COMMSCOPE

WHAT'S NEXT FOR THE DATA CENTER:

2021 Trends to Watch

Contents

Looking to the year ahead: What's impacting the data center?

There's no such thing as 'business as usual' in the data center, and 2020 was certainly no exception to this rule. With the volume of data pouring into the data center continuing to climb, driven by even greater connectivity demand - network planners are rethinking how they can stay a step ahead of these changes!

Looking back to 2014, when the 25G Ethernet Consortium proposed single-lane 25-Gbit/s Ethernet and dual-lane 50-Gbit/s Ethernet, it created a big fork in the industry's roadmap, offering a lower cost-per-bit and an easy transition to 50G, 100G and beyond.

In 2020, 100G hit the market en masse, driving higher and higher fiber counts - and larger hyperscale and cloud-based data centers confronted their inevitable leap to 400G. With switches and servers on schedule to require 400G and 800G connections, the physical layer must also contribute higher performance to continuously optimize network capacity.

The ability to evolve the physical layer infrastructure in the data center is ultimately key to keeping pace with demand for the low latency, high-bandwidth, and reliable connectivity that subscribers demand. Take a look at these top trends to watch as data center managers plan for 800G and the data mushroom effect that 5G will bring!

About the authors

James Young

James is the Director of CommScope's Enterprise Data Center division, overseeing strategy and providing leadership to product and field teams globally. Formerly, James has been involved in a variety of roles including sales, marketing and operations for communication solutions working with Tyco Electronics/AMP, Anixter, Canadian Pacific and TTS in Canada. James has gained extensive experience in the sale of OEM products, network solutions and value-added services through direct and indirect channel sales environments. James graduated with a Bachelor of Science from the University of Western Ontario. He is a registered Communication Distribution Designer (RCDD) and certified Data Center Design Professional (CDCP).

Jason Bautista

As Solution Architect for Hyperscale and Multi-Tenant Data Centers, Jason is responsible for data center market development for CommScope. He monitors trends in the data center market to help drive product roadmap strategies, solutions and programs for Hyperscale and Multi-Tenant data center customers. Jason has more than 19 years experience in the networking industry having held various customer facing positions in product development, marketing and support for a diverse range of networks and customers across the globe.

Ken Hall

Kenneth Hall is Data Center Architect for North America at CommScope, responsible for technology and thought leadership, as well as high speed migration planning for enterprise and related data centers. Ken has been with CommScope and other acquired companies for 32 years in global program and project management, technical sales, marketing and industry standards.

Matt Baldassano

Matt Baldassano supports the Northeast Region of CommScope as a Technical Director– Enterprise Solutions specializing in Data Center connectivity. He has served roles as Business Development Manager and Technical Marketing Engineer with CommScope's Data Center Business Unit. His experience also includes responsibilities as Account Engineer in New York City and Dallas TX for EMC2 corporation serving both wired data centers and in-building wireless systems and has written topics on wireless security. Matt holds a BS in Computer Science from St John's University and a MS in Technology from the University of Advancing Technology.

1 /

Adapting to higher fiber counts in the data center

The volume of digital traffic pouring into the data center continues to climb; meanwhile, a new generation of applications driven by advancements like 5G, AI and machine-to-machine communications is driving latency requirements into the single millisecond range. These and other trends are converging in the data center's infrastructure, forcing network managers to rethink how they can stay a step ahead of the changes.

Traditionally, networks have had four main levers with which to meet increasing demands for lower latency and increased traffic.

- Reduce signal loss in the link
- Shorten the link distance
- Accelerate the signal speed
- Increase the size of the pipe

While data centers are using all four approaches at some level, the focus—especially at the hyperscale level—is now on increasing the amount of fiber. Historically, the core network cabling contained 24, 72, 144 or 288 fibers. At these levels, data centers could manageably run discrete fibers between the backbone and switches or servers, then use cable assemblies to break them out for efficient installation. Today, fiber cables are deployed with as many as 20 times more fiber strands — in the range of 1,728-, 3,456- or 6,912 fibers per cable.

Higher fiber count combined with compact cable construction is especially useful when interconnecting data centers (DCI). DCI trunk cabling with 3,000+ fibers is common for connecting two hyperscale facilities, and operators are planning to double that design capacity in the near future. Inside the data center, problem areas include backbone trunk cables that run between high-end core switches or from meet-me rooms to cabinet-row spine switches.

Whether the data center configuration calls for point-to-point or switch-to-switch connections, the increasing fiber counts create major challenges for data centers in terms of delivering the higher bandwidth and capacity where it is needed.

The massive amount of fiber creates two big challenges for the data center. The first is, how do you deploy it in the fastest, most efficient way: how do you put it on the spool, how do you take it off of the spool, how do you run it between points and through pathways? Once it's installed, the second challenge becomes, how do you break it out and manage it at the switches and server racks?

Rollable ribbon fiber cabling

The progression of fiber and optical design has been a continual response to the need for faster, bigger data pipes. As those needs intensify, the ways in which fiber is designed and packaged within the cable have evolved, allowing data centers to increase the number of fibers without necessarily increasing the cabling footprint. Rollable ribbon fiber cabling is one of the more recent links in this chain of innovation.

Rollable ribbon fiber is bonded at intermittent points.
Source: ISE Magazine

Rollable ribbon fiber cable is based, in part, on the earlier development of the central tube ribbon cable. Introduced in the mid-1990s, primarily for OSP networks, the central tube ribbon cable featured ribbon stacks of up to 864 fibers within a single, central buffer tube. The fibers are grouped and continuously bonded down the length of the cable which increases its rigidity. While this has little affect when deploying the cable in an OSP application in a data center a rigid cable is undesirable because of the limited routing restrictions these cables require.

In the rollable ribbon fiber cable, the fibers are attached intermittently to form a loose web. This configuration makes the ribbon more flexible, allowing manufacturers to load as many as 3,456 fibers into one two-inch duct, twice the density of conventionally packed fibers. This construction reduces the bend radius making these cables easier to work with inside the tighter confines of the data center.

