



WHAT'S NEXT FOR THE DATA CENTER:

400G and Beyond

Contents

About the author	3
Introduction	4
Chapter 1: Options for optical transceivers	5
Chapter 2: Densification and campus architecture	7
Chapter 3: The evolving role of the data center in a 5G-enabled world	11
Conclusion	16



About the author



James Young

James currently serves as 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. His sales experience includes electronic transmission components, telephony systems, network systems, LAN infrastructure products and fibre transmission system products. James has garnered substantial experience in OEM and channel marketing, as well as network operations as assistant director of CP's computers and communications group.

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).

400G in the data center

Data centers continue to grow, with global internet traffic expected to increase up to three times by 2021 - and the key to meeting this challenge is the ability to evolve the physical infrastructure layer in the data center.

Back in 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.

As larger hyperscale and cloud-based data centers confront their inevitable leap to 400G, network managers face a multitude of challenges and decisions. They must adapt to higher and higher fiber counts, choose the right optical transceivers, and prepare for the data mushroom effect that 5G will bring.

In this series of briefings, CommScope experts provide their take on the technologies and trends behind the move to higher speeds.




1 | 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.

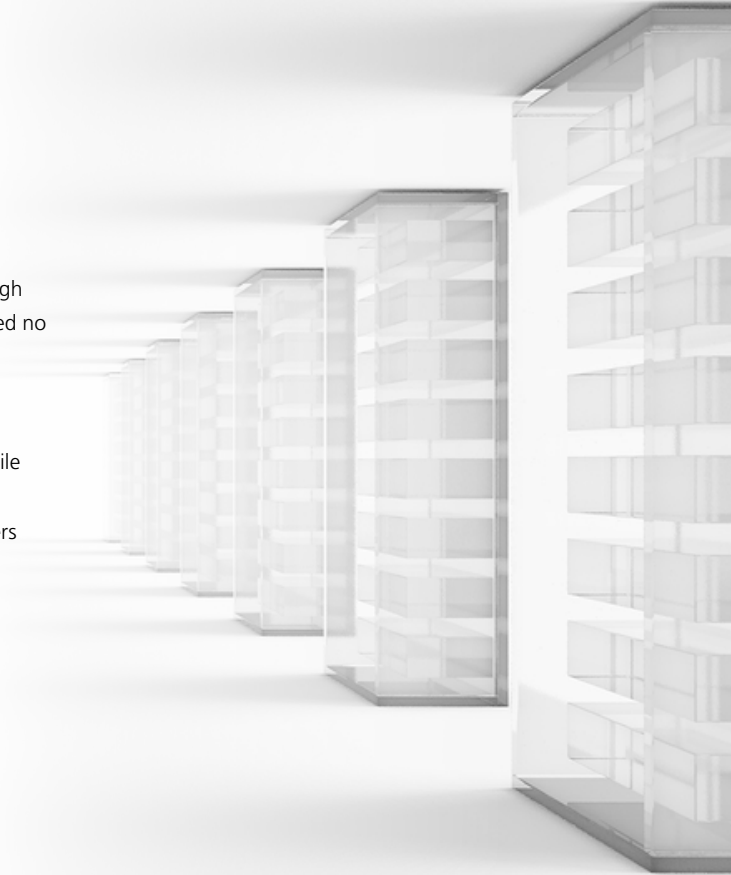
While the details involved with migrating to higher and higher speeds are daunting, it helps to put the process in perspective. As data center services evolve, storage and server speeds must also increase. Being able to support those higher speeds requires the right transmission media. In choosing the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

2 | 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.




Rollable ribbon fiber
Source: Fluke Networks

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.

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.



200-micron fiber vs 250-micron fiber
Source: ISE Magazine

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 optical 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.

