ethernet alliance

White Paper
Evolving Infrastructure Requirements for AI Data Centers

Contributors:
Bob Wagner, Sr. Business Development Manager, Panduit
Jose Castro, Distinguished Manager, Optical Research, Panduit
Brian Kelly, Sr. Solutions Architect, Panduit
Justin Blumling, Technical Sales Engineer, Panduit

December 2025

ethernet alliance

Table of Contents

Table of Contents

Introduction

Artificial intelligence is transforming every industry. No other technology since the advent of the internet has brought such a wave of change in how business is done and how data centers need to be built. This paper reviews the emerging data-center infrastructure requirements for AI deployments, focusing on power density, cooling approaches, network-design and cabling considerations, and protocol support.

With AI models growing in size and complexity and demanding the fastest accelerators, rapid changes in data-center infrastructure requirements are unfolding. Power demands continue to influence locations and cooling strategies, but designing high-bandwidth, low-latency network infrastructure has become equally critical. Once a peripheral consideration, the network is now central to data-center performance, and its needs can be met with physical infrastructure capable of supporting high power, dense fiber cabling, and scalable growth for even the largest AI facilities.

2. Power Density

In 2014, a group of data-center users was asked to predict their needs for power density in 2025, and 70% of respondents said the average would be >= 40kW per rack [1]. With time and hindsight, we can see this power envelope was certainly pushed by distinct users in scientific research, gaming, finance, and big data, as according to Uptime Institute's [2] recent data, the average kW per rack in 2024 was only 8kW. However, the 2014 forecast is prescient, because the advent of AI has created designs needing 40kW to 140kW today. An Nvidia roadmap shows the Kyber system at 600kW per rack by 2027. The increased power demand that we see today, coupled with the expectation that AI will exceed 50% of the total data center hardware spend by 2030, will lead to a rapid rise in average rack power.

The sudden increase in demand for AI has brought substantial stress to data-center owners over the past two years, as they had to pivot to new building designs while also looking at retrofitting even newly built data centers. Large-scale AI data centers often require hundreds of megawatts and scale to several gigawatts in the most advanced facilities. To put this in perspective, these power demands are comparable to small- or mid-sized cities. This new normal is challenging the way data center buildings get their power. Due to multiple year lead times to access more power from the grid, many data centers are deciding to generate their own power via renewable sources such as wind and solar, through gas-powered generators, and small modular nuclear reactor plants (SMR). In addition, the higher power demands have made 415v to the rack essential for any new data-center building and limiting most buildings that can only offer 208v from housing AI equipment.

In the rack, this means having either more higher power distribution units (PDUs) or opting for a DC busbar. Most traditional data centers have 2 PDUs per rack, as that provides sufficient left-side/right-side power redundancy. However, as shown in Figure 1, the quantity per rack, current capacity, and style of PDUs change significantly for AI.

