COMMSCOPE®

WHAT'S NEXT FOR THE DATA CENTER

2023 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 looking ahead to 2023 we can count on much of the same. 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 Gbps Ethernet, and dual-lane 50 Gbps 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

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.

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.

About the authors

Ken Hall

Ken Hall is data center architect for North America at CommScope, responsible for technology and thought leadership as well as optical infrastructure planning for global scale and related data centers. In this role he has been instrumental in the development and release of high speed, ultra low loss fiber optic solutions to efficiently enable network migration for data center operators.

Previously Ken worked with TE Connectivity/Tyco Electronics/ AMP in a variety of roles. His experience includes global Network OEM and Data Center program management and strategy, project management, marketing, industry standards and technical sales management. Ken was also responsible for industry standardization and proliferation of copper and fiber small form factor connectors and high-density interfaces for network electronics OEMs.

Ken has nine patents to date for fiber-optic connectors and infrastructure management systems.

Ken graduated with a Bachelor of Science from Shippensburg University. He is a registered Communication Distribution Designer (RCDD) and Network Technology Systems Designer(NTS).

Hans-Jürgen Niethammer

Hans-Jürgen joined CommScope's cabling division in July 1994 and has held several key positions in product management, technical services and marketing including Director of Program Management EMEA, Director of Marketing EMEA and Director of Technical Services & Sales Operations EMEA.

Since January 2013, Hans-Jürgen is responsible for CommScope's data center market development in EMEA, ensuring that CommScope's solutions enable customer's data center infrastructures to be agile, flexible and scalable in order to meet the requirements in this dynamic market segment, today and in future.

Hans-Jürgen is an international expert for data centers, fiber optics and AIM systems, a member of several ISO/IEC and CENELEC standardization committees and editor of several international standards.

Hans-Jürgen holds a chartered engineer degree in Electronic Engineering and is a graduated state certified business economist.

About the authors

Alastair Waite

Alastair Waite joined CommScope in September 2003 as a Product Manager for the company's Enterprise Fibre Optic division, since that time he has held a number of key roles in the business including Head of Enterprise Product Management, for EMEA, Head of Market Management and the Data Center business leader in EMEA.

Since January 2016, Alastair has had responsibility for architecting CommScope's data center solutions, ensuring that customer's infrastructures are positioned to grow as their operational needs expand in this dynamic segment of the market.

Prior to joining CommScope, Alastair was a Senior Product Line Manager for Optical Silicon at Conexant Semiconductor, where he had global responsibility for all of the company's optical interface products.

Alastair has a BSc in Electronic Engineering from UC Wales.

James Young

James is the director of CommScope's Enterprise Data Center division, overseeing strategy and providing leadership to product and field teams globally. 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).

