

New challenges and opportunities
await MTDCs at the network edge



Introduction

Whether the setting is the data center in general, or multi-tenant data centers (MTDCs) in particular, most any conversation regarding the future typically begins with network flexibility and speed of service—the overarching themes that define success. Within those themes is a range of more qualitative questions: Where is the edge? What are the storage and compute requirements in that area? What are the customer and application requirements?

As more activity is pushed to “the edge” (to be defined in a moment), the entire network ecosystem is having to adapt. We know, with some certainty, the unique strengths of edge-based devices and core data centers. Much less certain is the role MTDCs will play in this high-reliability, low-latency environment.

In this paper, we examine what CommScope believes to be the MTDC’s “Goldilocks zone,” the diverse requirements that must be addressed, and how MTDC network managers can prepare to take advantage of the opportunity.

Machine-to-machine communication is driving the change

Growth in data center traffic is being generated from the outside in. The source for much of the increase is the number of edge-based devices that are generating exponentially more machine-to-machine (M2M) data. Some of the data is being processed at the edge, but the volume coming into the data center represents a big increase—forcing data center managers to reconsider how they think about their network design and physical infrastructure. Consider the ubiquitous example of the self-driving car.

With regards to network requirements, the impact of autonomous vehicles can be viewed on two levels. On the most basic level, the vehicle must develop a static understanding of the world around

it. Where am I? What’s around me? What laws must I obey in this area? Based on industry estimates, that kind of static awareness consumes about four terabytes of data per day.

On a more advanced level, known as “vehicle under task,” the car must also be able to predict, adjust and react to its dynamic environment. This means sensing people and objects around the vehicle, detecting their movements, calculating speeds and trajectories, and navigating various traffic patterns. This second stage, where the car is having to negotiate and react, consumes about 40 TB a day, not including the video captured by multiple cameras in case visual documentation is needed.

Some of the data will be processed onboard the car and some at a data center further back in the network. As for the video, the




Figure 1: Growth of M2M traffic in the data center

Source: CISCO Annual Internet Report, 2018-2023

Automotive Edge Computing Consortium (AECC) estimates that about 30 percent will be uploaded to help refine models and train algorithms, and about 2 percent will be uploaded to retain for audit trails for accidents or traffic incidents. The bottom line: Autonomous cars could require data transfer offload ranging from 383 GB an hour to 5.17 TB an hour. When multiplied by the compute needs of hundreds or thousands of autonomous vehicles sharing the road, the network can quickly become overwhelmed.

Of course, handling the data is just one challenge. Satisfying the low and ultra-low latency requirements of a growing number of edge-based applications and devices may be an even tougher challenge for networks. This brings up the first fundamental question for understanding the role of the MTDC. Where, exactly, is the edge? (Hint, it's a trick question.)

MTDC at the edge

In defining the edge, there are actually two perspectives.

The first is the customer edge. It will be located on the customer's premises to support ultra-low latency measured in single-digit milliseconds. These include manufacturing applications like fully-automated robotics, enabled by 5G or wireless transport.

The second definition of the edge is the network edge. This will be located further back toward the network core, but close enough to support latency below 20 milliseconds. This will enable applications such as cloud-assisted driving, 360-degree

video and high-res gaming. It is here at the network edge that multi-tenant data centers will find a sweet spot. One reason is that so many MTDCs have adapted to thrive in densely populated areas. These adaptations provide an excellent fit for the high-density and good-not-great latency demands that must be addressed at the network edge.

To capitalize on this opportunity, MTDCs must be ready to accommodate a variety of customer configurations. Many customers will provide their own cabinets and specific requirements (ranging from a single cabinet to two cabinets up to a full suite) depending on performance and other factors. Whatever processes they have running in their server cabinets will dictate the size and number of the cabinets you'll be expected to roll into your suites to support that. In other cases, customers will arrive with nothing but a logical IT diagram indicating servers, switches and connectivity requirements. They look to the data center for the physical layer solutions.

