Across the Campus and into the Cloud: The Growing Need for Data Center Interconnect (DCI) Technology

Digitization is driving more traffic to more data centers across the network and between enterprise data centers, disaster recovery locations, and, eventually, to multiple cloud peering points. Data centers, therefore, no longer exist as individual and isolated islands but are now part of a highly interconnected ecosystem, making data center interconnects (DCIs) increasingly essential.

As shown in Figure 1, "DCI" refers to the infrastructure that connects two or more individual data centers involved in a common task. The geographic scope of a DCI varies greatly. Multiple data centers within a campus may define a campus-scale DCI. At the same time, clusters of data centers across several cities may constitute a regional DCI (also known as an "availability zone"). At a broader level, data center networks spread across the world are connected to create a global DCI. With more traffic needing to be transmitted between data centers, the DCI market is predicted to grow by over 16% CAGR between 2023 and 2028.[1]

Key transmission considerations

Traffic must be sent across the campus and out into the cloud rapidly and securely to maintain quality of service. There are a number of ways to transmit data between data centers in a DCI. The technology can involve sending high-speed Ethernet or optical signals over dedicated fiber (or wavelength services). A DCI usually requires a high-speed WAN link via multiprotocol label switching (MPLS), Ethernet, metro Ethernet, virtual private LAN services (VPLS), or other protocols.

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 requirements. When you get down to the component level, there are just as many options and variables. The following are just a few developments that may influence your decision.

Figure 1: Basic architecture of a DCI

Border edge (gateway) devices

Leaf-and-spine switches are integral to creating a data center fabric and are used to ensure efficient east-west data transmission. As a data center approaches cloud scale, 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, as shown in Figure 2.

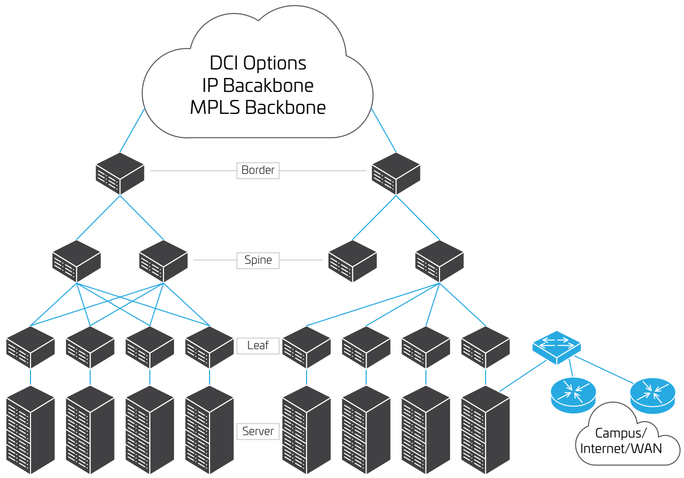

Figure 2: Placement of border-edge switches

Coherent optics

Coherent transmission has been used in long-haul and undersea links for about a decade, but this technology is now 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 signaling on links up to 10 km. Figures 3 and 4 show two of the main pluggable coherent optic modules—the QSFP-DD and the CFP2. The QSFP-DD has the smallest footprint and consumes less power for less demanding applications, while the CFP2 has a larger footprint and highest power dissipation for improved performance and reach.

Figure 3: QSFP-DD coherent optic module

Figure 4: CFP2 coherent optic module

As data centers continue to grow and become more interconnected, the bandwidth needs of the DCI network are also increasing, reaching 100 terabytes per second (Tbps) over multiple wavelengths. This bandwidth demand is supported by faster data rates—400 gigabits per second (Gbps) and 800 Gbps per wavelength—and will continue to escalate.

Enabled by the IEEE802.3ct standard, coherent optics are typically used for ultra-high bandwidth applications ranging 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 over just 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, as shown in Figure 5. DWDM is a versatile transmission technology that supports coherent optics and on-off key (ook) amplitude signaling. When DWDM is combined with coherent modulation, individual channel bandwidth can expand to 400 or 800 Gbps.