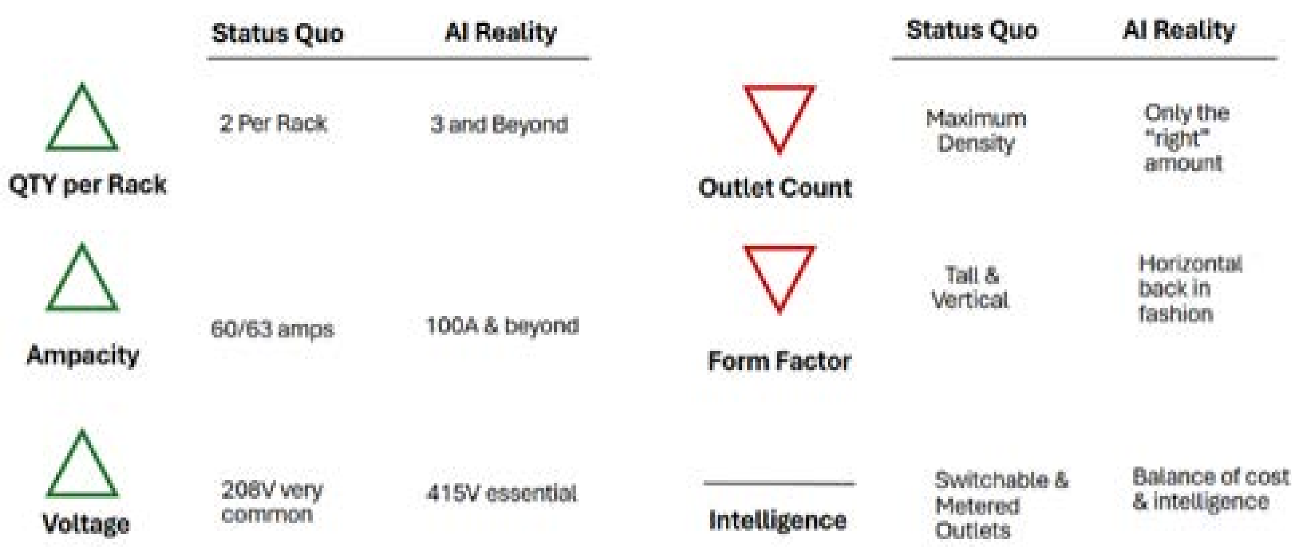

Figure 1: Changes to Data Center Power Requirements
Source: Panduit

When the latest generation of AI racks started requiring 120+ kW, rack PDUs were replaced by Open Compute Project (OCP) style power distribution, which uses power shelves and a 48V DC-powered busbar. While the v3 OCP rack was initially designed to be used to supply 30kW, the spec had to be quickly upgraded to enable the newest generation of GPUs. Such densities not only challenge the power delivery systems in a data center; they also require specialized cooling technology and techniques.

3. Cooling Approaches

Supercomputers from IBM, Cray, and others have successfully used liquid cooling since the 1970s. However, until recently, most liquid cooled systems have been confined to high-performance clusters in academia, government agencies, and US National Laboratories. Today, most data center companies are spending a lot of time researching the best way to supply liquid cooling.

With the advent of mainstream graphic processing unit- (GPU-) centric systems, new cooling designs like direct-to-chip (DTC), immersion, and re-imagined approaches to existing technologies like rear-door heat exchangers (RDHX) have become essential cooling methodologies. There continues to be disagreements about whether water or dielectric refrigerants are best, but nearly all customers are looking to add liquid cooling.

Per Figure 2 below, air cooling is generally considered to be insufficient at greater than 50kW per rack, above which RDHX becomes the preferred method. As systems like Nvidia's NVL72 exceed 120kW per rack, DTC cooling is needed. Often a RDHX is used in conjunction with DTC, which usually covers only the GPUs and possibly the central processing units (CPUs). Immersion-cooling technologies are even more efficient than DTC but require significant infrastructure changes and add further complexity to the data center.

Figure 2: Type of Cooling to use based on Rack Power
Source: Panduit

4. Network Design and Cabling Considerations

Along with the changes in power requirements and cooling approaches, AI is dramatically increasing fiber density by up to 8 times higher than traditional data centers, necessitating larger pathways and more efficient cable management strategies. This shift—driven by higher data rates, flat rail-optimized network topologies, and multiple networks within AI systems—makes robust physical infrastructure more critical than ever.

AI systems have four discrete networks, each with specialized functions. Each network requires its own cabling and runs at a potentially different data rate than the others:

- Compute network (back end), 400G -> 800G
- Storage network, 200G -> 400G
- In-band management (front end), 100G
- Out-of-band management, <10G

The highest data rates are used for the compute network, which connects the GPUs to each other. Most GPU-based servers use 400G ports (4×100G) for compute and fewer lanes for storage and in-band networks, requiring splitting the cabling for those connections. Currently, the compute back-end networks use 400G QSFP112 and OSFP transceivers, which connect via one eight-fiber multi-fiber push on (MPO) cable, or 800G OSFP transceivers, which connect via two eight-fiber MPO cables as shown in Fig. 3. However, Dell'Oro forecasts rapidly increasing data rates with 1.6Tb transceivers becoming dominant by 2027, then superseded by 3.2Tb in 2030. [3] The 1.6Tb modules will use either two 16-fiber MPO cables at 100G/lane or two eight-fiber MPO cables at 200G/lane, which is just now becoming available.