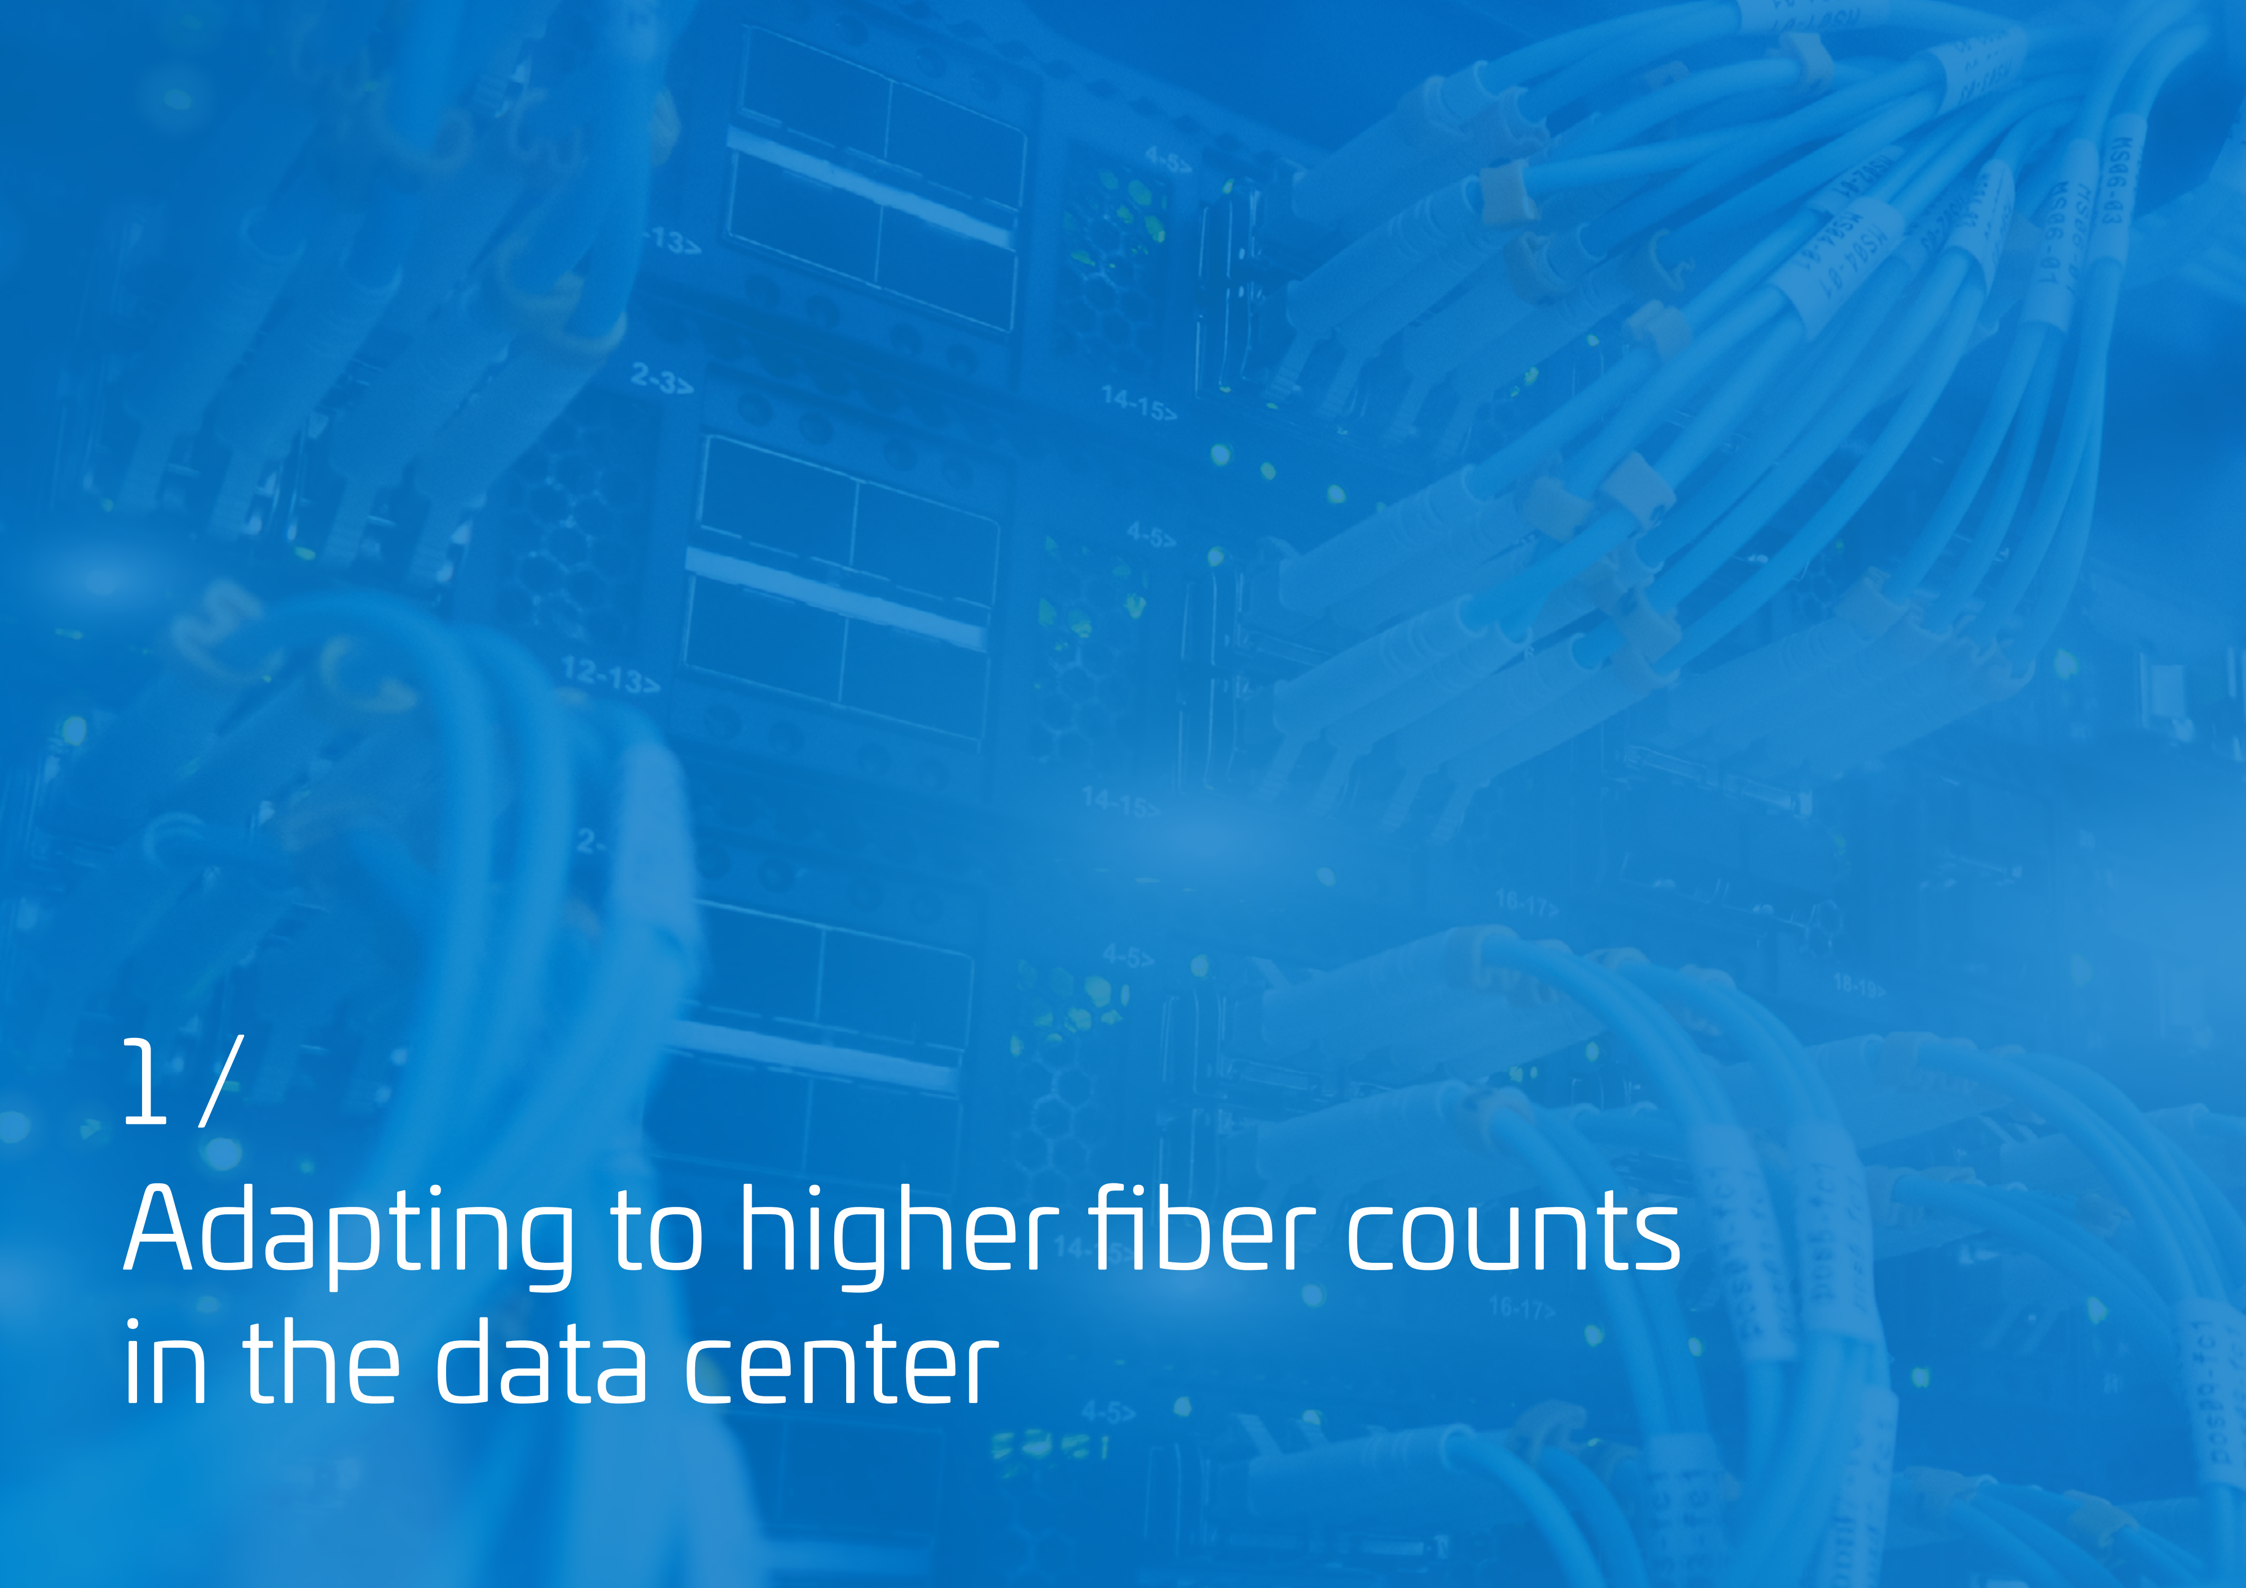

1 / Adapting to higher fiber counts in the data center

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. Data center interconnect (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 first: How do you deploy fiber 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: How do you break fiber out and manage it at the switches and server racks?

Rollable ribbon fiber cabling

The progression of fiber and optical network 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 in a cable construction without necessarily increasing the cabling footprint. Rollable ribbon fiber cabling is one of the more recent links in this chain of innovation.

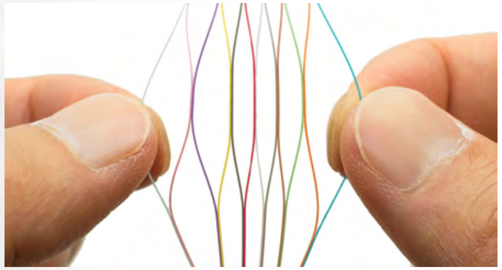

Rollable ribbon fiber is bonded at intermittent points. Source: ISE Magazine

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 effect 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 a rollable ribbon fiber cable, the optical 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 that 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.

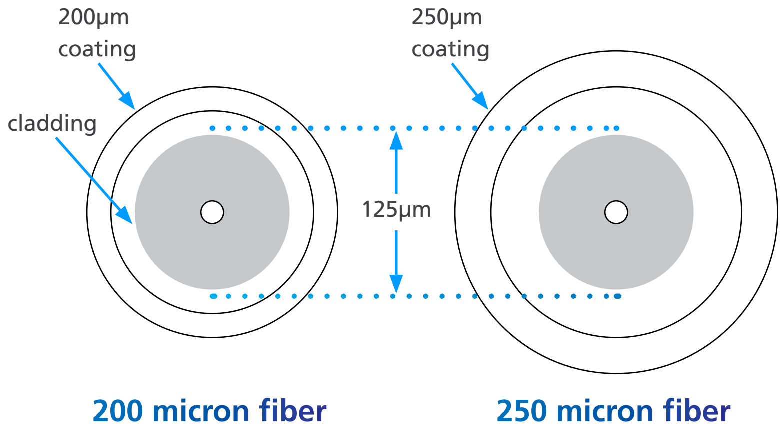

For optical performance and splice compatibility, 200-micron fiber features the same 125-micron core/cladding as the 250-micron alternative. Source: ISE Magazine

New chipsets are further complicating the challenge

All servers within a row are provisioned to support a given connection speed. In today's hyper-converged fabric networks, however, 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 bandwidth 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. Choosing a higher over subscription ratio might be tempting to reduce switch costs, however most modern Cloud and Hyperscale data center network designs target 3:1, or less, to deliver world class network performance.

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 with each switch to switch link contributing to the overall networks over subscription. 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 (TH4) switch—enabling 25.6 Tbps of Ethernet switching capacity in a single ASIC. This comes

less than two years after the introduction of the Tomahawk 3 (TH3), which clocked in at 12.8 Tbps per device.

These ASICs have not only increased lane speed; they have increased the number of ports they contain. Data centers can keep the oversubscription ratio in check. A switch built with a single TH3 ASIC 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 one pair, two pairs, four pairs, or eight pairs of fibers within the same 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 more suited to the demands of new workloads like machine learning (ML) or artificial intelligence (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 than 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. 