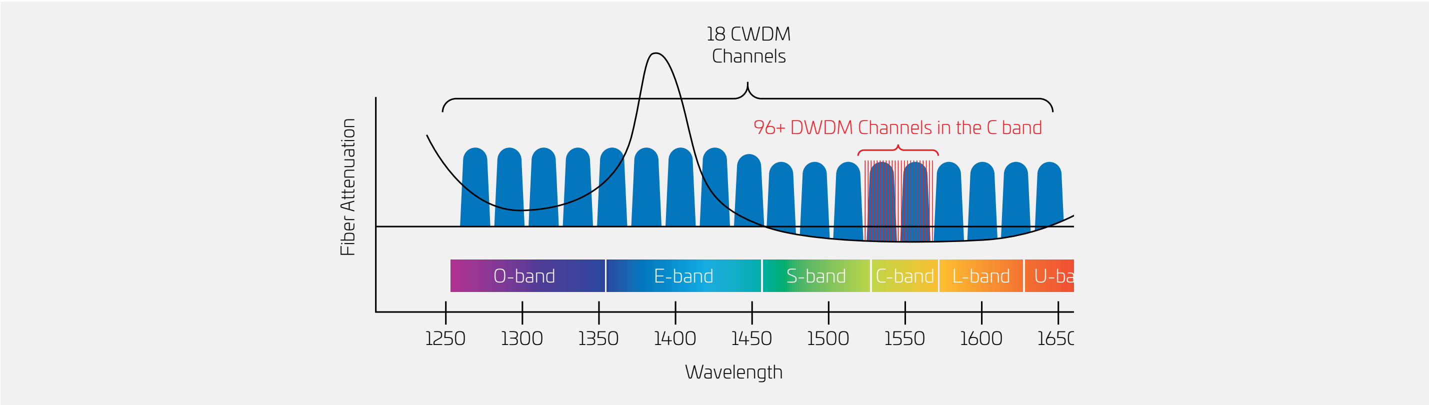

Figure 5: CWDM/DWDM band and channel configurations

QPSK coherent coding vs PAM4

Coherent systems are based around phase shift keying (PSK). This phase modulation technique allows 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), sometimes called 4 quadrature amplitude modulation (4QAM), 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. Figure 6 shows the advancement of modulation formats to support higher data rates.

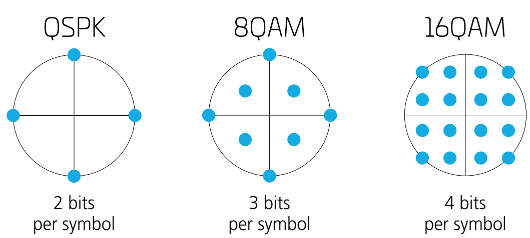

Transmit bit rate = [symbol rate] x [bits per symbol] x [polarization (x2)]

Figure 6: High-order modulation—Constellation™ diagrams

Pulse amplitude modulation (PAM4) is a four-level modulation scheme designed for short-haul fiber links. As shown in Figure 7, PAM4 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.

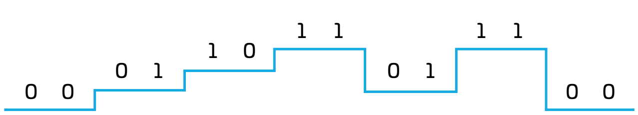

Figure 7: PAM4 signaling technology

Unlike coherent optics, PAM4 is highly susceptible to fiber impairments. This limits range to a maximum of 30 kilometers, whereas coherent optical fiber communication systems can potentially span thousands of kilometers on amplified links. Optical edge gateways convert between QPSK and PAM4 modulation formats.

New designs drive higher fiber-count DCIs

Expanding east-west data flows and moving to a low-latency leaf-and-spine switching architecture have created a tsunami of data inside and among data centers. 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 global cloud and multi-tenant data center (MTDC) providers.

Once the cable has been terminated inside the data center, it needs to 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.

Patch cord hygiene is of the utmost importance to ensure this zone is manageable and can support future growth. As noted earlier, the role of the DCI network across the campus is increasing; soon, they will need to support a throughput of 100T or more. This will require thousands of fibers, all converging at the MMR or MDA, which is why patch cord management is mandatory. Ensuring the ODF has both fiber patch cord routing and slack management is also vital to ensuring that the MDA and MMR can support all future growth needs. Figure 8 shows the infrastructure components for fiber entry into the data center or MTDC meet-me room and a typical campus-scale DCI.