For their part, the MTDC is looking to satisfy the variety of customer needs while continuing to drive deployment times lower and accelerate the return on their efforts. They're finding some unlikely but willing partners like AWS whose Wavelength solution features fully-equipped and connected cabinet rows. All the needed computing and switching capacity is built in and can be rolled into the customer's cage. As the customer requires more performance, more pre-configured racks can be rolled in and hooked up to the MTDC's infrastructure. Alternatively, MTDCs may be asked to provide cabinets tailored to support custom compute characteristics.




Figure 2: Customer edge vs the network edge

Source: AT&T

Supporting flexible connectivity

Located on the network edge, the MTDC will rely on a highly agile and redundant physical layer infrastructure to meet the fiber density and low latency demands of customers. Currently, the base topology, shown in Figure 3, is a leaf-and-spine network connected by structured cabling. The structured cabling is critical, as it enables the network to continually expand and migrate to higher speeds easily. Using high fiber-count trunk cables, like 24-fiber MPO, the backbone cabling can remain fixed. As optical networking technologies change from duplex to parallel optics, and back again, you simply have to change the module and optical fiber presentation entering or exiting the spine or the leaf cabinet. This eliminates the need to rip and replace trunk cabling. This approach to physical layer design using higher fiber-count trunks can also eliminate the need to navigate overcrowded raceways that may be already servicing multiple customers.

Another advantage of deploying a structured cabling architecture in a spine-and-leaf network is its support for equal-cost multi-path (ECMP) routing. ECMP is a routing strategy that enables packet forwarding to a single destination using multiple and equal performance “best paths” as determined by routing metric calculations. As a network load-balancing protocol, ECMP requires an any-to-any leaf-and-spine switch configuration. Without a consistent and structured approach to physical layer cabling, it is difficult to know the exact characteristics of each path in the network fabric—significantly impairing the ability to expand that fabric with new server cabinets as additional networking performance is demanded by customers.

SYSTIMAX® structured cabling

SYSTIMAX® is a comprehensive platform of advanced copper and fiber connectivity, with automated infrastructure management and high-efficiency interfaces that help you put more connections into increasingly crowded data center spaces. And, because SYSTIMAX is a CommScope solution, it's backed by a reliable, responsive global network of manufacturing, logistics and expertise. [LEARN MORE >](#)

As shown in Figure 3, the spine switches are at the edge of the network. Optical distribution frames connect the data center to the outside world and connect the myriad optical fibers within the fabric. Below the spine switches are the leaf switches—or top-of-rack switches in some configurations—that connect the server cabinets to the fabric. The resulting structured cabling network fabric is supported by MPO or LC connector-based technology for 100 GbE+ bandwidth. This configuration can provide massive processing power and is easily reconfigured to accommodate the expanding and fluid nature of an edge-based environment.

Network Edge Compute—Logical




Figure 3: Typical network-edge topology

Accommodating the need for faster network speeds

As activity at the edge ramps up, link speeds are accelerating. While a few MTDC customers still prefer 40 GbE, most discussions with network and IT managers are focused on enabling 100 GbE connectivity, with some customers talking about 200 GbE and 400 GbE for their backbone speeds.


In April 2020, the Ethernet Technology Consortium released the 800GBASE-R specification for 800 GbE, which implements a new media access control (MAC) and Physical Coding Sublayer (PCS). The takeaway for the MTDC network manager is the need to incorporate enough capacity and bandwidth now to support these higher data rates in the future.

The speed of the network fabric in leaf-and-spine deployments is dictated by the bandwidths required by each of the attached server cabinets. Once the servers are in place, the customer may need to increase server speed. The MTDC must future-proof the network to allow the customer to upgrade their speed without having to touch the physical infrastructure around the cabinets.