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 industry standards and multi-source agreements (MSA's) increase in number, and deliver accelerated lane speeds, the cabling partner plays a bigger role in enabling the data center's technology roadmap. Currently, the standards regarding 100G/400G and evolving 800G involve a dizzying array of alternatives. Within each option, there are multiple approaches available to transport the data, including duplex, parallel and wavelength division multiplexing—each with a particular optimized application in mind. A cabling infrastructure design should be engineered to support as many of these transport alternatives as possible throughout its life span.

It all comes down to balance

As fiber counts grow, the amount of available space in the data center will continue to shrink the amount of available space in the data center does not necessarily track this growth. 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 and cloud level—had better strap in. As bandwidth demands and service offerings increase and latency becomes more critical to the end user/machine, 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 /
The cost/benefit analysis behind OM5

To address the growing demand for faster network speeds, IEEE, the standardization committee for Ethernet, is applying 3 key technologies to increase the ethernet bandwidth:

- Increasing the number of data streams (lanes) by increasing the number of fibers for transmission. While traditionally every data lane used 2 optical fibers, today we see ethernet applications using eight, 16, or even 32 optical fibers. From a cabling perspective, the increasing number of optical fiber per application is handled by multi-fiber connectors (MPO).
- Increase the baud-rate modulation. More specifically, this involves stepping from a 25 Baud-NRZ scheme to a 50 Baud-PAM4 scheme. Of course, with the doubling of speed in PAM4, there are trade-offs in terms of signal quality and transceiver costs.
- Upgrade the per-fiber capacity. WDM technology can run multiple data streams using different wavelengths per fiber core—enabling network managers to support up to eight WDM channels per optical fiber.

While many applications apply one of the described technologies to increase speed, other applications use more than one. 400GBase-SR4.2, for example, combines the benefits of more parallel optical fibers (eight) and the use of short wavelength division multiplexing (SWDM; mostly adopted as Bi-Di technology).

The challenge for data center network managers is mapping out their journey to 400G/800G and beyond without knowing what twists and turns lie ahead and how or where the three paths will intersect, as with 400GBase-SR4.2. Therein lies the value of OM5 multimode optical fiber, a new multimode optical fiber designed and standardized to support multiple wavelengths in a single fiber core.