Inside the cable, the intermittently bonded fibers take on the physical characteristics of loose fibers which easily flex and bend making it easier to manage in tight spaces. In addition, rollable ribbon fiber cabling uses a completely gel-free design which helps reduce the time required to prepare for splicing, therefore reducing labor costs. The intermittent bonding still maintains the fiber alignment required for typical mass fusion ribbon splicing.

Reducing cable diameters

For decades, nearly all telecom optical fiber has had a nominal coating diameter of 250-microns. With growing demand for smaller cables, that has started to change. Many cable designs have reached practical limits for diameter reduction with standard fiber. But a smaller fiber allows additional reductions. Fibers with 200-micron coatings are now being used in rollable ribbon fiber and micro-duct cable.

It is important to emphasize that the buffer coating is the only part of the fiber that has been altered. 200-micron fibers retain the 125-micron core/cladding diameter of conventional fibers for compatibility in splicing operations. Once the buffer coating has been stripped, the splice procedure for 200-micron fiber is the same as for its 250-micron counterpart.

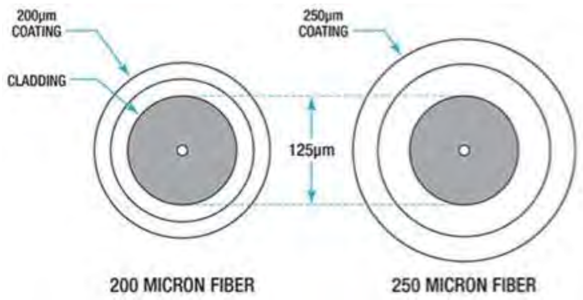

For optical performance and splice compatibility, 200-micron fiber features the same 125-micron core/cladding as the 250-micron alternative.
Source: ISE Magazine

New chipsets are further complicating the challenge

All servers within a row are provisioned to support a given connection speed. But in today's hyper-converged fabric networks it is extremely rare that all servers in a row will need to run at their max line rate at the same time. The difference between the server's upstream bandwidth required and the downstream capacity that's been provisioned is known as the oversubscription, or contention ratio. In some areas of the network, such as the inter-switch link (ISL), the oversubscription ratio can be as high as 7:1 or 10:1. A higher ratio is chosen to reduce switch costs, but the chance of network congestion increases with these designs.

Oversubscription becomes more important when building large server networks. As switch to switch bandwidth capacity increases, switch connections decrease. This requires multiple layers of leaf-spine networks to be combined to reach the number of server connections required. Each switch layer adds cost, power and latency however. Switching technology has been focused on this issue driving a rapid evolution in merchant silicon switching ASICs. On December 9, 2019, Broadcom Inc. began shipping the latest StrataXGS Tomahawk 4 switch, enabling 25.6 Terabits/sec of Ethernet switching capacity in a single ASIC. This comes less than two years after the introduction of the Tomahawk 3 which clocked in at 12.8Tbps per device.

These ASICs have not only increased lane speed, they have also increased the number of ports they contain. Data centers can keep the oversubscription ratio in check. A switch built with a single TH3 ASICs supports 32 400G ports. Each port can be broken down to eight 50GE ports for server attachment. Ports can be grouped to form 100G, 200G or 400G connections. Each switch port may migrate between 1-pair, 2-pair, 4-pairs or 8-pairs of fibers within the same amount of QSFP footprint.

While this seems complicated it is very useful to help eliminate oversubscription. These new switches can now connect up to 192 servers while still maintaining 3:1 contention ratios and eight 400G ports for leaf-spine connectivity! This switch can now replace six previous-generation switches.

The new TH4 switches will have 32 800Gb ports. ASIC lane speeds have increased to 100G. New electrical and optical specifications are being developed to support 100G lanes. The new 100G ecosystem will provide an optimized infrastructure which is more suited to the demands of new workloads like machine learning or AI.

The evolving role of the cable provider

In this dynamic and more complex environment, the role of the cabling supplier is taking on new importance. While fiber cabling may once have been seen as more of a commodity product instead of an engineered solution, that is no longer the case. With so much to know and so much at stake, suppliers have transitioned to technology partners, as important to the data center's success as the system integrators or designers.

Data center owners and operators are increasingly relying on their cabling partners for their expertise in fiber termination, transceiver performance, splicing and testing equipment, and more. As a result, this increased role requires the cabling partner to develop closer working relationships with those involved in the infrastructure ecosystem as well as the standards bodies.

As the standards surrounding variables such as increased lane speeds multiply and accelerate, the cabling partner is playing a bigger role in enabling the data center's technology roadmap. Currently, the standards regarding 100GE/400GE and evolving 800Gbs involve a dizzying array of alternatives. Within each option, there are multiple approaches, including duplex, parallel and wavelength division multiplexing – each with a particular optimized application in mind. Cabling infrastructure design should enable all of these alternatives.

All comes down to balance

As fiber counts grow, the amount of available space in the data center will continue to shrink. Look for other components, namely servers and cabinets to deliver more in a smaller footprint as well.

Space won't be the only variable to be maximized. Combining new fiber configurations like rollable ribbon fiber cables with reduced cable sizes and advanced modulation techniques, network managers and their cabling partners have lots of tools at their disposal. They will need them all.

If the rate of technology acceleration is any indication of what lies ahead, data centers, especially at the hyperscale cloud level, better strap in. As bandwidth demands and service offerings increase, and latency becomes more critical to the end-user, more fiber will be pushed deeper into the network.

The hyperscale and cloud-based facilities are under increasing pressure to deliver ultra-reliable connectivity for a growing number of users, devices and applications. The ability to deploy and manage ever higher fiber counts is intrinsic to meeting those needs.

The goal is to achieve balance by delivering the right number of fibers to the right equipment, while enabling good maintenance and manageability and supporting future growth. So, set your course and have a solid navigator like CommScope on your team.