One way to do this is through increasing fiber densities in the trunk cables, and supporting them at the switch and server layers with high- and ultra high-density fiber-optic patch panels. Deploying a structured cabling architecture, enables network managers to easily tap into more bandwidth as needed without having to touch the trunks. This is a good strategy for aggregating many cabinets in one area.

In addition to adding fiber, MTDCs can also choose to employ equipment to support wavelength division multiplexing as well as different forms of encoding to increase the QAM rate—thereby increasing the data carrying capacity of the physical infrastructure.

Fiber-optic panels



CommScope's enhanced high-density (EHD), ultra high-density (UD) and high-density (HD) fiber panels are uniquely designed to provide open access to individual fibers. At the same time, an innovative fiber containment and routing design protects every connection and keeps the fiber infrastructure accessible and manageable. The result? Easier and faster moves, adds and changes; accelerated mean time to repair; simplified installation and lower costs. [Learn more >](#)

In reality, data centers will need to use all three strategies to support the data demands at the network edge. Currently, this means providing speeds up to 25 GbE at the server port and 100 GbE from the leaf switch to spine switch, while tracking the progress of future network speeds of 200 GbE or 400 GbE.




Figure 4: 400 GbE ports on the rise

Simple Migration Using A 24F MPO Trunk Cable




Figure 5: Simple migration using 24-fiber MPO

24-fiber MPO connectivity

As MTDCs increase their bandwidth, network managers must decide how to deliver it to the compute and storage resources in the most efficient manner. MPO connectorization has become the default solution. Currently there are three types of MPO connectors: 8-fiber, 12-fiber and 24-fiber. All three can support virtually any multimode technologies as well as many singlemode technologies.

While there are good use cases for both the 8- and 12-fiber MPOs, the 24-fiber modules offer the lowest-cost duplex support for multimode applications and the highest panel density. More importantly, the 24-fiber MPO provides greater flexibility as MTDCs struggle to adapt to changing workloads. Figure 4, for example, illustrates the various configurations the technology supports during a simple migration. To maintain duplex connectivity from the spine to the leaf switch, the 24-fiber technology allows you to aggregate the 12 LC connections over one simple trunk cable. Three 40/100 GbE channels can be aggregated onto a single 24-fiber MPO trunk. This simplifies the architecture by reducing the overall number of trunk cables used in a fabric and eliminates the need to buy, stock, prep and install multiple cables. It also reduces the load in the cable raceways.

SDN and white-boxing

Another trend affecting the physical layer is the move to software-defined networking (SDN), also known as “white boxing.” Service providers, hyperscalers, and enterprises are replacing their branded, dedicated hardware with generic servers and switches (white boxes). As a result, the application-specific intelligence moves from the equipment to the software, via SDN. White boxing creates added efficiencies and flexibility, as well as some new challenges. Network capacity must support the dynamic nature of SDN. The cabling channels in that area must be able to withstand the next generation of active equipment, whose functionality is no longer fixed at the point of installation but is now dynamic.

Standardized Cabling Architecture for MTDC's




Figure 6: Standardized MTDC cabling architecture with dedicated ODFs

Customer diversity and the expanding role of ODFs

The need for physical layer flexibility is also being driven by the growing diversity of customers the MTDC must now support—and the list goes beyond retail and wholesale customers. With available space and geographic proximity to the larger markets, MTDCs are well positioned to service network-edge applications like social media and search engine activity. This creates an interesting opportunity to partner with hyperscalers looking for a cost-effective presence at the edge.

To accommodate these larger customers, MTDCs are interconnecting their cages using multiple high-fiber core count trunks. In a typical MTDC configuration, the service provider network terminates at a dedicated optical distribution frame (ODF) in the meet-me room, ready for presentation out to the other suites. One strategy many MTDCs are employing is to use additional ODFs as intermediate distribution points dedicated to serving the customer cages on each floor. This enables MTDCs to independently add capacity for new or expanding service providers as well as customers. Employing an additional local distribution point, configured from 19-inch fiber panels very close to the customer cages, often enables the final connections to be made quickly using shorter runs or even patch cords.