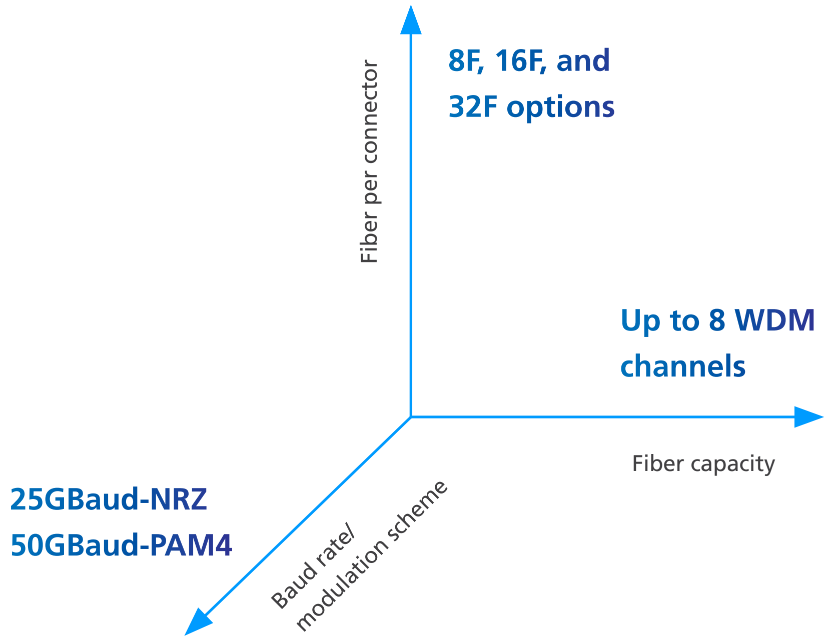

Three paths to higher Ethernet speeds

OM5 multimode fiber

Introduced in 2016, OM5 is the first approved WBMMF (wide band multimode fiber). The characteristics of OM5 are optimized to handle high-speed data center applications using several wavelengths per fiber (Bi-Di). The technical details and operational benefits of the OM5 technology are widely known:

- Reduces parallel fiber counts
- Lowers cabled fiber attenuation
- Enables wider effective modal bandwidth (EMB)
- Has 50 percent longer reach than OM4

Because OM5 shares the same geometry (50 µm core, 125 µm cladding) with OM3 and OM4, it is fully backward compatible with these optical fiber types.

OM5 vs. OM4: a closer look at the cost/benefit analysis

When compared side-by-side, OM5 offers some clear technical and performance advantages over OM4.

Yet, despite OM5's benefits, its adoption has met with resistance from some data center operators (in much the same way data centers were slow to replace OM3 when OM4 was introduced). One potential reason for the hesitancy to shift to OM5 is its higher price. However, a closer look at the cost/benefit analysis of OM5 vs. OM4 suggests a different story.

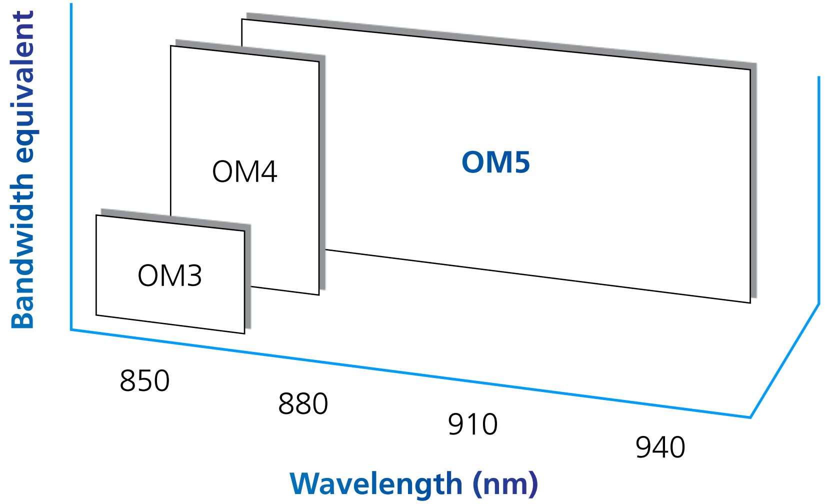

Difference in effective modal bandwidth

Costs

Opponents of OM5 optical fiber like to point to its 50-60 percent higher purchase price versus OM4 optical fiber. But to look only at the optical fiber price is to ignore the bigger picture in which data center managers must operate. First of all, with putting optical fiber in a fiber trunk cable, the price premium of an OM5 fiber cable shrinks to approximately 16 percent compared to an OM4 fiber cable. And secondly, by the time you add the cost of patch panels and cassettes on both ends of the trunk cabling, the original 50-60 percent price premium of the optical fiber is significantly diluted. In fact, when you compare the total cost of identically configured links, OM4 and OM5, the OM5 is only about 6.2 percent more expensive than the OM4.

Example scenario

Consider a real-world case involving a 144-fiber, 50-meter trunk cable connected to four 2xMPO12-to- LC modules and one high-density 1U panel on either end. Approximate costs are given for each set of components. Note that the total costs for OM4 and OM5 are identical for the panels and the cassettes, just the fiber trunk cable shows a difference of aapproximately 16 percent between OM4 and OM5.

When you calculate the overall cost for each end-to-end link ($10,310 for the OM4 and $10,947 for the OM5), the cost difference of $637 represents a bump of 6.2 percent.

HD panel	4 x 2x MPO12 to LC modules	50m, 144-fiber trunk	4 x 2x MPO12 to LC modules	HD panel		
OM4: **$318**	OM4: **$2,880**	OM4: **$3,914**	OM4: **$2,880**	OM4: **$318**	=	**$10,310**
OM5: **$318**	OM5: **$2,880**	OM5: **$4,551**	OM5: **$2,880**	OM5: **$318**	=	**$10,947**

Moreover, keep in mind that structured cabling represents only about 4 percent of the overall data center CapEx (including construction, supporting infrastructures like power and cooling and UPS and all IT equipment like switches, storage and servers). Therefore, switching to OM5 will increase the overall data center CapEx by 0.24 percent—less than one-quarter of 1 percent. In absolute dollars, this means an extra $2,400 for every $1,000,000 of data center CapEx.

Benefits

The question for data center managers is whether OM5's incremental cost increase outweighs its benefits. Here are just a few of the direct and indirect benefits.

OM5 provides higher capacity per fiber—resulting in fewer fibers and longer reach in a Bi-Di application. The extended reach in 100G and 400G applications with a Bi-Di is 50 percent farther than OM4, and it uses 50 percent fewer fibers. OM5 enables support for 100G (and even more, looking at 800G and 1.6T Bi-Di) using just two fibers. Plus, with the ability to span 150 meters versus just 100 for OM4, it provides greater design flexibility as your cabling architectures evolve.

Reducing the number of parallel fibers required, OM5 also makes better use of existing fiber pathways—creating space should additional fibers need to be added.

OM5: A hedge against uncertainty?

Perhaps most importantly, OM5 gives you the freedom to leverage future technologies as they become available. Whether your path to 400G/800G and beyond involves more fibers per connector, more wavelengths per fiber, or adoption of higher modulation schemes, OM5 provides the application support, extended bandwidth and longer lengths you need.

When it comes to addressing the continual challenges of higher speed migrations in a quickly evolving environment, keeping your options open is everything. You may not need all the advantages OM5 offers, or they may prove pivotal down the road. You can't know—and that's the point. OM5 enables you to hedge your bets with minimal risk. That's CommScope's view; we'd like to hear yours.

3 /
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 network fabric 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.

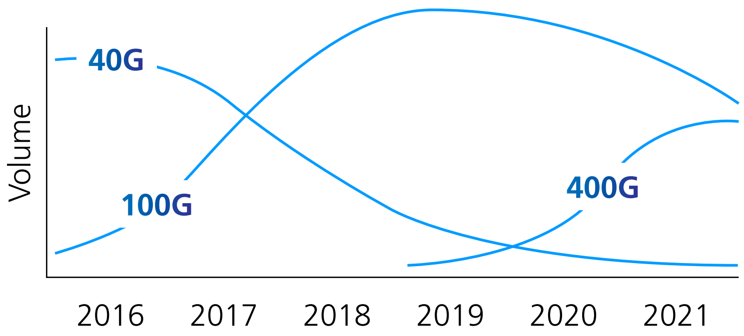

400G port numbers include both 8x50G and 4x100G implementations.
Source: NextPlatform 2018

400G 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 400G 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 away from that module technology to the second-generation, and size reduced, 400G 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 400G ports. Note that these modules 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 400G 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 400G; as 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 are already MSA group and standards committees 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, perhaps followed by 1.6T schemes 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.

4 /
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 higher. 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 even higher fiber core counts are being considered for the future.

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, fiber cable manufacturers are moving toward rollable-ribbon construction with 250- and/or 200-micron buffering.

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 leaving it to lay flat. On average, this type of design allows two 3,456 fiber cables to fit into a two-inch duct compared to a flat design that can accomodate only a single 1,728 fiber cable in the same space (using a 70 percent duct max fill rate).

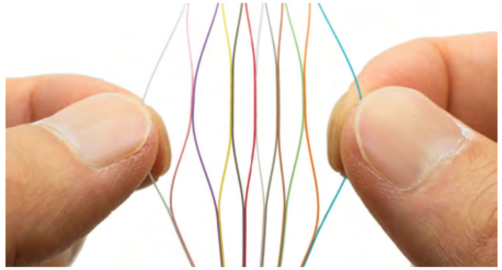

Rollable ribbon fiber is bonded at intermittent points. Source: ISE Magazine

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 manufacturers 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 400G, 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 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 above 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.

