



Data Center Interconnect Enables AI Data Centers to Scale Across

In the context of this article, we will understand data center interconnect (DCI) to cover the physical infrastructure that connects two or more individual data centers together—allowing us to explore the use cases for DCI and methods of transmission. The geographic scope of DCI varies greatly, as it may include multiple data centers across a connected campus; however, it is more commonly used to describe clusters of data centers across several cities (also known as an “availability zone”). At even larger scale, data center networks spread across the world are connected to create a global DCI.

Data center interconnect and the importance of scaling across for AI

Why has DCI become an important topic in today’s data center world? DCI has existed for a long time and has always been important; however, there’s been a significant recent change brought about by AI applications that has made it a critical parameter in the overall performance and economics of training large language models (LLMs)—and that change is centered on scaling.

To put this into context, a scale-up network is used to interconnect a large number of GPUs in a single cabinet with one another, forming a single block of compute. When the power, space or cooling to host more GPUs in that cabinet is exhausted, adjacent cabinets within the same data center or data hall are connected together to create a scale-out network. Scale-out networks lead to a significant increase of both compute capability and the amount of electrical power required to keep the GPUs running. Power availability is a common logistical limiting factor to the amount of infrastructure that can be housed in a single data center building. This drives the need to build new data centers, either on a campus, across a city,




Figure 1: Basic architecture of a DCI

or across a region to host the additional compute resources required for the training of large and hyperscale size models.

For these model sizes that require distributed compute, geographically discrete data centers must be orchestrated to work together, and this is where DCI plays a crucial role. In fact, it's become so important that there's a new term that has evolved to describe this DCI network: scale-across.

Data centers that are used to train the largest LLMs frequently no longer exist as individual isolated islands but have become part of an archipelago of buildings that have formed a highly interconnected scale-across fabric, expanding horizontally to seamlessly enable increased flows of east-west traffic patterns. This scale-across network is collectively known as an AI-optimized DCI.

According to the Dell'Oro Group, the annual capacity purchased for direct DCI is expected to show a threefold increase between 2025 and 2028, with the majority of interfaces shipped being 800G, 1.2T or 1.6T coherent optics, ZR and ZR+ transceiver standards.

New DCI traffic types

DCI traffic has become extremely varied due to the demands of AI. For example, today's DCI networks can expect to transport:

- AI model parameters (e.g., all-reduce, all-to-all)
- AI gradients
- GPU-GPU and node-node synchronization data

- Large latency-sensitive data flows
- Control and orchestration data for fault detection, scheduling, fabric telemetry, fault detection and recovery signaling.

What's important to understand is that the majority of this data is intermediate compute data that keeps the fabrics and distributed workloads synchronized, rather than user data (such as images or documents). Although these DCI traffic patterns are driven mainly by training activity, inference AI can also require large amounts of data transported via DCI to synchronize its operations depending on the architectural model being deployed. At present, AI training (rather than inference) is the biggest driver of traffic; however, this is expected to change in the future as the industry moves from primarily training to deploying more infrastructure to support inference.

AI presents multi data center DCI challenges

With large AI models having 70 billion parameters or more, and hyperscale models like ChatGPT 4 utilizing over one trillion parameters, there are key infrastructure challenges that come into play when managing compute across multiple data centers. For example, there will be high inter-region latency (often 20 milliseconds or more, [as explained by NVIDIA in this technical blog](#)), which can introduce performance bottlenecks during large-scale LLM training.

With this and many other factors in mind, DCI network infrastructure must be designed with sufficient optical fiber capacity to support low-latency and high-bandwidth channels, ensuring fast and unconstrained communication flow between data centers to maximize training efficiencies.

DCI transmission

There are a number of ways to transmit data between data centers over DCI. The technology can involve sending high-speed Ethernet via optical signals over dedicated fiber or wavelength services. In most cases, a DCI requires a high-speed WAN link, which could mean using MPLS, Ethernet, VPLS, Metro Ethernet, and so forth.