In a tiered-ODF design, a single customer cage can cross-connect to multiple service providers via a simple patch. Because many

MTDCs bill for the activity at the cross-connect, the faster those connections can be made, the higher the potential revenue. Additionally, the customer-to-service provider connection often generates monthly recurring revenue. Having the ability to easily hook up customers to more service providers makes this a valuable strategy.

The CommScope ODF portfolio

CommScope's ODF portfolio includes solutions for every area and application within the data center. The FACT ODF handles 5,376 LC connections and provides full front accessibility and tool-less installation in a modular, lightweight frame. The NG4access® ODF supports up to 3,456 connections and offers on-frame splicing through a full front and rear access platform. There are also several specialized ODF solutions available.

Future trends

400 GbE and 800 GbE will see the deployment of new optical modules, like the QSFP-DD and OSFP modules. These are being developed as second-generation 400GE, for high port-density data center switches.

The QSFP-DD is a new module and cage/connector system similar to current QSFP modules, but with an additional row of contacts that provide an eight-lane electrical interface. Each lane will support up to 25 Gbps NRZ modulation or 50 Gbps PAM4 modulation and provide up to 14.4 Tbps aggregate bandwidth in a single switch slot.




Figure 7: New connectivity options for 400 GbE

Like the QSFP-DD, the OSFP features 8x lanes of 50G PAM4 (electrical) and 8x lasers of 50G PAM4 or 4x lasers of 100G PAM4 (optical). One of the key differences between the two modules is power capabilities. The OSFP consumes more power than the QSFP-DD—15 watts compared to 12 watts. The QSFP-DD is also backward compatible with QSFP+ (40G), QSFP28 (100G) and QSFP56 (200G) while the OSFP requires a QSFP-to-OSFP converter module.

New flavors of connectors are going to come along with 400 GbE. In addition to the traditional 12-fiber MPO and two-row MPO 12, a 16-fiber MPO is also getting ready to hit the market. Two very small form factor connectors, the SN and MDC connector, will also soon be ready for prime time, as well.

The best use of these very high-density connectors will be the switch-to-switch or switch-to-server applications enabling the highest density of fibers.

Remaining adaptable will be the key to success

In the near future, as edge computing comes into focus, its effects on how MTDCs need to adapt will crystalize as well. This much we do know: The edge is redefining the relationships between MTDCs and their customers and attracting some non-traditional customers and competitors. MTDCs must be prepared to evolve to meet the new customer requirements while protecting their turf. That means adapting their infrastructure and compute resources to support applications and devices closer to the network edge. Exactly where those resources need to be located depends on the latency demands. CommScope predicts that the sweet spot will be those areas requiring 10–20 milli-second performance range.

With the rate of growth in bandwidth and fiber requirements constantly increasing, MTDC operators must build as much flexibility into their physical-layer infrastructure as possible. This requires a structured cabling approach, very high-density fiber panels and ODFs, and flexible MPO configurations that will be able to support new customer types, faster speeds and additional service provider circuits. MTDC operators would be well advised to begin preparing now for the changes that are on the way.

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](https://www.commscope.com)

COMMScope®

[commscope.com](https://www.commscope.com)

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

© 2020 CommScope, Inc. All rights reserved.

Unless otherwise noted, all trademarks identified by ® or ™ are registered trademarks or trademarks, respectively, of CommScope, Inc. This document is for planning purposes only and is not intended to modify or supplement any specifications or warranties relating to CommScope products or services. CommScope is committed to the highest standards of business integrity and environmental sustainability, with a number of CommScope's facilities across the globe certified in accordance with international standards, including ISO 9001, TL 9000, and ISO 14001. Further information regarding CommScope's commitment can be found at <https://www.commscope.com/corporate-responsibility-and-sustainability>.

WP-115083-EN (11/20)