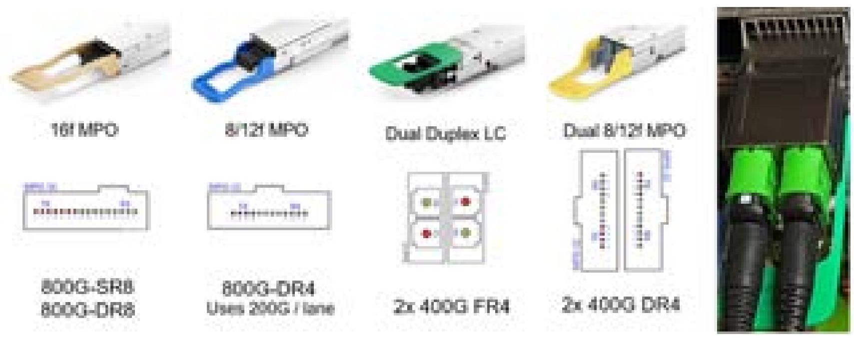

Figure 3: 800G Optical Transceiver Types
Source: Cisco

The compute back-end network uses a specialized Clos (scalable, non-blocking Spine-Leaf) topology called "rail-optimized topology" to minimize tail latency and network interference by flattening the network and reducing the number of switch hops between GPUs (as shown in Figure 4). This is important as the goal of any AI network is to minimize latency between the GPUs so that the entire system can act as one large brain. However, precise cable routing and mapping are required to correctly install and maintain a rail-optimized network.

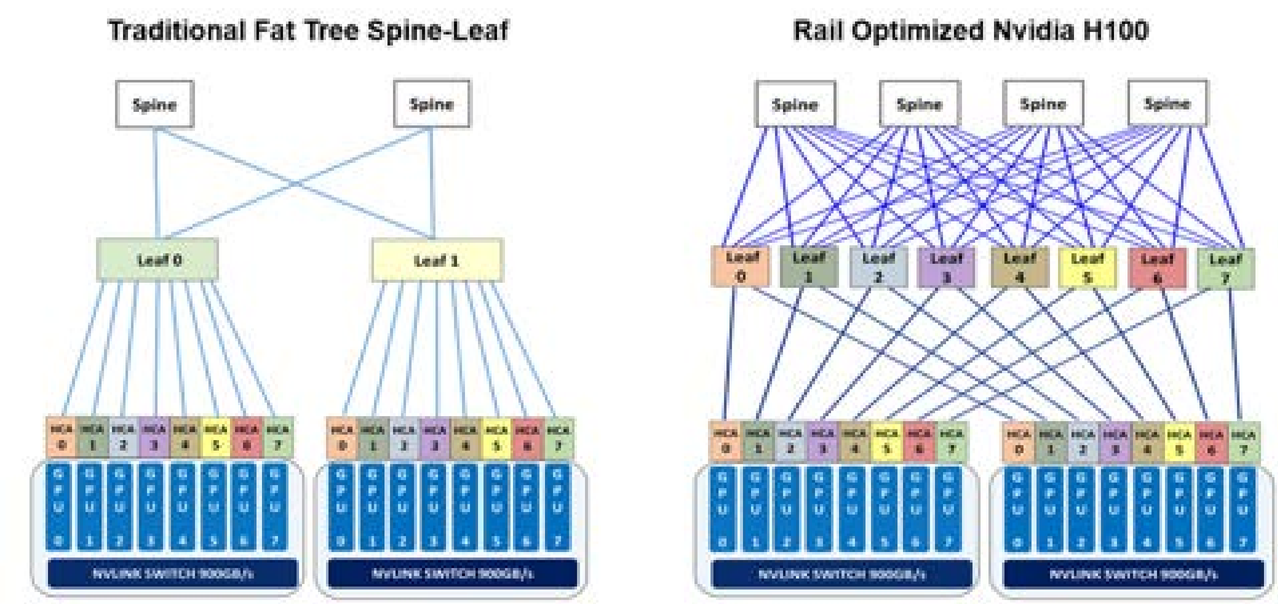

Figure 4: Comparison of AI Rail Optimized to Traditional Spine-Leaf Architecture
Source: Panduit