Selecting the right DCI infrastructure, architecture and topology depends on a wide range of variables, including the location of the data centers, the distance between data centers, bandwidth and availability requirements, the capabilities of local service providers, and security concerns.

There are just as many options and variables when you get down to the component level. Following are just a few of the solutions and developments that may figure into your decision.

Border edge (gateway) devices

Leaf-and-spine switches are integral to creating a data center fabric and are used to ensure efficient east-west transmission of data.

As a data center approaches cloud scale, the spine-layer switches are typically dedicated to facilitating the east-west data flow. A

new group of switches, known as “border edge switches,” has been added to the architecture to handle incoming and outgoing traffic at the edge of the data center fabric. In some cases, the job of handing off data to the DCI network can be performed at the leaf layer, using a border leaf switch.

Coherent optics

Coherent transmission has been used in long-haul and undersea links for about a decade, but we now see this technology migrating to the data center. Over time, the makers of coherent transceivers have reduced the size, power and cost of their optics to be more attractive for shorter and shorter links. IEEE Ethernet standardized the use of 100G and 400G coherent signaling on fiber links up to 80 km. A new project (802.3dj) will write standards for 800G coherent over just 10 km. The Optical Interworking Forum (OIF) has published the 400ZR and 800ZR standard specifying 800G transmission over a single wavelength with reaches of 80-120 km. Work on 1600ZR is currently underway at OIF.

As data centers continue to grow and become more interconnected, the bandwidth needs of the DCI network are also growing, reaching 100T over multiple wavelengths. This bandwidth

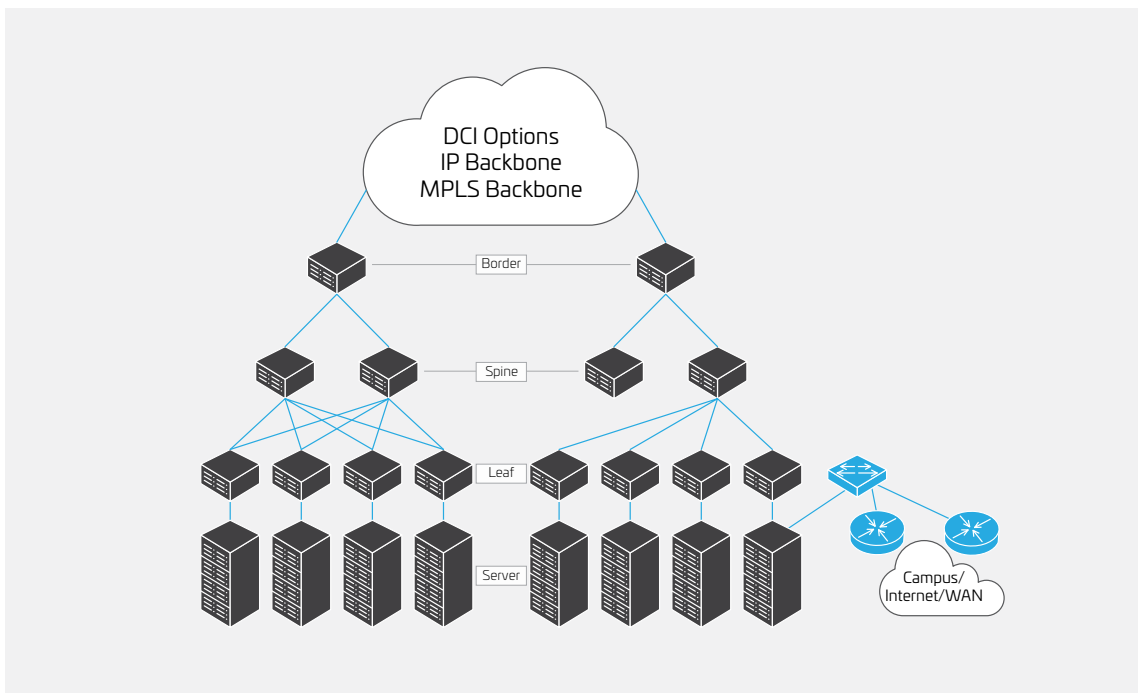


Figure 2: Example of a border edge switch facilitating DCI at the leaf layer




Figure 3: QSFP-DD coherent optical module

Figure 4: CFP2 coherent optical module

demand is supported by faster data rates—400G and 800G per wavelength—and will only continue to escalate.