2 /

400G in the data center: Options for optical transceivers

The first measure of an organization's success is its ability to adapt to changes in its environment. Call it survivability. If you can't make the leap to the new status quo, your customers will leave you behind.

For cloud-scale data centers, their ability to adapt and survive is tested every year as increasing demands for bandwidth, capacity and lower latency fuel migration to faster network speeds. During the past several years, we've seen link speeds throughout the data center increase from 25G/100G to 100G/ 400G. Every leap to a higher speed is followed by a brief plateau before data center managers need to prepare for the next jump.

Currently, data centers are looking to make the jump to 400G. A key consideration is which optical technology is best. Here, we break down some of the considerations and options.

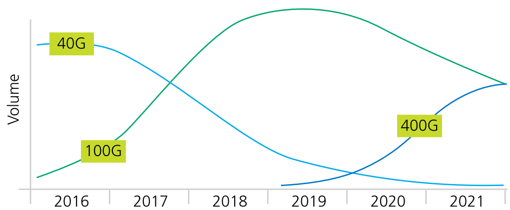

Note: 400G port numbers include both 8x50G and 4x100G implementations
Source: NextPlatform 2018

400GE optical transceivers

The optical market for 400G is being driven by cost and performance as OEMs try to dial into the data centers' sweet spot.

In 2017, CFP8 became the first-generation 400GE module form factor to be used in core routers and DWDM transport client interfaces. The module dimensions are slightly smaller than CFP2, while the optics support either CDAUI-16 (16x25G NRZ) or CDAUI-8 (8x50G PAM4) electrical I/O. Lately, the focus has shifted to the second-generation 400GE form factor modules: QSFP-DD and OSFP.

Developed for use with high port-density data center switches, these thumb-sized modules enable 12.8 Tbps in 1RU via 32 x 400GE ports and support CDAUI-8 (8x50G PAM4) electrical I/O only.

While the CFP8, QSFP-DD and OSFP are all hot-pluggable, that's not the case with all 400GE transceiver modules. Some are mounted directly on the host printed circuit board. With very short PCB traces, these embedded transceivers enable low power dissipation and high port density.

Despite the higher bandwidth density and higher rates per channel for embedded optics, the Ethernet industry continues to favor pluggable optics for 400GE; they are easier to maintain and offer pay-as-you-grow cost efficiency.

Start with the end in mind

For industry veterans, the jump to 400G is yet another waystation along the data center's evolutionary path. There is already an MSA group working on 800G using 8 x 100G transceivers. CommScope—a member of the 800G MSA group—is working with other IEEE members seeking solutions that would support 100G-per-wavelength server connections using multimode fiber. These developments are targeted to enter the market in 2021, followed by 1.6T schemes, perhaps in 2024.

While the details involved with migrating to higher and higher speeds are daunting, it helps to put the process in perspective. As data center services evolve, storage and server speeds must also increase. Being able to support those higher speeds requires the right transmission media.

In choosing the optical modules that best serve the needs of your network, start with the end in mind. The more accurately you anticipate the services needed and the topology required to deliver those services, the better the network will support new and future applications.

3 /

400G in the data center Densification and campus architecture

400G creates new demands for the cabling plant

Higher bandwidth and capacity demands are driving fiber counts through the roof. Fifteen years ago, most fiber backbones in the data center used no more than 96 strands, including coverage for diverse and redundant routing.

Current fiber counts of 144, 288, and 864 are becoming the norm, while interconnect cables and those used across hyper- and cloud-scale data centers are migrating to 3,456 strands. Several fiber cable manufacturers now offer 6,912-fiber cables, and 7,776 fibers are on the horizon.

New fiber packaging and design increases density

The higher fiber-count cabling takes up valuable space in the raceways, and their larger diameter presents performance challenges regarding limited bend radii. To combat these issues, cable OEMs are moving toward rollable-ribbon construction and 200-micron fiber.

Whereas traditional ribbon fiber bonds 12 strands along the entire length of the cable, rollable ribbon fiber is intermittently bonded —allowing the fiber to be rolled rather than lay flat. On average, this type of design allows 3,456 strands to fit into a two-inch duct compared to a flat design that can accommodate only 1,728 in the same space.

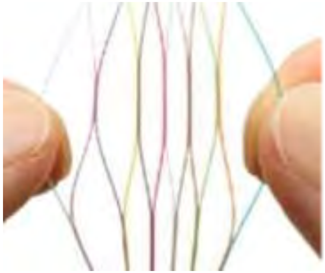

Rollable ribbon fiber
Source: Fluke Networks

The 200-micron fiber retains the standard 125-micron cladding, which is fully backward compatible with current and emerging optics; the difference is that the typical 250-micron coating is reduced to 200 microns. When paired with rollable ribbon fiber, the decreased fiber diameter enables cabling OEMs to keep the cable size the same while doubling the number of fibers compared to a traditional 250-micron flat ribbon cable.

Technologies like rollable ribbon and 200-micron fiber are deployed by hyperscale data centers to support the increased demand for inter-data center connectivity. Within the data center, where leaf-to-server connection distances are much shorter and densities much higher, the primary consideration is the capital and operating cost of optic modules.

For this reason, many data centers are sticking with lower cost vertical-cavity surface-emitting laser (VCSEL) transceivers, which are supported by multimode fiber. Others opt for a hybrid approach—using singlemode in the upper mesh network layers while multimode connects servers to the tier one leaf switches. As more facilities adopt 400GE, network managers will need these options to balance cost and performance as 50G and 100G optic connections to server become the norm.

80 km DCI space: Coherent vs. direct detection

As the trend to regional data center clusters continues, the need for high-capacity, low-cost data center interconnect (DCI) links becomes increasingly urgent. New IEEE standards are emerging to provide a variety of lower-cost options that offer plug-and-play, point-to-point deployments.

Transceivers based on traditional four-level pulse amplitude modulation (PAM4) for direct detection will be available to provide links up to 40 km while being directly compatible with the recent 400G data center switches. Still other developments are targeting similar functionality for traditional DWDM transport links.

