ICT TODAY

THE OFFICIAL TRADE JOURNAL OF BICSI

April/May/June 2020
Volume 41, Number 2

BICSI

TOMORROW'S MODERN DATA CENTER

PLUS:

+ Data Center Optical Fiber Cabling Decisions Based on Scale
+ Definition of an Edge Data Center—404 NOT FOUND
+ The Connector: Revolutionizing Connected Infrastructure

TOMORROW'S MODERN DATA CENTER

COVER ARTICLE By Seán Adam

It is said that the only constant in life is change and few in the ICT industry would disagree that the pace of changes is increasing at a dramatic rate. Hand-in-hand with this pace of change is the ever-increasing consumption of data. Cisco claims that by 2022 cellular data usage will eclipse 77 exabytes per month. That is the number 77 followed by 18 zeroes!

In part, this will be enabled by emerging technologies, such as 5G and the next generation WiFi 802.11ax standard, also known as Wi-Fi 6. These technologies bring a promise of increased speeds, faster response (i.e., latency), and the reality of true internet of things (IoT). The number of connected devices is expected to exceed 25 billion by 2025.[1]

Consider what the transition from 4G to 5G looks like and why this is important to the impact on the data center of tomorrow. Table 1 includes some of the key specifications.

The step function improvement in latency and data rates makes more and larger content more easily available. This efficiency in accessing data will be a self-fulfilling prophecy driving the need for a massive increase in storage.

However, technology does not drive growth; it simply enables it. It will be the end applications that define the reality of 5G and the infrastructure necessary to support it. Many people already have at least one wearable, such as a smartwatch or fitness tracker. Combine that with smartphones and the beginning of a more connected on-person existence is achieved. Now merge that with surroundings that integrate traffic and security systems in smart cities and smart buildings. Overlay with augmented reality, and the digital world and real world merge into one. Dreams of fully autonomous cars promise to transform roadways and, for many people, their greatest singular time investment in a day, their commute. Nearer on the horizon will be fleets of drones, delivering everything from packages, to pizza, to drug prescriptions. UPS Flight Forward, a subsidiary of UPS,

	From the 4G World	To the 5G World
Latency	> 10 ms	< 1 ms
Data Traffic	7.2 exabytes / month	> 50 exabytes / month
Peak Data Rates	1 Gb/s	20 Gb/s
Device Connection Density	100,000 connections/km^2	> 1 million connections/km^2

TABLE 1: Key specification comparison between 4G and 5G.

received certification from the U.S. government to operate its own drone fleet. Healthcare deliveries are its first target application.

The ability to transmit data quickly and at a high rate leads to the "four FRs" of tomorrow's communication infrastructure: Redundancy, Resiliency, Reliability, and Regionalization. It is here that the data center emerges as the heart and mind; not as a handful of large data centers but rather a vast, distributed network of data centers.

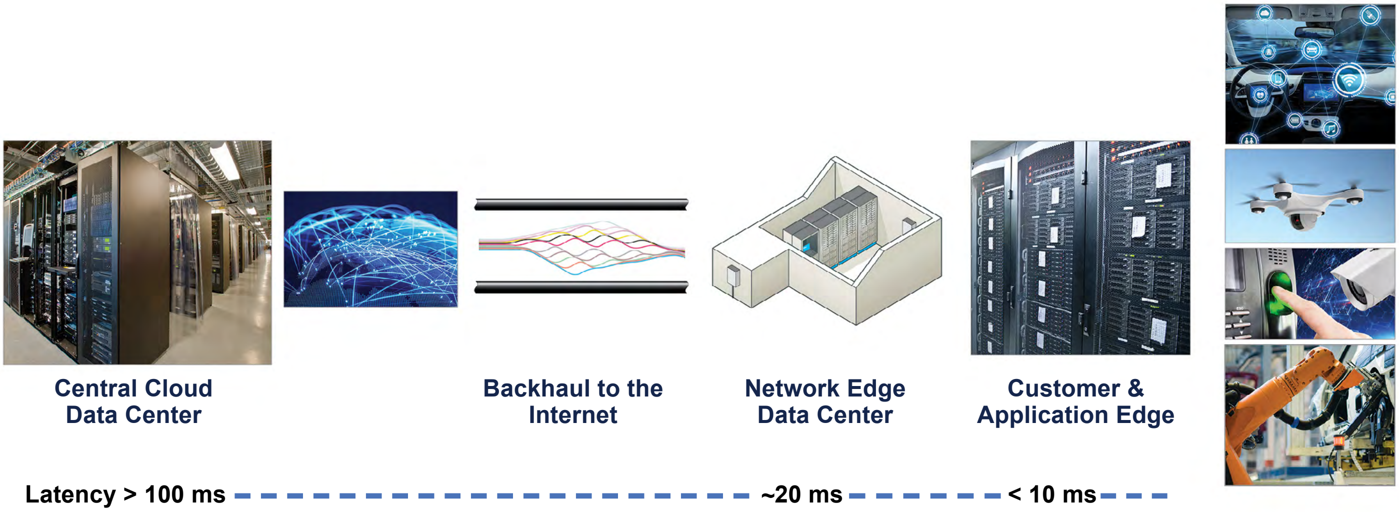

FIGURE 1: Latency to the edge.

