MOORE INSIGHTS

The AI Infrastructure Buildout: From the Power Grid to the Rack

Artificial intelligence may feel like a software story, but its growth depends on an expanding physical network of data centers, power systems, and cooling equipment. At the end of 2025, hyperscale operators ran about 1,360 large data centers worldwide. In the U.S., data centers used roughly 4.4% of the country's electricity in 2023, and the Department of Energy estimates that share could rise to between 6.7% and 12% by 2028.

That helps explain why the AI buildout reaches far beyond chips and servers. Electricity has to be generated, moved to the site, converted into a usable form, and protected from interruptions. Cooling systems must continuously remove the heat produced by increasingly dense computing equipment. Inside the building, racks of processors, memory, storage, and networking equipment convert that electricity into computing capacity.

Together, these systems form a continuous chain that stretches from the power source to the individual processor. The buildout can be understood in terms of two interconnected layers: data center electrification and data center racks.

Data Center Electrification: Bringing Power to the Facility

The first layer begins outside the building. Most data centers connect to a regional grid supplied by some mix of natural gas, nuclear power, renewables, hydroelectric facilities, and other sources. Some operators also use batteries, fuel cells, on-site generation, or microgrids to improve reliability.

Datacenter Electrification

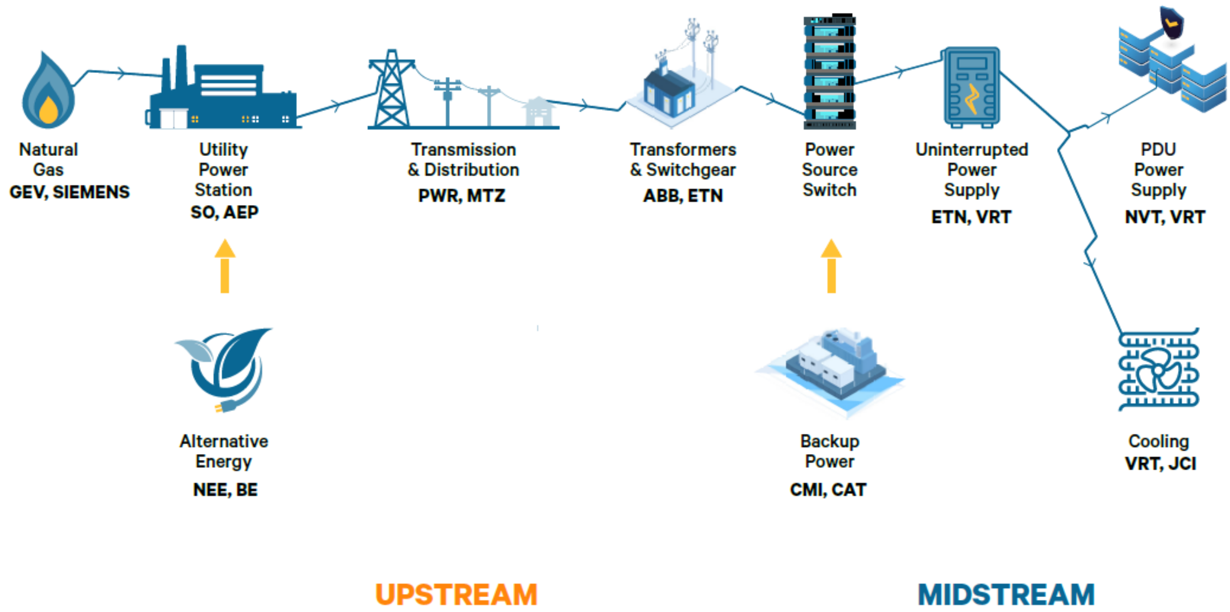

Source: Ned Davis Research

Utilities and power producers supply electricity, while engineering and construction firms build the connections needed to move it. Notable participants here include GE Vernova (GEV), American Electric Power (AEP), Quanta Services (PWR), and MasTec (MTZ).

For a large data center, access to electricity can be as important as access to land or fiber. These facilities may require power on the scale of a major industrial site, while transmission lines, substations, and utility connections can take years to complete. A project may have financing, demand, and equipment in place but still wait months to years for power.

Once electricity reaches the property, several systems prepare it for the computing equipment. Transformers adjust the voltage. Switchgear protects the facility from faults and allows parts of the electrical system to be isolated safely. Transfer equipment moves the load between power sources. ABB (ABB) and Eaton (ETN) are active in these areas.