Optical networks typically combine multimode fiber (MMF) for short server-to-leaf switch links and single-mode fiber (SMF) for longer inter-switch connections. Cabling can be deployed as point-to-point, using a single cable per connection, which is cost-effective and suitable for small deployments, or as structured cabling, which is preferred for larger or expandable networks. Structured cabling helps organize and protect cables while simplifying maintenance by clearly segmenting the network. This allows technicians to isolate faults, swap components, and perform upgrades with minimal disruption. It also consolidates multiple smaller fiber cables into high-count trunks, reducing pathway congestion. Figure 5 illustrates how trunks can increase fiber capacity in a pathway by 400%.

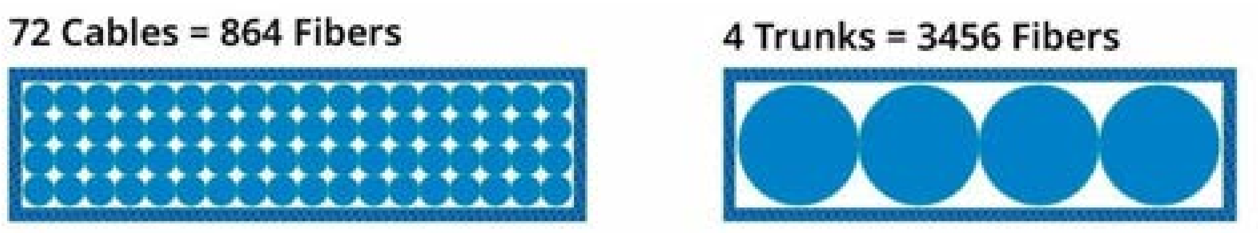

Figure 5: Pathway density comparison of 12 fiber cables to 864 fiber trunks
Source: Panduit

Structured cabling, which uses patch panels or enclosures, also offers areas for storing excess cable slack which is a common problem in larger data centers where it is harder to correctly estimate cabling length. One method for reducing installation time is to hang the patch panels from the pathways above the racks so the cable in the overhead pathways can be installed well before the rack arrives. This leaves only the simple tasks of plugging in the power and connecting short patch cords from the server to the overhead panel to have the new racks up and running (see Figure 6).

