SMRs. Each provide their own distinct advantages and the engineers planning the facility will evaluate each along with a life cycle cost analysis. These sources can either supplement grid power or replace it in a microgrid configuration.

The large load interconnection processes are evolving across the country. ERCOT in Texas provides a nice example of an ISO that has set in place a specific process for large loads. Other ISOs and utilities are putting forth similar processes where proposed large loads will go through a series of large load studies, with a keen focus on grid reliability. These studies require a study deposit usually exceeding $100,000. To require further financial commitment, data center developers are also required to provide financial security at the execution of an interconnection agreement. This provides further “skin in the game” where if the developer decides to back out, the utility can draw funds against the financial security that was posted (cash or letter of credit). The goal of these interconnection processes is to ensure grid reliability and require that data centers are financially invested and committed to the projects.

The regulatory and policy landscape for AI data centers continues to evolve and must balance grid reliability versus the growing national priorities associated with AI data center growth. The issue has drawn attention at the Federal, State and local levels. FERC will continue to propose policy aimed to properly manage large load interconnections across the interstate transmission system in the U.S. NERC has the challenge of creating new reliability initiatives to ensure data centers being added to the system have adequate telemetry and protection equipment and operational protocols. Near term policies will require increased standardization and stricter financial and disclosure rules. Developers and utilities will find success from proactive planning, transparency in data sharing, and early planning.