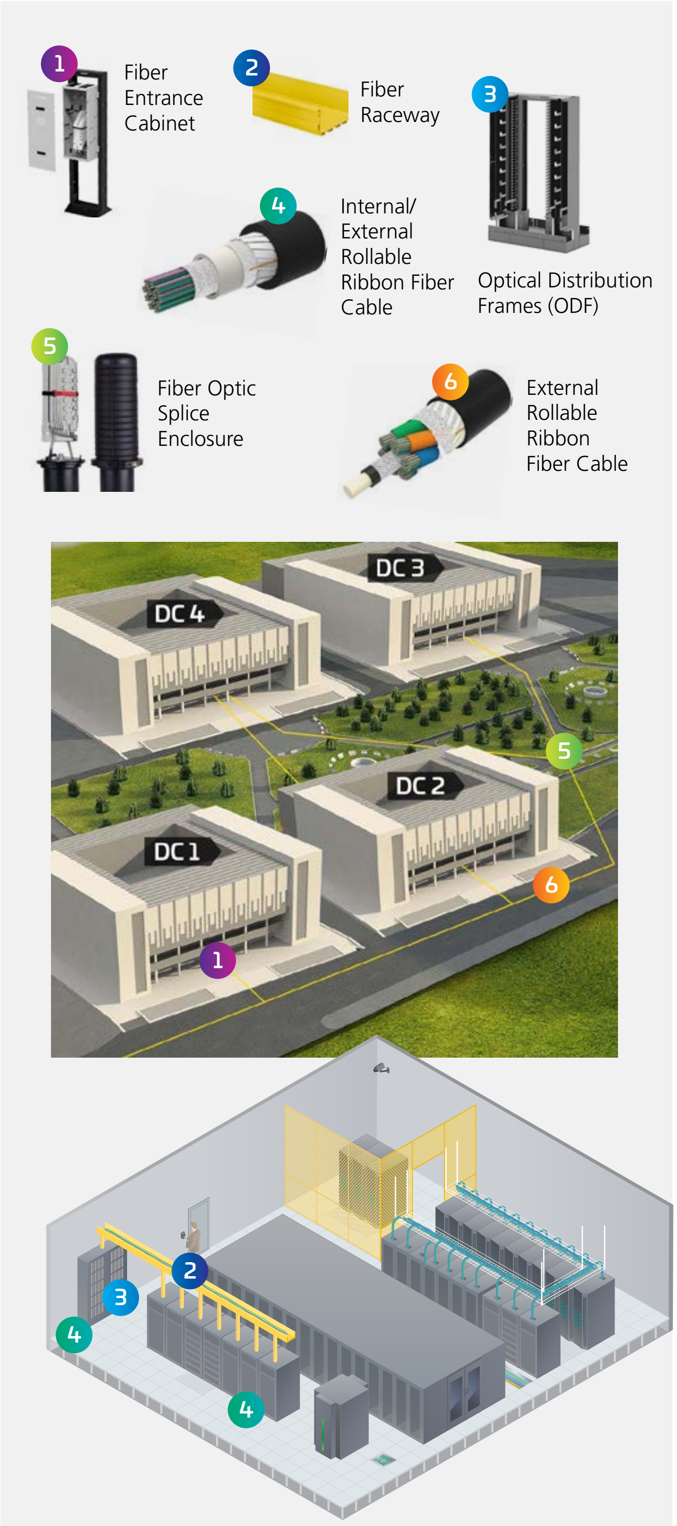

Figure 8: Entry into the data center building or MTDC meet-me room and infrastructure components for a typical campus-scale DCI

Increasing deployment and expansion of DCI networks across the campus and into the cloud will continue unabated. As switch architectures flatten in support of machine learning and artificial intelligence, increasingly more data will be created and shared, mostly by machines that—unlike humans—never take a break or grow tired. As a result, data loads and DCI network requirements will continue to increase. This presents physical challenges for the campus network designer, especially if the network has to 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 by adding more fibers per cable via size-reduced cladding or by introducing a commercially viable alternative glass technology.

[1] Data Center Interconnect Market to Grow at CAGR of 16.03% through 2028, The Brainy Insights, April 2023

CommmScope 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.

CO-118690-EN (03-24)