Some industry analysts predict that the global population of hyperscale facilities, today numbering in the high hundreds, could grow to over 10,000. Driven by the need for low latency (Figure 1) and redundancy, what will emerge is a mass mesh of interconnected clusters of cloud data centers.

It is not only hyperscale data centers that will underpin this future. Retail and colocation providers act as today's edge for hyperscalers, bringing their end customers closer to their end applications. Wholesale colocations provide a way for hyperscalers and other large content providers to achieve a regional self-aligned footprint. And of course, there are the mixed strategies enabled by colocation facilities where hyperscalers, smaller enterprises, and service providers come together in a targeted location. Common to all of these is the need to provide a level of interconnection through a vast, deep, and rich optical fiber infrastructure.

THE OPTICAL FIBER INFRASTRUCTURE FOR THE MODERN DATA CENTER

The fiber infrastructure, made up of cables with thousands of fibers and millions of connections, ties together clusters of data centers, often across hundreds of kilometers. Within each cluster reside multiple data centers, thereby providing resiliency, redundancy, and regional storage.

To achieve this level of high-density interconnection, installers need to work with cables where a per-cable fiber count of 6,912 fibers is not out of the ordinary. Laying cables of this size and density is a challenge in and of itself, coupled with the need for potentially hundreds of splices and thousands of connectors. Limited duct space between clusters and buildings creates an installation nightmare and can consume a significant amount of time. Installers' key differentiator becomes their ability to deliver fast, reliable installations for Day 1 success. To address these challenges, optical fiber and cable manufacturers spend hundreds of millions of dollars in research to develop new solutions that decrease overall cable size and increase density-per-square-meter, while focusing on improving ease-of-use.

For the densities required by the data center, optical fiber ribbon solutions are key because they enable easier overall cable management and the ability to perform mass fusion splicing. Some modern cable designs combine

The fiber infrastructure, made up of cables with thousands of fibers and millions of connections, ties together clusters of data centers, often across hundreds of kilometers.

the benefits of ribbon and loose tube designs. Choosing a gel-less ribbon solution removes the typical mess associated with gel-filled cable solutions. This, alone, can remove hours from installation time. These modern, flexible cable designs enable the fiber to be pulled together easily in groups of 12 and as individual fibers.

Because size reduction is one of the greatest differentiators, cable manufacturers know that it is critical to reduce the size and weight of the cabling infrastructure. As shown in Figure 2, today's flexible ribbon solutions can provide 864-fiber cable in the similar space once required for 144-fiber loose tube or 288-fiber traditional ribbon solutions. With these new innovations, 100 to 200 percent density increase in the same space can be reached. This translates to less used space in ladder racking, less weight for easier handling and installation, and smaller holes and conduits between data center halls and campus buildings due to the smaller diameter cable. Without these breakthroughs in optical fiber and cable manufacturing, the demands of tomorrow's data center could not be achieved.

New flexible ribbon designs, the technical enabler of 200 μm singlemode fiber (SMF), has allowed for the significant breakthrough in density-per-square-inch. The 200 μm fiber provides over a 35 percent reduction in overall fiber cross-section. This translates to the ability to support higher density, less overall weight, and a smaller bend radius that supports an overall easier installation. Choosing a ribbon solution that maintains a 250 μm pitch enables backward compatibility to existing infrastructures built on 250 μm SMF with 250 μm pitch (Figure 3).

This is most important within the data center facilities where the need to maximize connections is pitted against the need to minimize overall infrastructure size. The smaller diameter cabling translates directly into

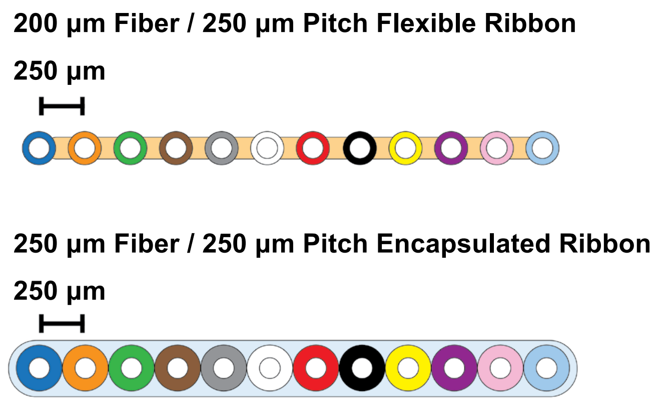

FIGURE 3: Comparison of 200 μm and 250 μm ribbon fiber.