As link distances increase beyond 40 km to 80 km and beyond, coherent systems offering enhanced support for long-haul transmission are likely to capture most of the high-speed market.

Coherent optics overcome limitations like chromatic and polarization dispersion, making them an ideal technical choice for longer links. They have traditionally been highly customized (and expensive), requiring custom "modems" as opposed to plug-and-play optic modules.

As technology advances, coherent solutions likely will become smaller and cheaper to deploy. Eventually, the relative cost differences may decrease to the point that shorter links will benefit from this technology.

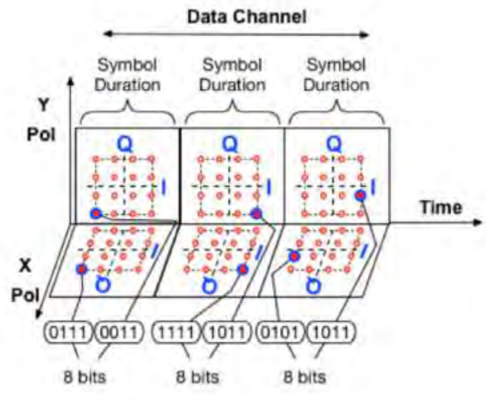

Source: https://www.cablelabs.com/point-to-point-coherent-optics-specifications

Taking a holistic approach to continual high-speed migration

The continual journey to higher speeds in the data center is a step-process; as applications and services evolve, storage and server speeds must also increase. Adopting a patterned approach to handle the repeated periodic upgrades can help reduce the time and cost needed to plan and implement the changes. CommScope recommends a holistic approach in which switches, optics and fiber cabling operate as a single coordinated transmission path.

Ultimately, how all these components work together will dictate the network's ability to reliably and efficiently support new and future applications. Today's challenge is 400G; tomorrow, it will be 800G or 1.6T. The fundamental requirement for high-quality fiber infrastructure remains constant, even as network technologies continue to change.

4 /

Don’t look now — here comes 800G!

100G optics are hitting the market en masse and 400G is expected sometime next year. Nevertheless, data traffic continues to increase and the pressure on data centers is only ramping up.

Balancing the three-legged table

In the data center, capacity is a matter of checks and balances among servers, switches and connectivity. Each pushes the other to be faster and less expensive. For years, switch technology was the primary driver. With the introduction of Broadcom's StrataXGS® Tomahawk® 3, data center managers can now boost switching and routing speeds to 12.8 Tbps and reduce their cost per port by 75 percent. So, the limiting factor now is the CPU, right? Wrong. Earlier this year, NVIDIA introduced its new Ampere chip for servers. It turns out the processors used in gaming are perfect for handling the training and inference-based processing needed for artificial intelligence (AI) and machine learning (ML).

The bottleneck shifts to the network

With switches and servers on schedule to support 400G and 800G, the pressure shifts to the physical layer to keep the network balanced. IEEE 802.3bs, approved in 2017, paved the way for 200G and 400G Ethernet. However, the IEEE has only recently completed its bandwidth assessment regarding 800G and beyond. Given the time required to develop and adopt new standards, we may already be falling behind.

So, cabling and optics OEMs are pressing ahead to keep momentum going as the industry looks to support the ongoing transitions from 400G to 800G, 1.2 Tb and beyond. Here are some of the trends and developments we're seeing.

Switches on the move

For starters, server-row configurations and cabling architectures are evolving. Aggregating switches are moving from the top of the rack (TOR) to the middle of the row (MOR), and connecting to the switch fabric through a structured cabling patch panel. Now, migrating to higher speeds involves simply replacing the server patch cables instead of replacing the longer switch-to-switch links. This design also eliminates the need to install and manage 192 active optical cables (AOCs) between the switch and servers.

Transceiver form factors changing

New designs in pluggable optic modules are giving network designers additional tools, led by 400G-enabling QSFP-DD and OSFP. Both form factors feature 8x lanes, with the optics providing eight 50G PAM4. When deployed in a 32-port configuration, the QSFP-DD and OSFP modules enable 12.8 Tbps in a 1RU device. The OSFP and the QSFP-DD form factor support the current 400G optic modules and next-generation 800G optics modules. Using 800G optics, switches will achieve 25.6 Tbps per 1U.

New 400GBASE standards

There are also more connector options to support 400G short-reach MMF modules. The 400GBASE-SR8 standard allows for a 24-fiber MPO connector (favored for legacy applications) or a single-row 16-fiber MPO connector. The early favorite for cloud scale server connectivity is the single-row MPO16. Another option, 400GBASE-SR4.2, uses a single-row MPO 12 with bidirectional signaling—making it useful for switch-to-switch connections. IEEE802.3 400GbaseSR4.2 is the first IEEE standard to utilize bidirectional signaling on MMF, and it introduces OM5 multimode cabling. OM5 fiber extends the multi-wavelength support for applications like BiDi, giving network designers 50 percent more distance than with OM4.

But are we going fast enough?

Industry projections forecast that 800G optics will be needed within the next two years. So, in September 2019, an 800G pluggable MSA was formed to develop new applications, including a low-cost 8x100G SR multimode module for 60- to 100-meter spans. The goal is to deliver an early-market low-cost 800G SR8 solution that would enable data centers to support low-cost server applications. The 800G pluggable would support increasing switch radix and decreasing per-rack server counts.

Meanwhile, the IEEE 802.3db task force is working on low-cost VCSEL solutions for 100G/wavelength and has demonstrated the feasibility of reaching 100 meters over OM4 MMF. If successful, this work could transform server connections from in-rack DAC to MOR/EOR high-radix switches. It would offer low-cost optical connectivity and extend long-term application support for legacy MMF cabling.

So, where are we?