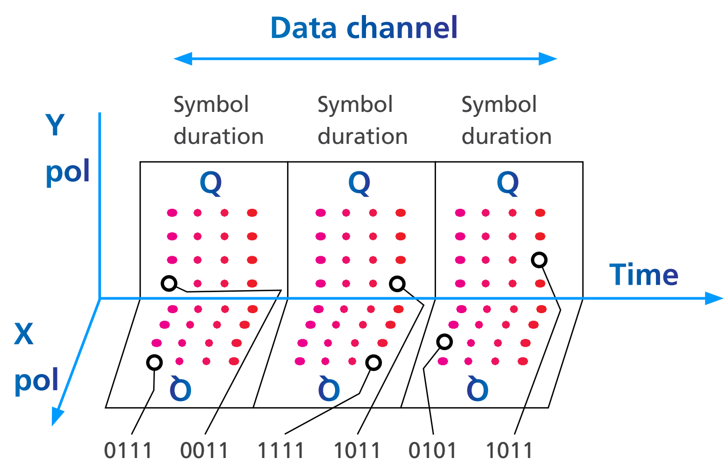

Source: www.cablelabs.com/point-to-point-coherent-optics-specifications

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 systematic 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.

5 / 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 AI and 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 manufacturers are pressing ahead to keep momentum going as the industry looks to support the ongoing transitions from 400G to 800G, 1.6Tb 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 (each of which are application, and therefore speed, specific).

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 with 16 fibers utilized) 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.

The demand for more capacity in enterprise data centers continues to escalate, and new strategies are required to scale the speed of the large installed base of multimode fiber (MMF) cabling infrastructures. In the past, adding multiple wavelengths to MMF has very successfully increased network speeds.

The Terabit Bidirectional (BiDi) Multi-Source Agreement (MSA) group—building on the success of 40G BiDI—formed to develop interoperable 800 Gbps and 1.6 Tbps optical interface specifications for parallel MMF. As a founding member of this BiDi MSA group, CommScope has led the introduction of multimode fibers, OM5, which are optimized to support applications that use multiple wavelengths like this MSA proposes.

MMF has been very popular with data center operators due to its support for short-reach high-speed links that aim to reduce the network hardware CapEx as well as (due to a lower power requirement) reduce OpEx. OM5 further enhances the value of MMF by extending the distance support for BiDi applications. In the case of IEEE 802.3400G BASE4.2, OM5 provides 50 percent more reach than does OM4 cabling. In the future, as we introduce next steps to 800G and 1.6T BiDi, the benefits of OM5 will become even more dramatic.

Using the technologies developed in IEEE802.3.db and IEEE802.3.cm, this new BiDi MSA will provide standards-based networks that will also enable single-fiber 100G BiDi, duplex fibers supporting 200G BiDi, and additional fibers added to reach 800G and 1.6T based on the evolving QSFP-DD and OSFP-XD MSAs with eight and 16 lanes, respectively.

On February 28, 2022, the MSA put it this way:

> "Leveraging a large installed base of 4-pair parallel MMF links, this MSA will enable an upgrade path for the parallel MMF based 400 Gb/s BiDi to 800 Gb/s and 1.6 Tb/s. BiDi technology has already proven successful as a way of providing an upgrade path for installed duplex MMF links from 40 Gb/s to 100 Gb/s. The Terabit BiDi MSA specifications will address applications for critical high-volume links in modern data centers between switches, and server-switch interconnects."

> "As a result of this MSA, the same parallel fiber infrastructure will be able to support data rates from 40 Gb/s up to 1.6 Tb/s. The MSA participants are responding to an industry need for lower cost and lower power solutions in 800 Gb/s and 1.6 Tb/s form factors that BiDi multimode technology can provide. For more information about the Terabit BiDi MSA, please visit terabit-bidi-msa.com."

Source: terabit.bidi.msa.com

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.

6 / 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 because it can be 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 support the low latency needed for applications like cloud-assisted driving and high-resolution gaming. It is at the network edge where MTDCs thrive.

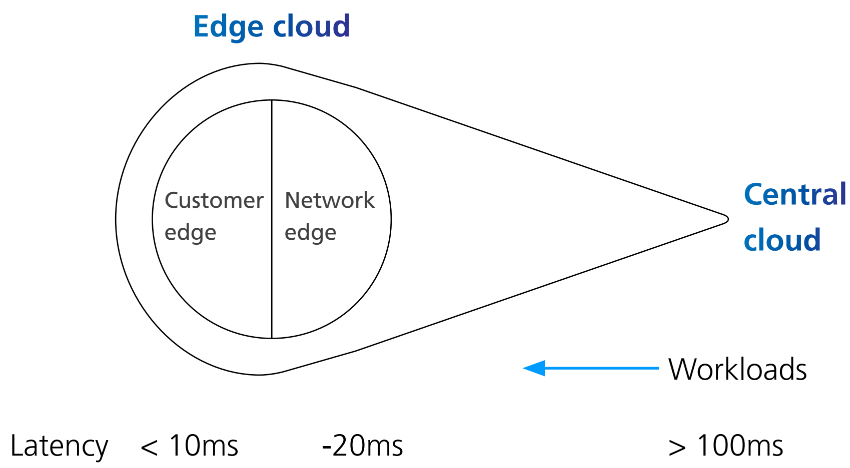

Flexible and accommodating

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-premises to an MTDC will require expert guidance to support their applications. A successful MTDC should be ready to accommodate both scenarios.

Operational 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 foundations i.e., the structured cabling. The recommended architecture for flexibility within the customer cage is based around a leaf-and-spine architecture. Using high fiber-count trunk cables, like 24- or 16-fiber MPO, allows the backbone cabling between the leaf-and-spine switches to remain fixed, because they have sufficient quantities of fibers to support future generations of network speeds.

For example, as Ethernet optics change from duplex to parallel ports, 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 is in place, there are additional considerations to take into account to ensure the MTDC can easily accommodate future speeds and bandwidth demands to and within in the cage. To achieve this, one must look to the server cabinets and their components and decide if the cabling pathways to those racks have sufficient space to support future moves, adds and changes—especially as new services and customers are introduced. Also, keep in mind that additions and alterations must be made simply and swiftly and possibly from a remote location. In such an instance, an automated infrastructure management system can monitor, map and document passive connectivity across an entire network. As more applications and services come to market, it soon becomes impractical to monitor and manage the cabling network manually.

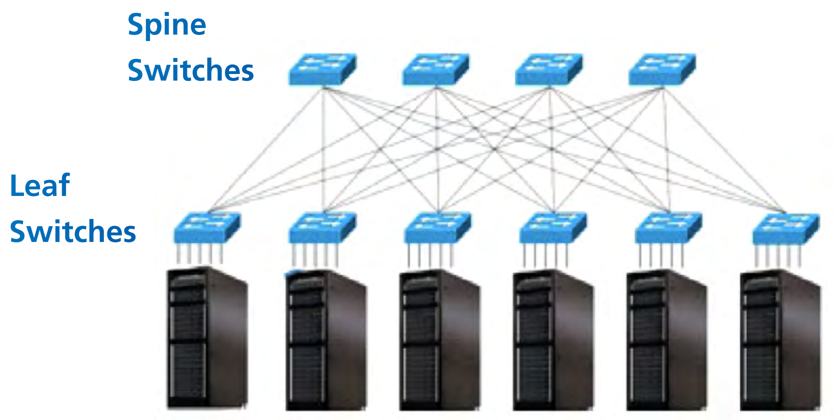

For a deeper dive into how MTDCs can optimize for capitalizing at the edge, check out CommScope's white paper: **"New challenges and opportunities await MTDCs at the network edge."**

7 / 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, and 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 4T 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.

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?

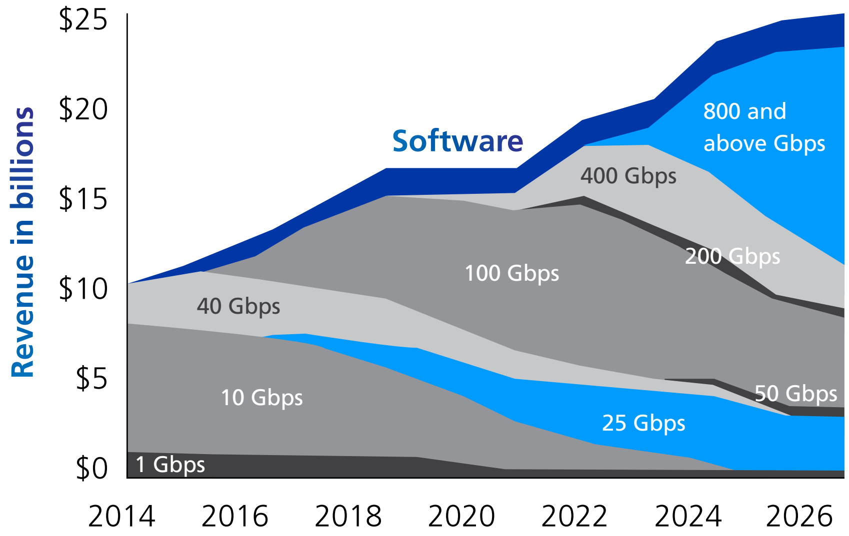

Source: 650 Group, Market Intelligence Report December 2020

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

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 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 larger 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-premises 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. Edge data center manufacturers are designing hosted data centers with standardized compute, storage and networking resources. Smaller customers—a gaming company, for example—can rent a virtual machine to host their customers and the data center operator will charge you on a revenue sharing model. For a small business competing for access to the edge, this is an attractive model (maybe the only way for them to compete).

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 50G per lane to 100G; switching bandwidth will increase to 25.6T; and migration to 100G technology will take us to 800G 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 traditional Ethernet optics and coherent processing technologies—that will fuel capacity. New switches featuring co-packaged optics and silicon photonics will drive more network efficiencies. And, of course more fiber everywhere—packaged in ultra-high-count, compact cabling—that will underpin the network performance evolution.