CABLE TYPE	FIBER COUNT					
	144F	288F	432F	864F	1,728F	3,456F
Loose Tube Cable 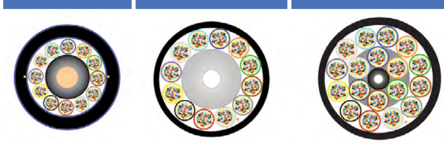	16.0 mm	18.9 mm	21.0 mm	—	—	—
Ribbon Loose Tube Cable 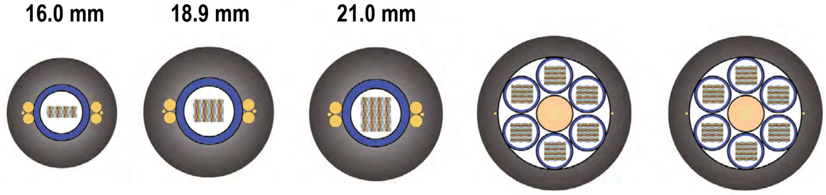	13.9 mm	19.8 mm	19.8mm	25.1 mm	35.4 mm	—
Flexible Ribbon Cable Example 	10.5 mm	12.0 mm	13.5 mm	17.5 mm	23.0 mm	30.0 mm

FIGURE 2: Cross-section comparison of different cable technologies.

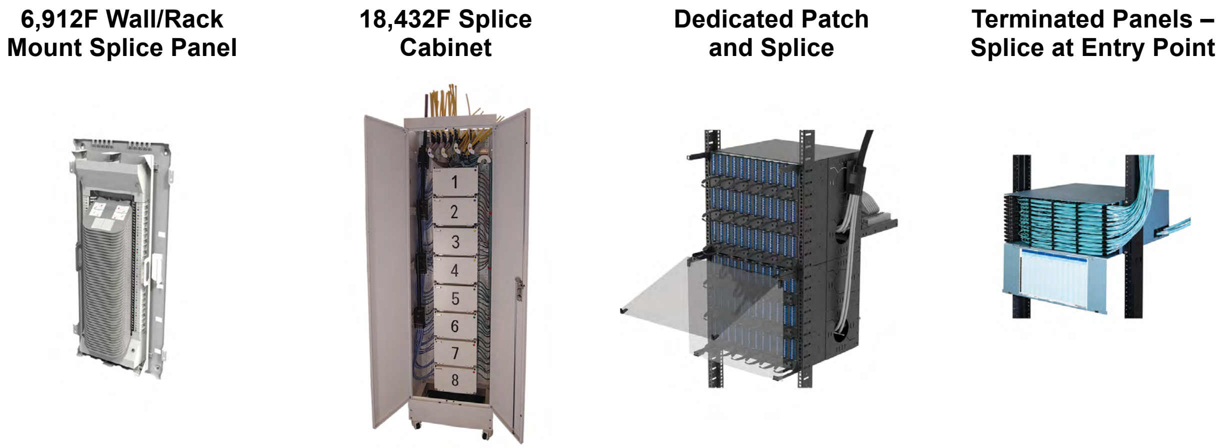

FIGURE 4: High-density splice patch cabinets and panels.

ease-of-installation, and its configurability makes today's flexible ribbon solutions a critical foundational piece to the high-speed highways that comprise the floor of a modern data center.

Inside the data center, managing this level of fiber can be a daunting challenge; it starts at the entry point and the transition from the outside fiber plant. Multiple cables, each with thousands of fibers, need to be patched or spliced. Pre-terminated cables allow for a faster overall installation, but they are not always an option. Often, the duct work leading up to the facility does not support the ability to pull through pre-terminated cabling. Therefore, another solution is needed. One approach is through high-density splice enclosures and cabinets (Figure 4).

High-Density Connectivity and Saving Data Center Space

Achieving a high level of density in connectivity is not trivial and manufacturers, who specialize in this space, spend significant engineering investments focused on innovating how best to support not only routing and containment of thousands of fibers, but also in easing the dressing and preparation of the fiber in the splice trays. When both the fiber technology advancements referenced earlier are brought together with companion high-density fiber management platforms, a true end-to-end solution emerges. These connectivity solutions, which are typically MPO or duplex-LC depending on the client's architecture and preference, provide a bridge from the outside fiber plant to the inside trunks. This is often done through patching frames. If not working with pre-terminated cabling, then field-installable splice-on/fuse-on connectors can provide a fast and convenient way to achieve factory-level connectivity (Figure 5).

