

An Intro to Data Center Interconnects

Data center interconnects (DCIs) refer to the networking technologies and solutions that enable seamless communication and data exchange between geographically dispersed data centers. As organizations increasingly rely on distributed computing resources and adopt cloud services, the need for efficient and high-speed connections between data centers becomes crucial. DCIs facilitate the transfer of data, applications, and workloads across multiple data center locations, ensuring optimal performance, redundancy, and scalability.

The Impact of Data Center Interconnects

The impact of robust data center interconnects on the operation of data centers is profound. Firstly, DCIs enhance overall reliability and availability by creating a resilient network infrastructure. In the event of a hardware failure or unexpected outage in one data center, DCIs enable seamless failover to another data center, minimizing downtime and ensuring continuous operations. This redundancy is vital for mission-critical applications and services.

Secondly, DCIs contribute to improved performance and reduced latency. By connecting data centers with high-speed, low-latency links, organizations can efficiently distribute workloads and resources, optimizing response times for users and applications. This is particularly important for real-time applications, such as video streaming, online gaming, and financial transactions.

Furthermore, DCIs support efficient data replication and backup strategies. Data can be synchronized across geographically distributed data centers, ensuring data integrity and providing effective disaster recovery solutions. This capability is crucial for compliance with regulatory requirements and safeguarding against data loss.

Types of Data Center Interconnects

As shown in the figure below, we can think about three categories of data center interconnects based on their reach

- Intra-data center interconnects (< 2km)
- Campus data center interconnects (<10km)
- Metro data center interconnects (<100km)

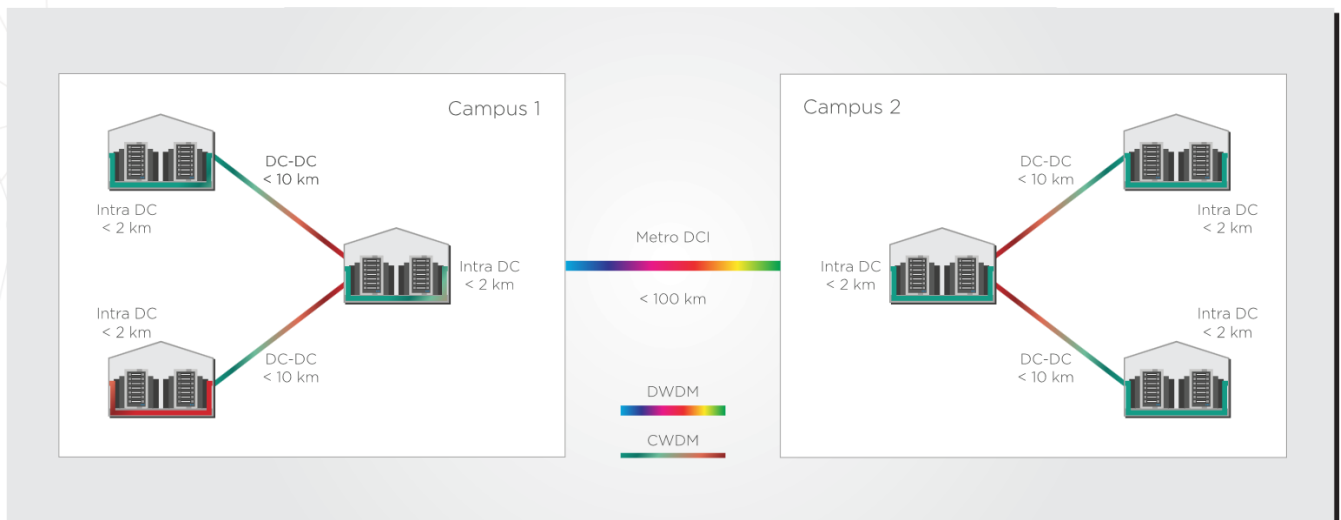


Figure 1: Example of different data center interconnect (DCI) links. Intra-data center interconnects usually have lengths up to 2kms, campus interconnects up to 10km, and metro DCIs up to 100km, although there can be longer DCIs.

Intra-datacenter interconnects operate within a single data center facility. These interconnects are designed for short-distance communication within the same data center building or complex. Intra-DCIs are optimized for high-speed, low-latency connections between servers, storage systems, and networking devices within a single data center. They are crucial for supporting the internal communication and workload distribution required for efficient data center operations.

Campus DCIs connect multiple data centers but are typically limited to a campus area, which may include multiple buildings or locations within a relatively close proximity. This type of interconnect is suitable for organizations with distributed computing resources that need to be interconnected for redundancy, load balancing, and seamless failover within a campus environment.

Metro DCIs connect data centers that are located in different metropolitan areas or cities. They cover longer distances compared to intra-datacenter and campus interconnects, typically spanning tens of kilometers to a few hundred kilometers.