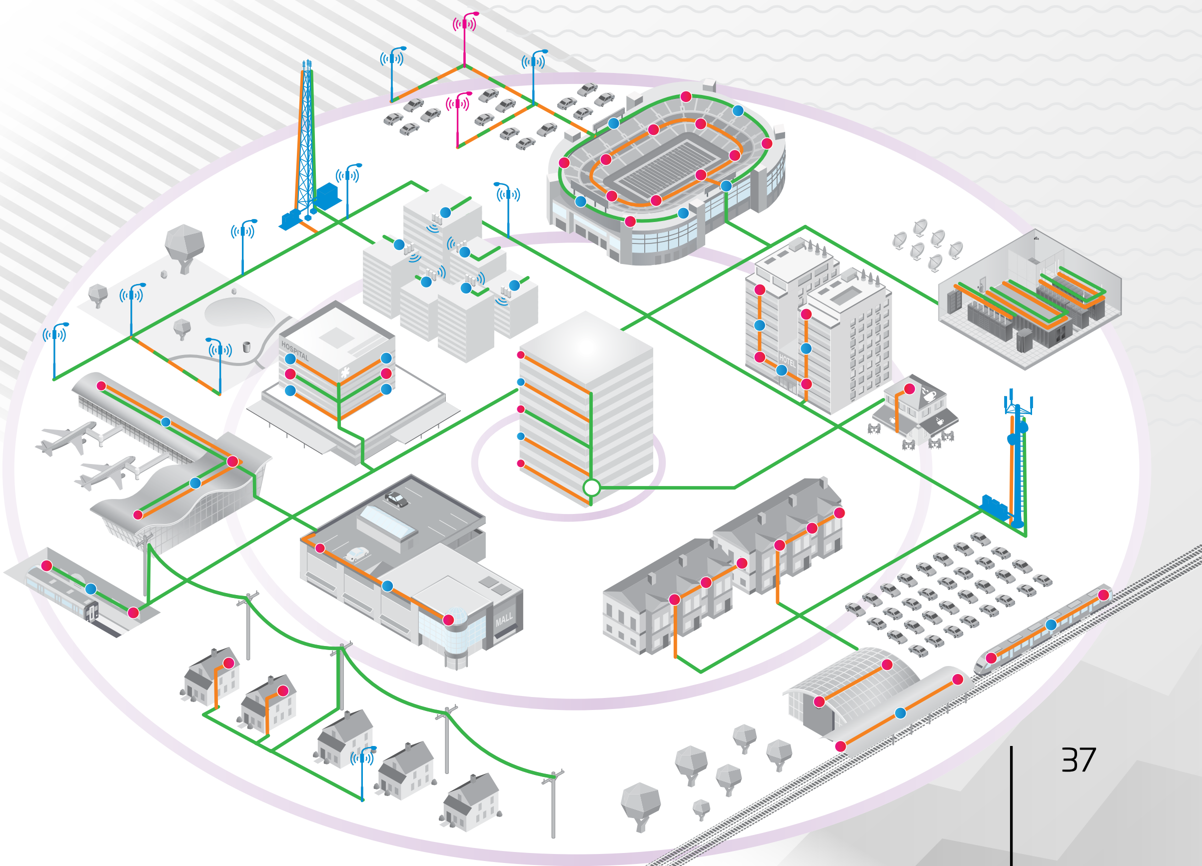

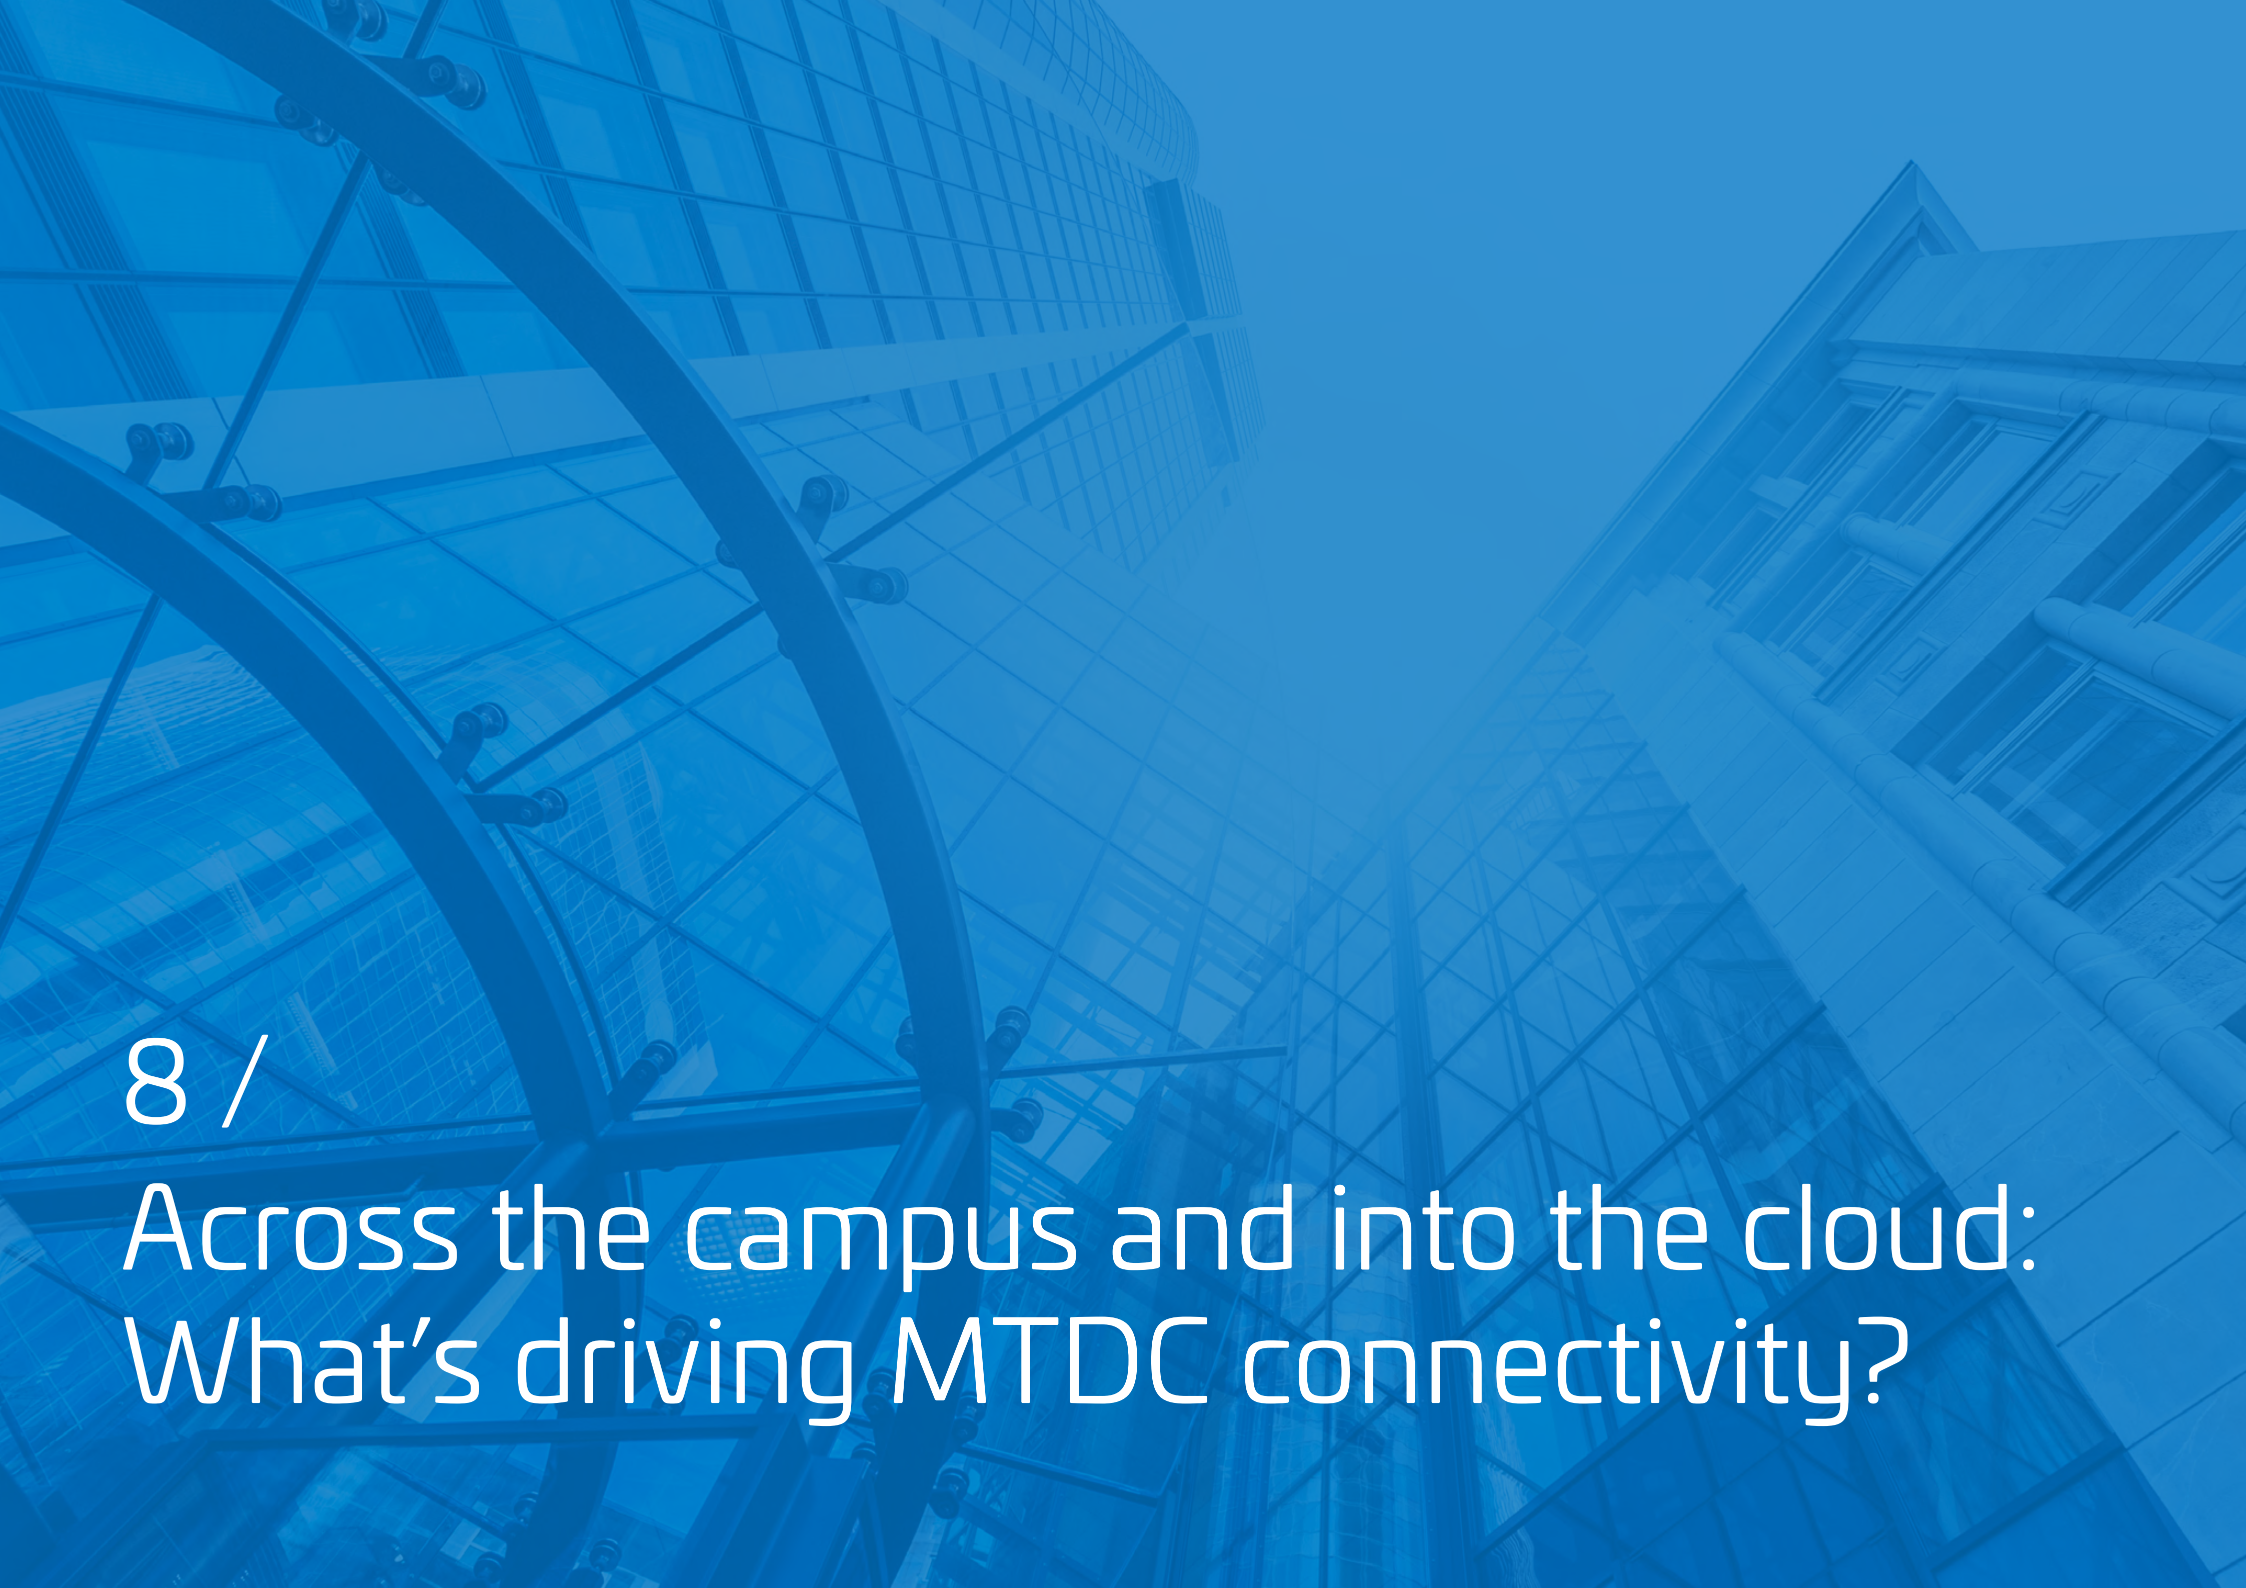

8 /
Across the campus and into the cloud: What’s driving MTDC connectivity?

It’s an incredible time to be working in the data center space—and specifically multitenant 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 AI 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 and diverse routing enable a mesh between current and future buildings. Prior to entering the facility, these outside plant (OSP) cables can be spliced to internal/external cables using a splicing closure for distribution within the data hall. This is applicable when panels and frames at the entrance facility have been pre-terminated with fiber-optic cables. Alternatively, OSP can be spliced immediately inside each building's entrance facility (EF) using high fiber count fiber entrance cabinets (FECs).