Coherent optics are typically used for ultra-high bandwidth applications ranging anywhere from 100G to 1T over very long distances. Powerful digital signal processing chips (DSPs) are embedded within these systems to mitigate linear effects caused by fiber impairments, including chromatic dispersion and polarization mode dispersion.

Coherent fiber optics utilize the natural properties of light to optimize digital modulation practices and fiber-optic carrying capacity in long-range applications. However, coherent transmission will change to suit DCI applications. Supported by a proposed new standard (IEEE802.3dj), this will attempt to enable 800G links greater than 10 km.

Dense wave division multiplexing (DWDM)

DWDM enables multiple wavelengths of light to travel over the same fiber simultaneously, with each wavelength carrying a discrete signal.

Tight wavelength spacing can enable up to 96 channels on a single fiber. DWDM is a versatile transmission technology able to support coherent optics and on-off keying (OOK) amplitude signaling. When combined with coherent modulation, individual channel bandwidth can expand to 400 or 800 gigabits (see Figure 5).

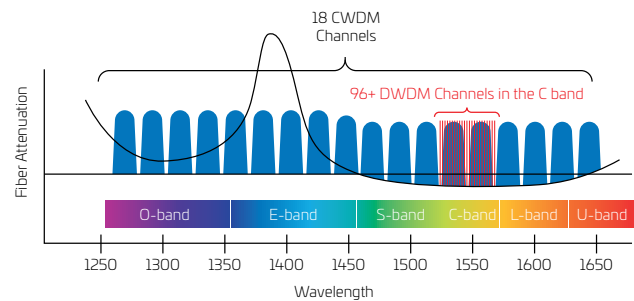


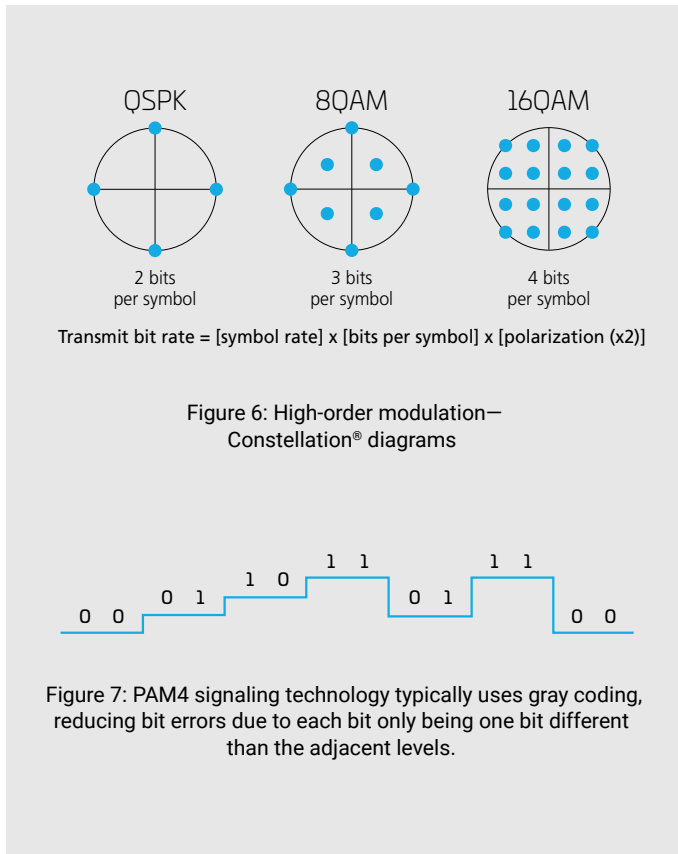
Figure 5: CWDM/DWDM band and channel configurations

QPSK coherent coding vs. PAM4

Coherent systems are based around phase shift keying (PSK)—phase modulation techniques that allow multiple symbols per bit to be encoded based on four phase shift orientations (e.g., 0°, 90°, 180°, and 270°). Many systems use quadrature phase shift keying (QPSK) to encode two bits per symbol. Dual polarization QPSK (DP QPSK) uses horizontal and vertical polarization along with QPSK to represent twice as many bits (see Figure 6).