Things are moving fast, and—spoiler alert—they're about to get much faster. The good news is that, between the standards bodies and the industry, significant and promising developments are underway that could get data centers to 400G and 800G. Clearing the technological hurdles is only half the challenge, however. The other is timing. With refresh cycles running every two to three years and new technologies coming online at an accelerating rate, it becomes more difficult for operators to time their transitions properly—and more expensive if they fail to get it right.

There are lots of moving pieces. A technology partner like CommScope can help you navigate the changing terrain and make the decisions that are in your best long-term interest.

5 / MTDCs at the network edge

"Edge computing" and "edge data centers" are terms that have become more common in the IT Industry as of late. Multitenant data centers (MTDCs) are now living on the edge to capitalize on their network location. To understand how and why, we first need to define the "edge."

What is the "edge" and where is it located?

The term edge is somewhat misleading as it is located closer to the core of the network than the name might suggest – and there is not one concrete edge definition, but two.

The first definition is that of the customer edge, located on the customer's premises to support ultra-low latency applications. An example would be a manufacturing plant that requires a network to support fully automated robotics enabled by 5G.

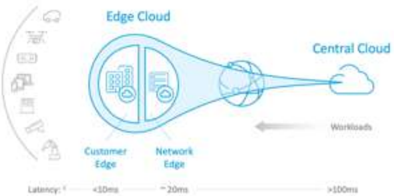

The second definition is that of the network edge, located toward the network core. This paradigm helps to support the low-latency needed for applications like cloud-assisted driving and high-resolution gaming. It is at the network edge where MTDCs thrive.

Flexible and accommodating

MTDCs that are flexible and ready to accommodate a variety of customer configurations can fully take advantage of their location at the edge of the network, as well as proximity to areas of dense population. Some MTDC customers will know what their requirements are and provide their own equipment. Other customers moving their operations off-premise to an MTDC will require expert guidance to support their applications. A successful MTDC will be able to accommodate both scenarios.

Flexibility is needed not only within the initial setup; the connectivity within the MTDC must be flexible on day one and two as well. To enable this flexibility, you need to consider your structured cabling. The recommended architecture for flexibility within the customer cage is based around a leaf-and-spine architecture. When using high fiber-count trunk cables, like 24-fiber MPO, the backbone cabling between the leaf and spine switches 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.

Once the leaf-and spine architecture are in place, there are additional considerations to take into account to ensure that the MTDC can easily accommodate future speeds and bandwidth demands in the cage. To achieve this, one must look to the server cabinets and their components, while keeping in mind that additions and alterations must be made simply and swiftly.

For a deeper dive into how MTDCs can optimize for capitalizing at the edge, check out CommScope's recent white paper entitled "New challenges and opportunities await MTDCs at the network edge."

6 /

The Evolving Role of the Data Center in the 5G-Enabled World

For decades, the data center has stood at or near the center of the network. For enterprises, telco carriers, cable operators and, more recently, service providers like Google and Facebook, the data center was the heart and muscle of IT.

The emergence of the cloud has emphasized the central importance of the modern data center. But listen closely and you'll hear the rumblings of change.

As networks plan for migration to 5G and IoT, IT managers are focusing on the edge and the increasing need to locate more capacity and processing power closer to the end users. As they do, they are re-evaluating the role of their data centers.

According to Gartner[1], by 2025, 75 percent of enterprise-generated data will be created and processed at the edge—up from just 10 percent in 2018.

At the same time, the volume of data is getting ready to hit another gear. A single autonomous car will churn out an average of 4,000 GB of data per hour of driving.

Networks are now scrambling to figure out how best to support huge increases in edge-based traffic volume as well as the demand for single-digital latency performance, without torpedoing the investment in their existing data centers.

[1] What Edge Computing Means for Infrastructure and Operations Leaders; Smarter with Gartner; October 3, 2018

A heavy investment in east-west network links and peer-to-peer redundant nodes is part of the answer, as is building more processing power where the data is created. But what about the data centers? What role will they play?

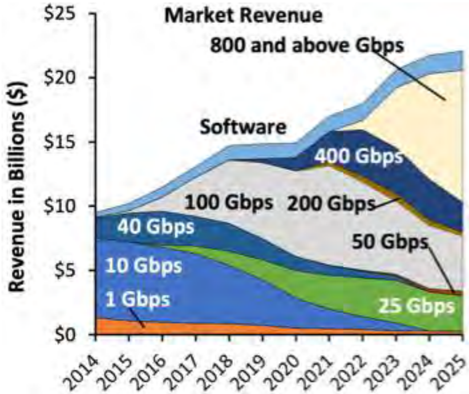

Source: 650 Group, Market Intelligence Report December 2020

The AI/ML feedback loop

The future business case for hyperscale and cloud-scale data centers lies in their massive processing and storage capacity. As activity heats up on the edge, the data center's power will be needed to create the algorithms that enable the data to be processed. In an IoT-empowered world, the importance of AI and machine learning (ML) cannot be understated. Neither can the role of the data center in making it happen.

Producing the algorithms needed to drive AI and ML requires massive amounts of data processing. Core data centers have begun deploying beefier CPUs teamed with tensor processing units (TPUs) or other specialty hardware. In addition, the effort requires very high-speed, high-capacity networks featuring an advanced switch layer feeding banks of servers—all working on the same problem. AI and ML models are the product of this intensive effort.

On the other end of the process, the AI and ML models need to be located where they can have the greatest business impact. For enterprise AI applications like facial recognition, for example, the ultra-low latency requirements dictate they be deployed locally, not at the core. But the models must also be adjusted periodically, so the data collected at the edge is then fed back to the data center in order to update and refine the algorithms.

Playing in the sandbox or owning it?

The AI/ML feedback loop is one example of how data centers will need to work to support a more expansive and diverse network ecosystem—not dominate it. For the largest players in the hyperscale data center space, adapting to a more distributed, collaborative environment will not come easily. They want to make sure that, if you're doing AI or ML or accessing the edge, you're going to do it on their platform, but not necessarily in their facilities.