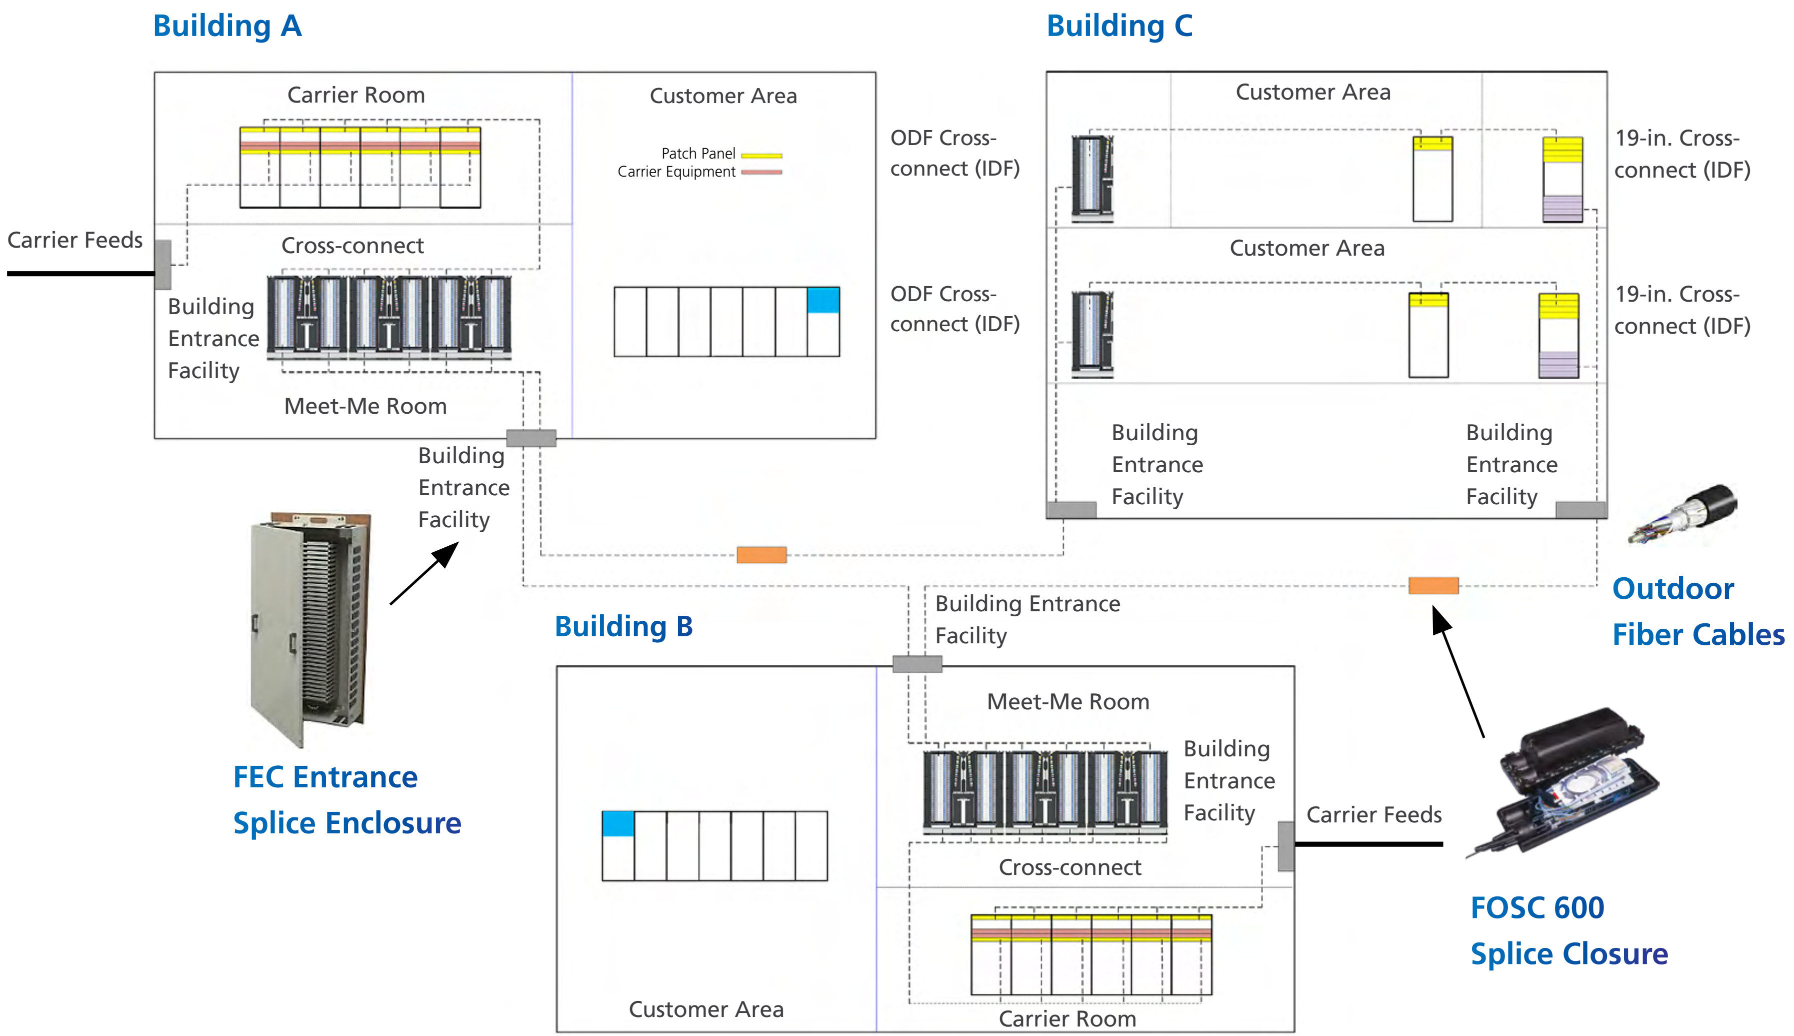

As additional buildings on the campus are constructed, they are fed by data center 1. The net result is that the network traffic between any two layers in any building can be easily re-directed around the campus—increasing availability and reducing the potential for network downtime.

These building interconnects are increasingly being fed by high- density rollable ribbon fiber cables. The unique web-like configuration makes the overall cable construction both smaller and 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 thousands of OSP fibers come together and connect to the ISP fiber, a focus on manageability has led to significant improvements in FECs and optical distribution frames (ODFs).

ODFs 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.

FEC options from CommScope 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 are critical to the smooth operation of a modern meet-me room (MMR), and have also come a long way since they were first developed for telco and broadcast networks. For example, ODFs with built-in intuitive routing can be ganged together in a row to support more than 50,000 fibers with a single patch cord length. Mechanically, ODFs also provide excellent front-side patch cord management–simplifying 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-premises and into the public and private cloud realms. Cloud providers and large enterprises, due to their international operations, will require various cable constructions and fire ratings to satisfy national regulations across regions. They will also demand different connector types and fiber counts to match to their network infrastructure architectures—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 may request lower density connectivity to and within the cage (or suite).

To connect the cages/suites from MMR, MTDCs are now deploying fiber in increments of 12 and 24 SMF on day one, as standard. 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 or 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.

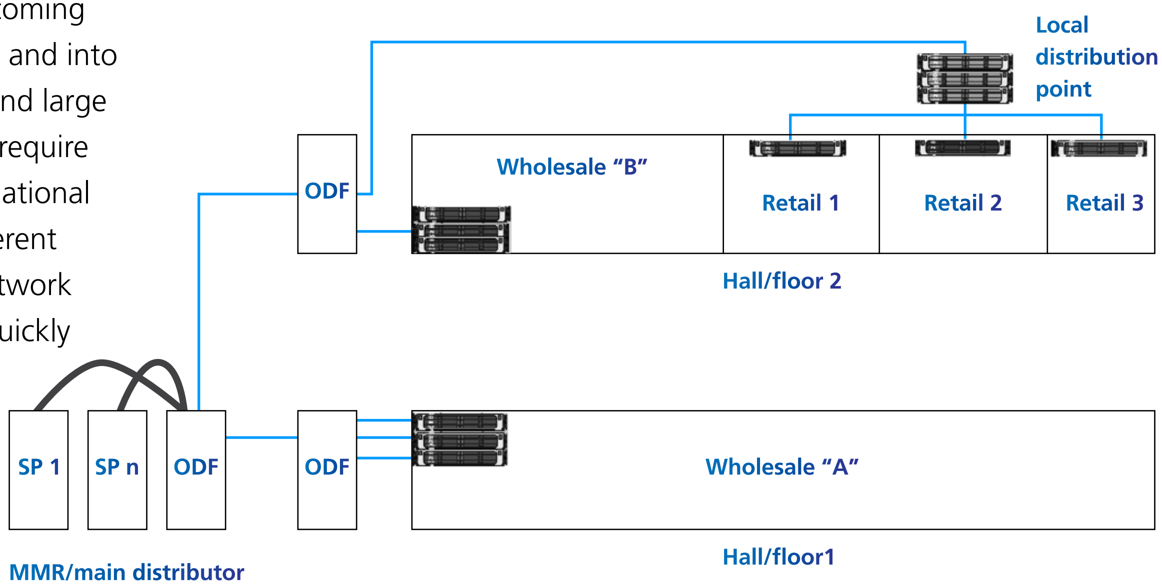

Cloud service providers, on the other hand, have extensive and highly volatile connectivity requirements. Fiber counts to these cages are generally much higher than enterprises, and sometimes cages can be tied together directly across a campus. These providers are deploying new physical infrastructure 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-terminated 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 transceiver options. New, higher density, application-specific connectors such as the SN connector are now being deployed in cloud scale networks. Therefore, it makes the most sense to select transceivers with the greatest interoperability among connector footprints and fiber counts.

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. This forward-looking view only scratches the surface of an increasingly complex and sprawling topic.

For more information on trends and to keep abreast of the fast-moving developments, rely on CommScope. It's our job to know what's next.

9 / The path to 1.6T begins now

The challenge in planning for exponential growth is that change is always more frequent and more disruptive than we expect. Hyperscale and multitenant data center managers are experiencing this firsthand. They are just beginning to migrate to 400G and 800G data speeds, but the bar has already been raised to 1.6T. The race is on, and everyone could win—assuming data center operators succeed in increasing application capacity and reducing the cost of services. In doing so, they can drive end-user costs lower while helping make internet more energy efficient. As with any leap forward, every success breeds another challenge. Higher capacity gives rise to new, more data-demanding, power-hungry applications, which require more capacity. And so the cycle repeats.

Driven by players like Google, Amazon and Meta (née Facebook), the explosion of cloud services, distributed cloud architectures, artificial intelligence, video, and mobile-application workloads will quickly outstrip the capabilities of 400G/800G networks. The problem isn't just bandwidth capacity; it's also operating efficiency. Data networking overhead is becoming an increasingly large part of overall delivery cost. Those costs, in turn, are driven by power consumption, which leads to the next-generation design objectives. The end goal is to reduce the power per bit and make this impossibly explosive growth a sustainable possibility.

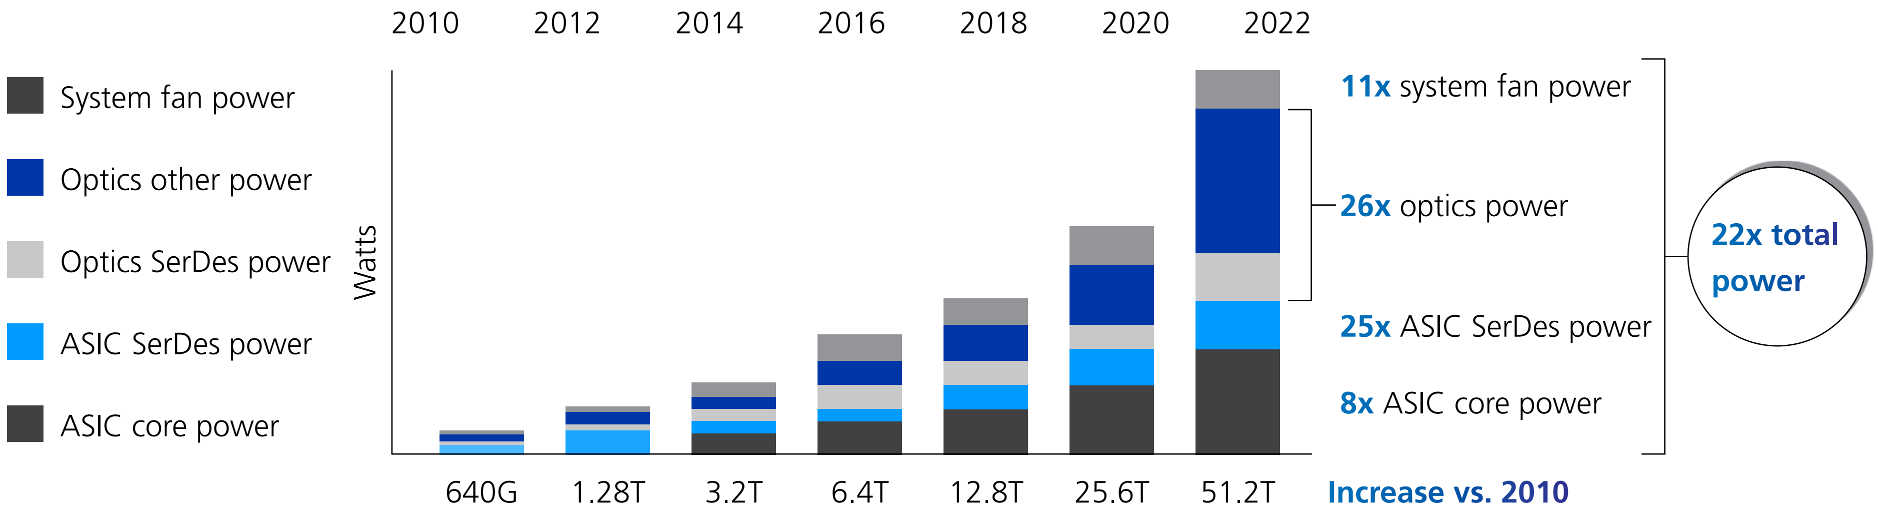

Source: www.ethernetalliance.org/wp-content/uploads/2021/02/TEF21.Day1_.Keynote.RChopra.pdf, January 25, 2021; Rakesh Chopra, Mark Nowell, Cisco Systems

Thinking power and networks

Expanding network capacity with the current generation of network switches would mean the power requirements will become unsupportable. (Mind you, this issue is also unfolding at a time when every corporate decision is scrutinized against the backdrop of environmental sustainability.)