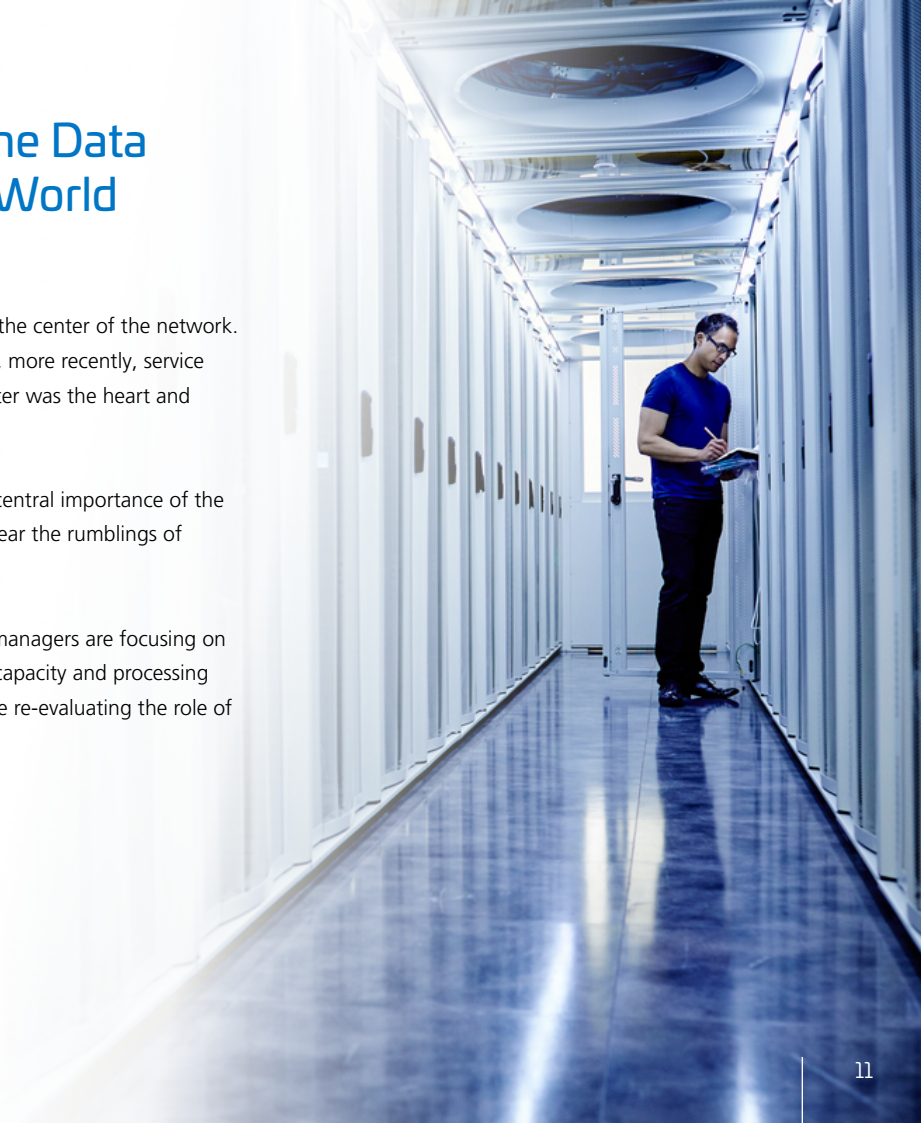


3 | The Evolving Role of the Data Center in 5G-Enabled World

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 Gartner¹, 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. 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?

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

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

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.



400G Takeaways

Whether you're a hyperscale data center, a player focused on the edge, or an infrastructure provider, in the next-generation network, there will be plenty of room—and work—for everybody. The slices aren't getting smaller; the pie is getting larger.

The need to futureproof the data center will continue to be at the top of the list for data center managers as technologies such as 5G, the IoT and AI drive the need for more data and higher speeds.


The evolution to 400G is a key step along the data center's expansion trajectory - and will lay the foundation for 800G, 1,6T schemes, and beyond!

When choosing the right transmission media, it's best to anticipate the topology required for future service delivery. Plan with the future in mind.

Selecting the appropriate optical transceivers, fiber packaging and design for increased density, and ensuring that AI and Machine Learning modules are located close enough to users, will all be an important consideration for the 5G evolution. However, when it comes to the data center, there's rarely a one-size-fits-all solution.

At CommScope, it's our job to know what's next; we stay on top of the ever-evolving data center landscape to support you and your networks' growth.

Contact us if you'd like to discuss your options when migrating to 400G.





COMMSCOPE®

commscope.com

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

© 2020 CommScope, Inc. All rights reserved.

Unless otherwise noted, all trademarks identified by ® or TM are registered trademarks, respectively, of CommScope, Inc. This document is for planning purposes only and is not intended to modify or supplement any specifications or warranties relating to CommScope products or services. CommScope is committed to the highest standards of business integrity and environmental sustainability with a number of CommScope's facilities across the globe certified in accordance with international standards, including ISO 9001, TL 9000, and ISO 14001.

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

EB-114804-EN