Data centers also need protection from outages. An uninterruptible power supply, or UPS, provides immediate power if the grid becomes unstable or goes down. Backup generators can then support the facility during a longer interruption. Vertiv (VRT) and Eaton are well-known suppliers of UPS systems, while Cummins (CMI) and Caterpillar (CAT) are associated with backup generation.

Power then moves through the building to individual racks. Larger distribution systems carry electricity across the facility, while rack-level power distribution units, or PDUs, deliver it to servers and other devices. Eaton, Vertiv, and nVent (NVT) all participate in this part of the chain.

Cooling Becomes Part of the Power Story

Nearly all the electricity used by computing equipment eventually becomes heat. Traditional data centers have relied mainly on cooled air, but denser AI systems are increasing the use of liquid cooling. These systems draw heat away from processors through cold plates and coolant lines.

A conventional rack may use less than 10 kilowatts. By comparison, NVIDIA says its GB200 NVL72 system can draw about 120 kilowatts at full load. That difference affects more than the electric bill. It changes the size of the utility connection, the required backup power, and sometimes the building design.

Power and cooling therefore have to be planned together. A higher-density rack requires more electricity and a more effective means of removing the resulting heat. Once those systems are in place, the focus moves inside the facility, where electrical capacity becomes computing work.

Data Center Racks: Turning Electricity Into Computing

Once electricity reaches the rack, it becomes computing capacity. A modern AI rack is more than a cabinet filled with separate servers. Its processors, memory, storage, networking, power, and cooling are designed to work as one coordinated system.