Metro DCIs are essential for creating a network of interconnected data centers across a metropolitan region. They facilitate data replication, disaster recovery, and business continuity by enabling seamless communication and resource sharing between data centers that may be geographically dispersed but still within a reasonable proximity.

The Rise of Edge Data Centers

Various trends are driving the rise of the edge cloud:

- **5G technology and the Internet of Things (IoT):** These mobile networks and sensor networks need low-cost computing resources closer to the user to reduce latency and better manage the higher density of connections and data.

- **Content delivery networks (CDNs):** The popularity of CDN services continues to grow, and most web traffic today is served through CDNs, especially for major sites like Facebook, Netflix, and Amazon. By using content delivery servers that are more geographically distributed and closer to the edge and the end user, websites can reduce latency, load times, and bandwidth costs as well as increasing content availability and redundancy.
- **Software-defined networks (SDN) and Network function virtualization (NFV).** The increased use of SDNs and NFV requires more cloud software processing.
- **Augment and virtual reality applications (AR/VR):** Edge data centers can reduce the streaming latency and improve the performance of AR/VR applications.

Several of these applications require lower latencies than before, and centralized cloud computing cannot deliver those data packets quickly enough. A data center on a town or suburb aggregation point could halve the latency compared to a centralized hyperscale data center. Enterprises with their own data center on-premises can reduce latencies by 12 to 30 times compared to hyperscale data centers.

Cisco estimates that 85 zettabytes of useful raw data were created in 2021, but only 21 zettabytes were stored and processed in data centers. Edge data centers can help close this gap. For example, industries or cities can use edge data centers to aggregate all the data from their sensors. Instead of sending all this raw sensor data to the core cloud, the edge cloud can process it locally and turn it into a handful of performance indicators. The edge cloud can then relay these indicators to the core, which requires a much lower bandwidth than sending the raw data.

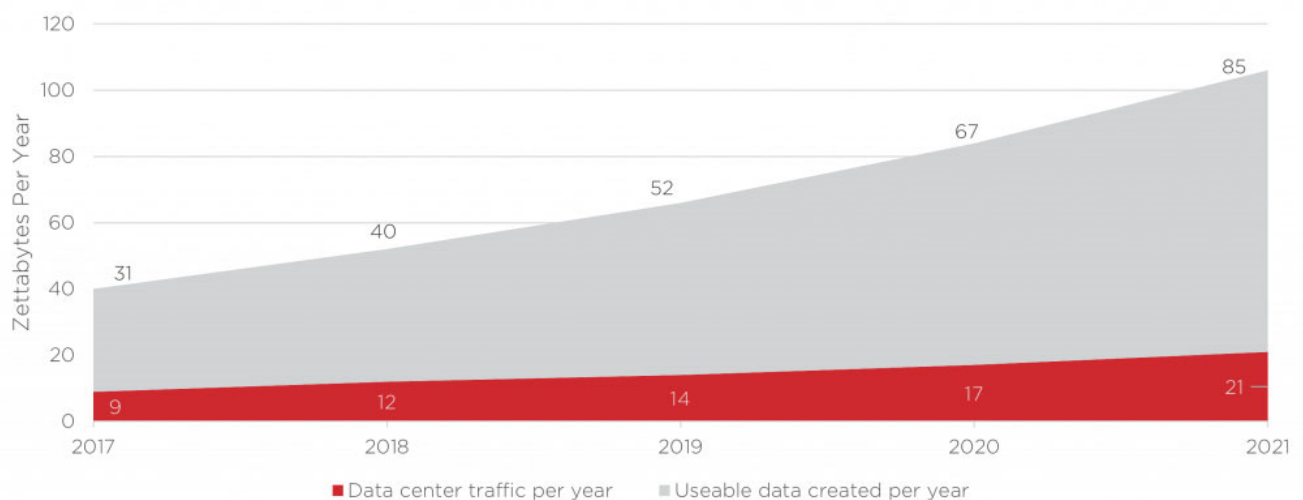


Figure 2. Global data center traffic vs. useable data created per year. Graphic sourced from Cisco Global Cloud Index.

Takeaways

In conclusion, Data Center Interconnects (DCIs) play a pivotal role in fostering reliable, available, and high-performing data center operations, ensuring seamless communication and workload distribution. The categorization of DCIs into intra-data center, campus, and metro interconnects reflects their adaptability to various communication needs. The emergence of edge data centers, driven by 5G, IoT, CDNs, SDNs, NFV, and AR/VR applications, addresses the demand for low-latency computing resources.

Positioned strategically at aggregation points, edge data centers efficiently process and relay data, contributing to bandwidth optimization and closing the gap between raw data generation and traditional data center capacities. As organizations navigate a data-intensive landscape, the evolution of DCIs and the rise of edge data centers underscore their critical role in ensuring the seamless, efficient functioning of distributed computing ecosystems.

[More Articles from EFFECT Photonics](#)