Pulse amplitude modulation (PAM4) is a four-level modulation scheme designed for short-haul fiber links. PAM4 (see Figure 7) uses four amplitude pulses, each containing two bits, to double the bandwidth of conventional binary signaling. The simplicity and low power requirements of PAM4 make it a popular option for 100G and 400G Ethernet applications.

Unlike coherent optics, PAM4 is highly susceptible to fiber impairments. This limits range to ≤ 10 km, whereas coherent optical fiber communication systems can potentially span thousands of kilometers on amplified links.



New designs driving higher fiber-count DCI networks

The expansion of east-west data flows and the move to a low-latency leaf-and-spine switching architecture has created a tsunami of data inside and among data centers, particularly AI data centers where training LLMs is the key function. This has pushed the development of data center campuses with multiple buildings into overdrive. This trend isn't isolated to a few hyperscalers in the U.S.; it also impacts cloud and multi-tenant data center (MTDC) providers globally.

The physical layer

Once the external cable has been terminated inside the data center, it must be presented for splicing to an onward destination or connected to patching and cross-connection equipment. This is done using optical distribution frames (ODFs) located in the meet-me room (MMR) or main distribution area (MDA), where all network cabling comes together and is prepared for distribution.

To ensure this zone can support future growth and is manageable, patch cord hygiene should be mandatory. As noted earlier, the role of the DCI network across the campus is increasing; in the near future, they will have to support throughput of 100T or more. This will require thousands of fibers—all converging at the MMR or MDA, underpinning the importance of good patch cord management. So, too, ensuring the ODF has both fiber patch cord routing and slack management will be key to ensuring the MDA and MMR can support all future growth needs.




Figure 8: Infrastructure components for a typical campus-scale DCI




Figure 9: Entry into the data center building or MTDC meet-me room

To the future

Increasing deployment and expansion of DCI networks across the campus and into the cloud will continue unabated. As AI demands more scaling in networks, more and more data is created and shared by GPUs that never take a break or grow tired. As a result, data loads and scale-across DCI network requirements keep increasing. This presents physical challenges for the campus network designer, especially if the network must interconnect data centers across a metro area or land owned by others.

If the data is to traverse these geographies seamlessly, the data-carrying capacity of the fiber cables must increase. This will be done either by adding more fibers per cable—via size-reduced cladding—or introducing a commercially viable alternative glass technology, like hollow core fiber for example. But regardless of specific technology, DCI will continue to grow in importance to the growth and evolution of AI data centers.

**By Alastair Waite, Senior Manager,
Global Data Center Market Development, CommScope**



This article was originally published as a chapter in CommScope's 2026 Data Center eBook, [available for free here](#).

CommScope pushes the boundaries of communications technology with game-changing ideas and ground-breaking discoveries that spark profound human achievement. We collaborate with our customers and partners to design, create and build the world's most advanced networks. It is our passion and commitment to identify the next opportunity and realize a better tomorrow. Discover more at commscope.com.

COMMSCOPE®
an Amphenol company

commscope.com

Visit our website or contact your local CommScope representative for more information.

© 2026 CommScope Technologies LLC, an Amphenol company. All rights reserved. CommScope and the CommScope logo are registered trademarks of CommScope and/or its affiliates in the U.S. and other countries. For additional trademark information see <https://www.commscope.com/trademarks>. All product names, trademarks and registered trademarks are property of their respective owners.

EB-117988.3-EN (03/26)