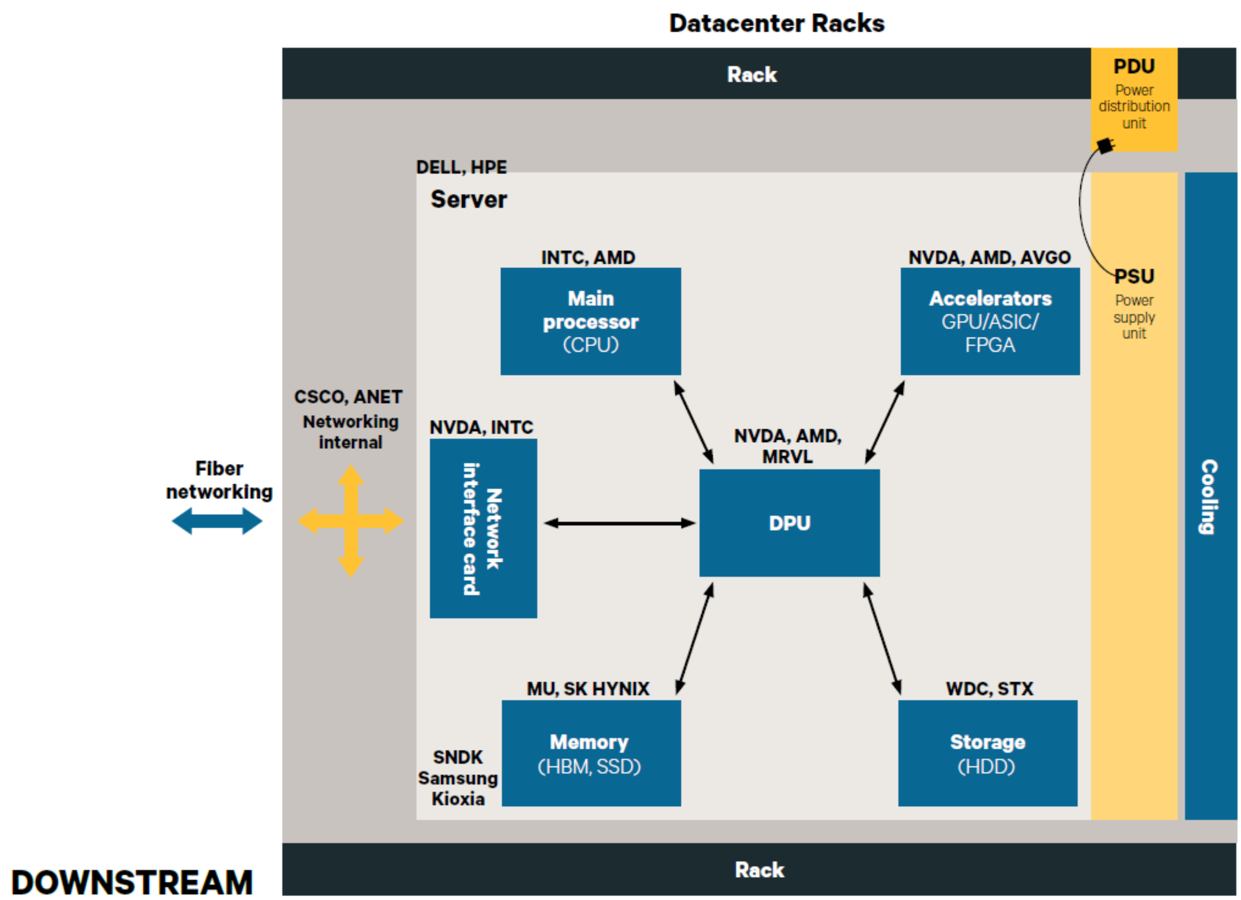

Source: Ned Davis Research

The central processing unit, or CPU, manages general computing tasks and coordinates the broader workflow. Graphics processing units, or GPUs, and other accelerators handle much of the intensive mathematical work involved in training and running AI models. Companies such as Dell Technologies (DELL) and Hewlett Packard Enterprise (HPE) assemble these components into server and rack systems, using processors from companies including Intel (INTC), AMD (AMD), and NVIDIA (NVDA).

NVIDIA's GB200 NVL72 shows how much computing can now be concentrated in one place. The liquid-cooled system combines 36 CPUs and 72 GPUs in a single rack-scale design. At this level of density, computing equipment, power delivery, networking, and cooling cannot be planned separately.

The processors also need a steady flow of data. High-bandwidth memory, or HBM, keeps information close to the accelerator while calculations are underway. Micron Technology (MU), SK Hynix, and Samsung are major suppliers in this area.

Storage holds larger amounts of information for later use. Solid-state drives provide faster access to active datasets and saved models, while hard-disk drives offer lower-cost capacity for information that is used less frequently. Put simply, memory helps the processor work in the moment; storage preserves information over time.

Networking connects the system. High-speed switches, network cards, and fiber links move data among servers, racks, and larger computing clusters. Arista Networks (ANET) and Cisco Systems (CSCO) are among the companies participating in this layer.

Power and cooling remain essential inside the rack. Electricity must be distributed and converted for the individual components, while the heat they produce must be removed continuously.

The important thing to understand here is that every component affects the others. Faster accelerators require more memory bandwidth and network capacity. More computing creates greater power demand and heat. The rack can therefore perform only as well as its most constrained component.

How the Two Layers Work Together

Data center electrification and rack infrastructure are useful for understanding the buildout, but they cannot be planned separately. Available power limits how much computing equipment a site can support. At the same time, the selected racks determine the transformers, backup systems, power distribution, and cooling capacity the facility requires.

In other words, the building must be designed around the computing load, while the computing equipment must operate within the building's physical limits.

A bottleneck anywhere in this chain can slow the entire project. A development may initially wait for a grid connection, then face shortages of transformers or cooling equipment. Once operational, the constraint may shift to memory, storage, or networking that cannot keep pace with the processors.

The Physical Foundation of AI

The AI infrastructure buildout begins well before the chip. It includes power generation, transmission, electrical equipment, backup systems, and cooling. Inside the data center, the chain continues through servers, accelerators, memory, storage, networking, and rack-level power systems.

Viewed end to end, it is one connected process: electricity reaches the facility, becomes computing work inside the rack, and ultimately becomes heat that must be removed. Every layer must support the next.

Tom Moore is the founder of Moore Invested, a fee-only investment management firm based in Atlanta, Georgia. With more than 25 years of experience, he has helped hundreds of investors prepare for and live comfortably in retirement. Moore Invested provides fiduciary investment management for individuals serious about retirement planning. Learn more at www.mooreinvested.com.

Moore Invested Disclosure

Advisory Services offered by Arkadios Wealth. Moore Invested and Arkadios are not affiliated through any ownership.

Past performance does not guarantee nor is indicative of future results. This summary of statistics, price, and quotes has been obtained from sources believed to be reliable but is not necessarily complete and cannot be guaranteed. All securities may lose value, may not be insured by any federal agency and are subject to availability and price changes. Market risk is a consideration if sold prior to maturity. Information and opinions herein are for general informational use only and subject to change without notice.