Providers like AWS, Microsoft and Google are now pushing racks of capacity into customer locations—including private data centers, central offices and on-premise within the enterprise. This enables customers to build and run cloud-based applications from their facilities, using the provider's platform. Because these platforms are also imbedded in many of the carriers' systems, the customer can also run their applications anywhere the carrier has a presence. This model, still in its infancy, provides more flexibility for the customer while enabling the providers to control and stake a claim at the edge.

Meanwhile, other models hint at a more open and inclusive approach. For example, Vapor IO has built a business model featuring hosted data centers with standardized compute, storage and networking resources. Smaller customers—a gaming company, for example—can rent the literal machine in a Vapor IO data center near their customers and run their applications on the Vapor IO platform. And they'll charge you a revenue sharing model. For a small business trying to get access to the edge for their services, that's an attractive model.

Foundational challenges

As the vision for next-generation networks comes into focus, the industry must confront the challenges of implementation. Within the data center, we know what that looks like: Server connections will go from 50 Gb per lane to 100 Gb, switching bandwidth will increase to 25.6 Tb, and migration to 100 Gb technology will take us to 800 Gb pluggable modules.

Less clear is how we design the infrastructure from the core to the edge—specifically, how we execute the DCI architectures and metro and long-haul links, and support the high-redundancy peer-to-peer edge nodes. The other challenge is developing the orchestration and automation capabilities needed to manage and route the massive amounts of traffic. These issues are front and center as the industry moves toward a 5G/IoT-enabled network.

Getting there together

What we do know for sure is that the job of building and implementing next-generation networks will involve a coordinated effort.

The data center—whose ability to deliver low-cost, high-volume compute and storage cannot be duplicated at the edge—will certainly have a role to play. But, as responsibilities within the network become more distributed, the data center's job will be subordinate to that of the larger ecosystem.

Tying it all together will be a faster, more reliable physical layer, beginning at the core and extending to the furthest edges of the network. It will be this cabling and connectivity platform—powered by PAM4 and coherent processing technologies, featuring co-packaged and digital coherent optics and packaged in ultra-high stranded, compact cabling—that will provide the continuous thread of consistent performance throughout.