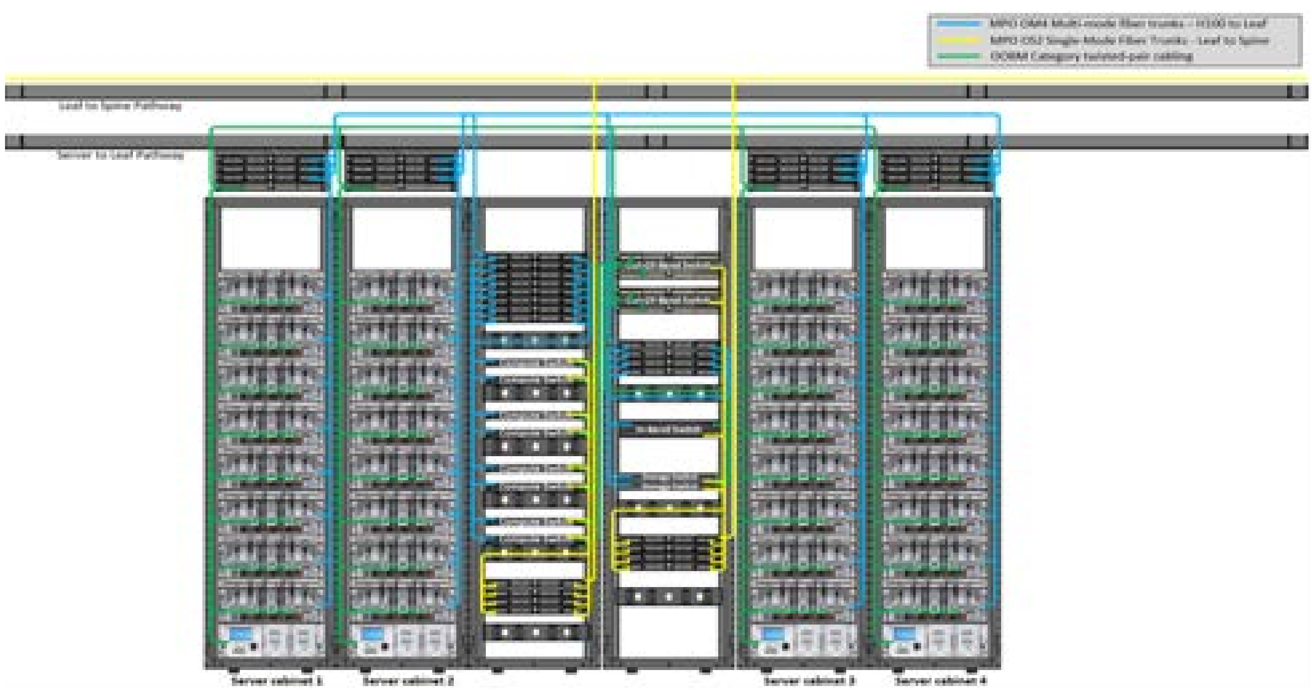

Figure 6: Liquid Cooled B200 SuperPOD with Overhead Patch Panels
Source: Panduit

5. Protocol Support

Front End traditional networks including communication, storage and management have successfully relied on Ethernet for many years. However, when the first Back End Networks were introduced, they utilized InfiniBand due to their legacy of coming from High Performance Computing (HPC) which values low latency. Since most IT employees have significantly more experience running Ethernet and there is a much larger eco-system of Ethernet-based products, the challenge became how to make Ethernet work more efficiently on Back End Networks.

Back End Networks are used for handling training and inferencing. This data critical communication between GPUs relies on Remote Direct Memory Access (RDMA), a core feature of the InfiniBand protocol, which allows one computer to read or write directly into another's memory without involving the CPU to reduce latency. To operate on Ethernet Networks, RDMA over Converged Ethernet (RoCE) was developed in 2010 and used on many HPC systems. RoCE enables data centers to provide a high bandwidth, low-latency, lossless fabric comparable to InfiniBand, while benefiting from the larger, more widely used Ethernet systems.

In 2023, in response to the AI revolution started by the release of Chat GPT, the Ultra Ethernet Consortium (UEC), was formed by the hyperscalers and major electronics vendors, with the goal of developing an Ethernet-native transport that preserves RDMA's benefits of low latency, while addressing RoCE's weaknesses for massive scaling to millions of GPUs. The UEC system will also introduce more efficient traffic management, supporting communication among large numbers of accelerators and enabling AI network expansion.

In parallel, the IEEE 802.3 standard is driving major increases in Ethernet transceiver speeds enabling 200G per lane which enables 800G over 8 fibers and 1.6Tb over 16 fibers. Following the lane assignment by IEEE 802.3dj, breakout signal configurations called shuffles can be produced to reduce the number of layers of switches creating flatter, more efficient network topologies that offer lower latency. Based on these developments and customer feedback, Dell O'ro's president Sameh Boujelbene recently stated, "As the industry moves toward 800 Gbits/sec and beyond, we believe Ethernet is now firmly positioned to overtake InfiniBand in these high-performance deployments."[4]

6. Conclusion

AI is dramatically changing the data-center landscape via its need for much higher power and new, rapidly evolving network architectures. Data-center owners are realizing that many of their existing buildings are no longer well positioned for the future and will be looking for support in creating purpose-built AI facilities that will use innovative network infrastructure for power, cooling and network cabling. The good news is that the Ethernet Alliance and many industry groups are working to help smooth this transition through industry education and new standards to improve efficiency and interoperability.

ethernet alliance

Citations

1. *[1] Emerson Network Power "Data Center 2025: Exploring the Possibilities" (2014)*
2. *[2] Uptime Institute "Uptime Institute Global Data Center Survey Results 2025" (2025)*
3. *[3] and 4.[4] Dell'Oro "Ethernet is Winning the War Against InfiniBand in AI Back-End Networks" (2025)*