Networks are under increasing pressure to reduce their power-per-bit ratio (the most common efficiency metric)—with targets eventually decreasing to 5 pJ/bit. Increasing the density (radix) of network switches is the demonstrated path to attack this problem. The result is greatly enhanced switch capacity and efficiency.

At a high level, overall switch power consumption is a growing concern. The keynote address presented at the 2021 Technology Exploration Forum showed switch power consumption rising 22x from 2010 to 2022. Looking deeper, the main component of the power increase is associated with electrical signaling between the ASIC and optical transmitter/receiver. Since electrical efficiency decreases as switching speed increases, switching speed is limited by electrical speed. Currently, that practical limit is 100G.

The path to lower power consumption lies in continuing the trend of larger, more efficient switching elements, more signaling speed, and more density. Theoretically, this path eventually leads to 102.4T—a goal that seems very challenging given a projection of the current switch designs. Therefore, some argue for a strategy based on point solutions. This would address the electrical signaling challenge (flyover cables vs PWB) and enable the continued use of pluggable optics. Increasing the signaling speed to 200G is also an option, while others suggest doubling the lane counts (OSFP-SD). Still, another camp advocates for a platform approach to move the industry toward a longer-term solution. A more systematic approach to radically increase density and reduce power per bit involves co-packaged optics (CPO).

The role of co-packaged and near-packaged optics (CPO/NPO)

Advocates of CPO and near-packaged optics (NPO) argue that achieving the needed power-per-bit objectives for 1.6T and 3.2T switches will require new architectures, and that CPO/NPO fits the bill. They make a good case in that CPO technologies limit electrical signaling to very short reaches—eliminating re-timers while optimizing FEC schemes. Taking these new technologies to market at scale would require an industry-wide effort to re-tool the networking ecosystem. New standards would greatly enhance this industry transformation.

One challenge with CPO is that they contain no field-serviceable optics and require very low failure rates (FIT)—something CPO must achieve compared to field-serviceable, pluggable optics. The bottom line with CPO is that it will take time to mature.

The industry will need new interoperability standards, and the supply chain must also evolve to support CPO. Many argue that, considering the risk factors associated with CPO, pluggable modules seem to make sense through 1.6T.

Switch designers and manufacturers have also proposed pluggable optics for 1.6T (based on 100G or 200G electrical SERDES speeds). This path would not require sweeping change, hence lowering the risk and shortening the time to market for this option. It's not a risk-free path, but proponents contend it poses far fewer challenges than the CPO path.

200G electrical signaling

Getting to the next switching node (doubling of capacity) can be done with more I/O ports or higher signaling speeds. The application and system-level drivers for each alternative are based on how the bandwidth will be used. More I/Os can be used to increase the number of devices a switch supports, whereas higher aggregate bandwidth combinations can be used for longer reach applications to reduce the number of fibers required to support the higher bandwidth.

In December 2021, the 4x400G MSA suggested a 1.6T module with options of 16x100 or 8x200G electrical lanes and a variety of optical options mapped through the 16-lane OSFP-XD form factor. A high-radix application would require 16 duplex connections (32 fibers)[1] at 100G (perhaps SR/DR 32) while longer reach options would meet up with previous generations at 200G/400G.

While suppliers have demonstrated the feasibility of 200G lanes, customers have concerns regarding the industry's ability to manufacture enough 200G optics to bring the cost down. Reproducing 100G reliability and the length of time needed to qualify the chips are also potential issues.[2]

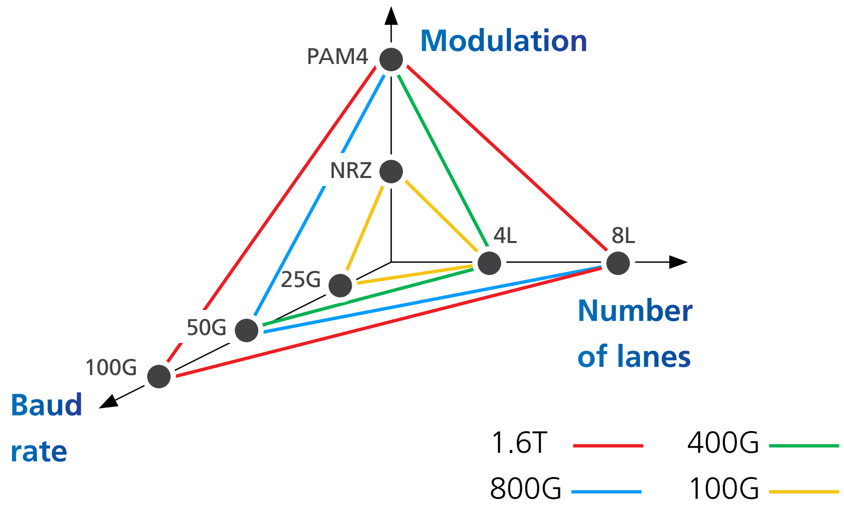

Potential paths to 200G lanes. Source: Marvell

1 OSFP-XD MSA included options for two MPO16 connectors supporting a total of 32 SMF or MMF fibers
2 The right path to 1.6T PAM4 optics: 8x200G or 16x100G; Light Counting; December 2021

Ultimately, any 1.6T migration route will involve more fiber. MPO16 will likely play a key role as it offers wider lanes with very low loss and high reliability. It also offers the capacity and flexibility to support higher radix applications. Meanwhile, as links inside the data center grow shorter, the equation tips toward multimode fiber with its lower-cost optics, improved latency, reduced power consumption and power/bit performance.

So, what about the long-anticipated predictions of copper's demise? At these higher speeds, look for copper I/Os to be very limited, as achieving a reasonable balance of power/bit and distance isn't likely. This is true even for short-reach applications that eventually will be dominated by optical systems.

What we do know

All this to say that, while the best path to 1.6T is uncertain, aspects of it are coming into focus. Higher capacity, higher speeds and significant improvement in efficiency will certainly be needed in a few short years. To be ready to scale these new technologies, we need to start designing and planning today.

Read more about steps you can take today to ensure your fiber infrastructure is ready for this future at **commscope.com.**

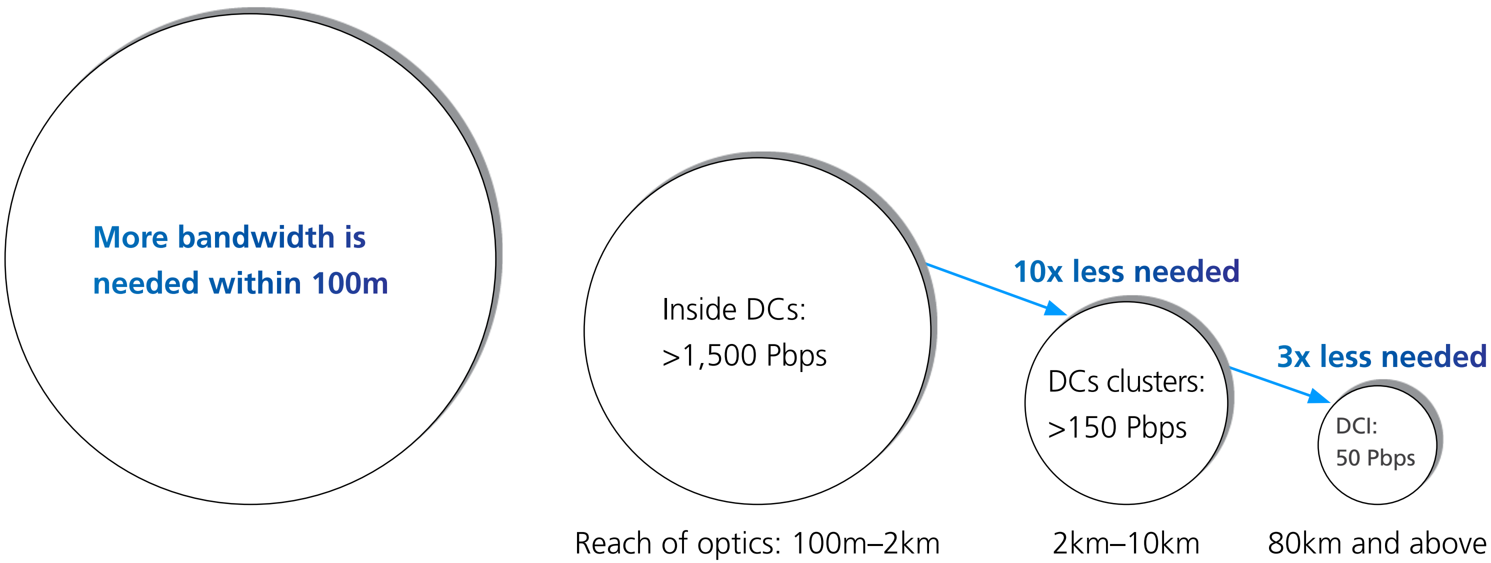

AOC and SR optics will occupy ToR, EoR, MoR for data rate at 100G, 400G and beyond. Source: LightCounting Mega Datacenter Optics Report

What's next?

Things are moving fast, and they're about to get much faster! 2021 and 2022 was unpredictable 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 2023 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 400G optical transceivers, network providers are developing future-proof solutions that will help 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 multitenant 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.com

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

Further information regarding CommScope's commitment can be found at www.commscope.com/About-Us/Corporate-Responsibility-and-Sustainability.

EB-115375.1-EN (07/22)