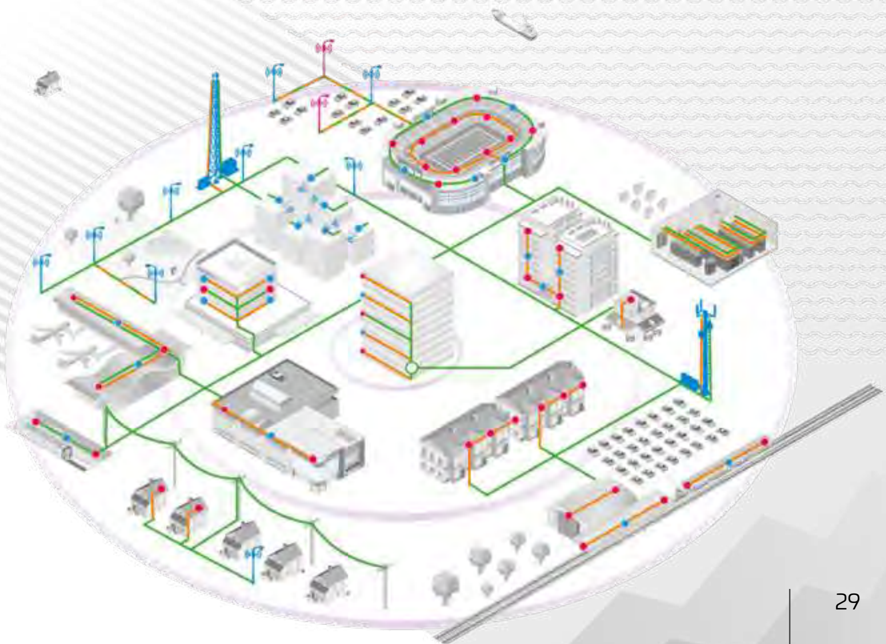

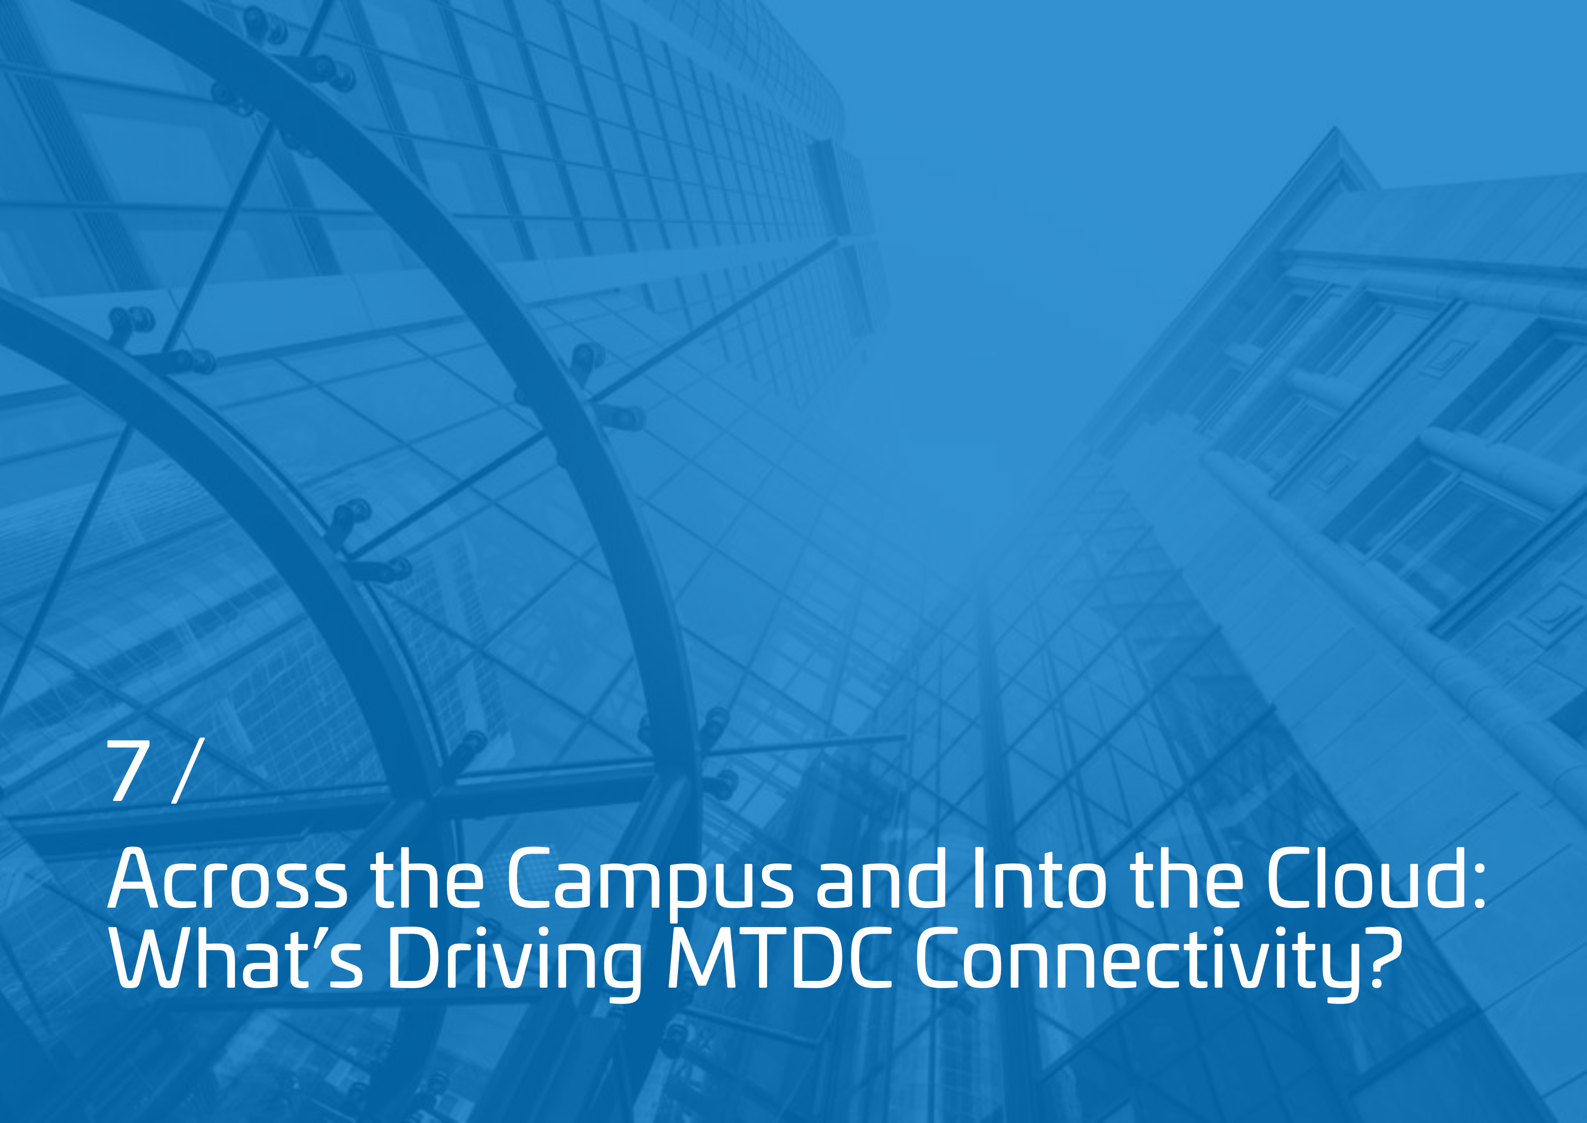

7 / Across the Campus and Into the Cloud: What's Driving MTDC Connectivity?

It's an incredible time to be working in the data center space and specifically the multi-tenant data centers (MTDCs). So much progress has been made recently in mechanical, electrical and cooling designs. The focus now shifts to the physical layer connectivity that enables tenants to quickly and easily scale to and from cloud platforms.

Inside the MTDC, customer networks are quickly flattening and spreading out east and west to handle the increase in data-driven demands. Once disparate, cages, suites and floors are now interconnected to keep pace with applications like IoT management, augmented reality clusters and artificial intelligence processors. However, connectivity into and within these data centers has lagged.

To address these gaps in connectivity, MTDC providers are using virtual networks as cloud on-ramps. Designing cabling architectures to connect within and between public, private, and hybrid cloud networks is challenging. The following highlights just a few of the many trends and strategies MTDCs are using to create a scalable approach to cloud interconnections.

Connecting the MTDC campus

The challenges of cloud-connectivity begin in the outside plant. High fiber-count cabling enables a mesh between current and future buildings. Prior to entering the facility, these outside plant (OSP) cables can be spliced to inside plant (ISP) fibers using a NEMA-rated splicing closure.

Alternatively, they can be spliced inside each building's entrance facility (EF) using high fiber count indoor-rated fiber entrance cabinets (FECs). Figure 1 shows a high-availability campus design featuring redundancy between DC 1 and DC 2 - 4.

As additional buildings on the campus are constructed, they are fed by DC 1. The net result is that the network traffic between any two layers in any building can be striped across the campus, increasing availability and reducing the potential for network downtime.

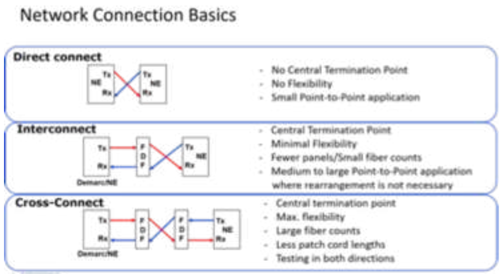

These building interconnects are increasingly being fed by high-density rollable ribbon fiber cables. The unique web-like configuration makes the ribbon more flexible, allowing manufacturers to load 3,456 fibers or more into an existing innerduct, or max out new larger duct banks created for this purpose. Rollable ribbon cables offer twice the density of conventionally packed fibers. Other benefits include:

- Smaller, lighter cables simplify handling, installation and subunits breakouts
- No preferential bend reduces the risk of installation error
- Easy separation and identifiable markings facilitate prep/splice and connectorization
- The smaller cable has a tighter bend radius for closures, panels and hand holes

Improved entrance facility connectivity

Inside the EF, where the OSP fiber connects to the ISP fiber, a focus on manageability has led to significant improvements in FECs and optical distribution frames (ODFs).

FECs are often overlooked as a strategic point of administration for the fiber plant. However, the ability to precisely identify, secure, and re-use stranded capacity can be the difference between days and months to turn up campus-wide connectivity.

Cabinet options include floor mount, wall mount, and rack mount designs capable of scaling to over 10,000 fibers. Other advantages include:

- Greater tray density for their mass fusion splicing
- Orderly transition from OSP to ISP cable
- Ability to break high-fiber cable down to a smaller cable counts

ODFs have also come a long way since they were first developed for telco and broadcast networks. For example, ODFs can now be ganged together in a row to support over 50k fibers with a single patch cord. CommScope's FLEX, FACT and NG4 ODFs are reflective of how far this technology has come. Mechanically these frames provide excellent front-side patch cord management and use a single-length patch cord to connect any two ports. This simplifies both inventory management and installer installation practices.

Capabilities for splicing high fiber count pre-terminated cables are engineered into the assemblies as demand for single-ended connector cabling continues to grow.

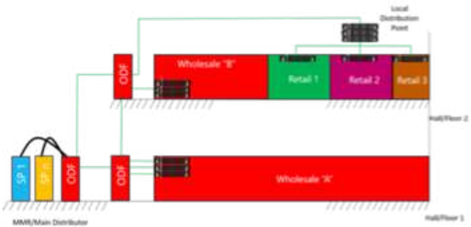

Supporting cloud connectivity within the MTDC

Access to cloud providers on the MTDC campus is becoming more critical as IT applications are moved off-premise and into the public and private cloud realms. Cloud providers and large enterprises require various construction and fire ratings between global regions, connectors types, and fiber counts to match to their network infrastructure, enabling them to scale quickly and with consistency regardless of installer skillset.

Of course, cloud connectivity requirements will vary based on the types of tenants. For example, traditional enterprises using private and hybrid cloud require basic connectivity to and within the cage or suite. In many cases, this means running bundles of single duplex jumpers, as needed, to connect providers to a tenant rack.

To connect the cages/suites from the meet-me-room (MMR), MTDCs are now deploying fiber in increments of 12 and 24 SMF. Once the tenant has moved out, de-installing doesn't require heavy cable mining. The MTDC can re-use the "last meter" runs into reconfigured white space by simply coiling it up and redeploying it to another cage demarc location. The structured cabling inside these cages—generally less than, but not limited to 100 cabinets—allows scalable connectivity to private and public providers.

Using cables above 24f can also provide significantly more density and less maintenance with similar labor costs.

Cloud service providers, on the other hand, have extensive and highly volatile connectivity requirements. Fiber counts to these cages are generally much higher than enterprises and sometimes cages can be tied together directly. These providers are deploying new physical infrastructure cabling several times per year and are constantly evaluating and refining their design based on CapEx considerations.

This involves scrutinizing the cost-effectiveness of everything from optical transceivers and AOCs, to fiber types and pre-term components.

Generally, the cloud provider cabling links into the MTDCs use higher fiber counts with diverse cable routing to support fewer points of failure. The end goal is to deliver predictable building blocks at varying densities and footprints. Uniformity can be hard to achieve, because, counterintuitively, as transceivers become more specialized, finding the right match of optics and connectors often becomes harder instead of easier.

For example, today's transceivers have varying requirements regarding connector types and loss budgets. Duplex SC and LC connectors no longer support all optical transeiver options. New, higher density, application-specific connectors such as the SN connector are now being deployed in cloud scale networks. Therefore, selecting transceivers with the greatest interoperability among connector footprints and fiber counts make the most sense.

Stay connected, keep informed

Across the MTDC campus, the need to interconnect the various buildings and provide the cloud-based connectivity that is vital to the success of retail and wholesale clients is driving changes in network architectures, both inside and out. Admittedly, this blog only begins to scratch the surface of an increasingly complex and sprawling topic.

For more information on the trends, I encourage you to check out the recent CommScope webinar, Connecting into the Cloud: Enabling a scalable approach to cloud interconnections in your MTDC. And to keep abreast of the fast-moving developments, rely on CommScope. It's our job to know what's next.

What's Next?

Things are moving fast, and they're about to get much faster! 2020 was a year of unpredictability and adaptation for everyone, but in the face of unforeseen challenges, data centers have experienced new levels of expansion and growth to accommodate rising connectivity demands. And as we look to 2021 and beyond, this growth is only set to increase.

The emergence of technologies like 5G and AI are key steps along the data center's expansion trajectory, and will lay the foundation for 800G, 1,6T schemes, and more! As networks ramp up their support for 5G and IoT, IT managers are focusing their efforts on the edge and the increasing need to locate more capacity. From rollable ribbon fiber cables to 400GE optical transceivers - network providers are developing future-proof solutions that will help to lead the way to a future of seamless end-to-end connectivity at every touchpoint.

Whether you're a player focused on the edge, a hyperscaler, a multi-tenant provider or a system integrator, there is plenty of room for everybody as the industry continues to grow. At CommScope, we're always looking at what's next and what's at the forefront of the ever-evolving data center landscape; contact us if you'd like to discuss your options when preparing for migration to higher speeds.

COMMSCOPE®

commscope.com

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

EB-115375-EN