The primary currency inside the data center is floor space. A lead network engineer with one of the largest hyperscalers said "All I care about is floor space and how I can extract the maximum amount of business per-square-

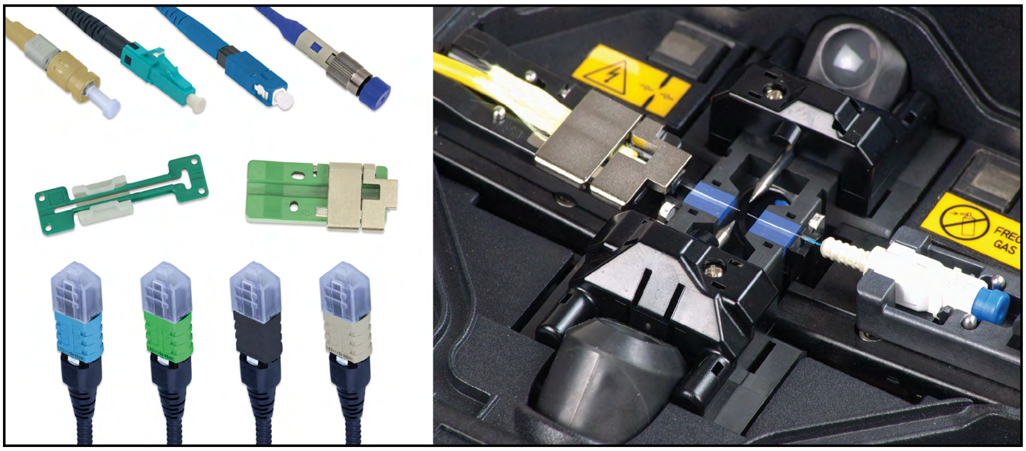

FIGURE 5: The field installable splice-on connectors have pre-polished ferrules that eliminate the need for polishing, adhesives, and crimping in the field, which minimizes the potential for operator error and expensive connector scrap, while providing fast and convenient factory-level connectivity.

Enterprise		Hyperscale	
1.	1 GbE to 10 GbE to 40 GbE	2.	25 GbE to 100 GbE and beyond
3.	Up to 150 m	4.	500 m to 2 km
5. 6.	Multimode meets most existing needs (generally up to 40 GbE) Singlemode may be used to provide future proofing	7.	Singlemode meets the current needs and can meet future requirements
8. 9.	Parallel optics to meet increased bit rates MPO connectors (pre-terminated or field-installable)	10. 11.	Increase serial speed to 100 GbE+ and parallel speed to 1 TbE MPO connectors (pre-terminated or field-installable)
12.	Shortwave wavelength division multiplexing (SWDM) with OM5 presents new growth path (extend to 100 GbE)	13. 14.	Course and dense wavelength division multiplexing (CWDM and DWDM) Duplex LCs and emerging nano connectors

TABLE 2: Traditional enterprise data center versus hyperscale infrastructure comparison.

inch." This is why being innovative in cable and fiber management design is so critical. Furthermore, it is important to think about not only revenue generation on Day 1 but Day 1001. Technology does not stand still. The data center floor footprint needs to support not only today's technology, but also tomorrow's. Understanding the direction clients envision for the future of their networks is of the utmost importance. Today a client's network may be at 10 GbE with plans to one day move to 40 GbE (most general enterprises) or it may be at 40 GbE and driving forward through 100 GbE and 400 GbE (the land of the hyperscaler).

Table 2 provides a comparative view of some of the key attributes defining the infrastructure in a typical enterprise, hyperscale or large colocation provider.

Providing a solution to achieve long-term success requires the selection of the right technology and a solution that provides forward mobility. The first step is choosing the right structured cabling approach—SMF or multimode (MMF). As table 2 shows, MMF can support 40 GbE; it may even support 100 GbE especially if using OM5, but then the infrastructure will be tapped out. Singlemode optical fiber has the greatest forward readiness, and for clients looking to leverage their fiber plant through multiple technologies, it is the best solution. The next step is choosing a fiber management solution that provides both flexibility and scalability. High-density connectivity platform solutions on the market today are designed to address this challenge through modularity.

Modular Solutions

Providing a rack-mounted, modular solution allows for easy reconfigurability. Platforms such as this provide an ability to evolve the floor space and drive continued

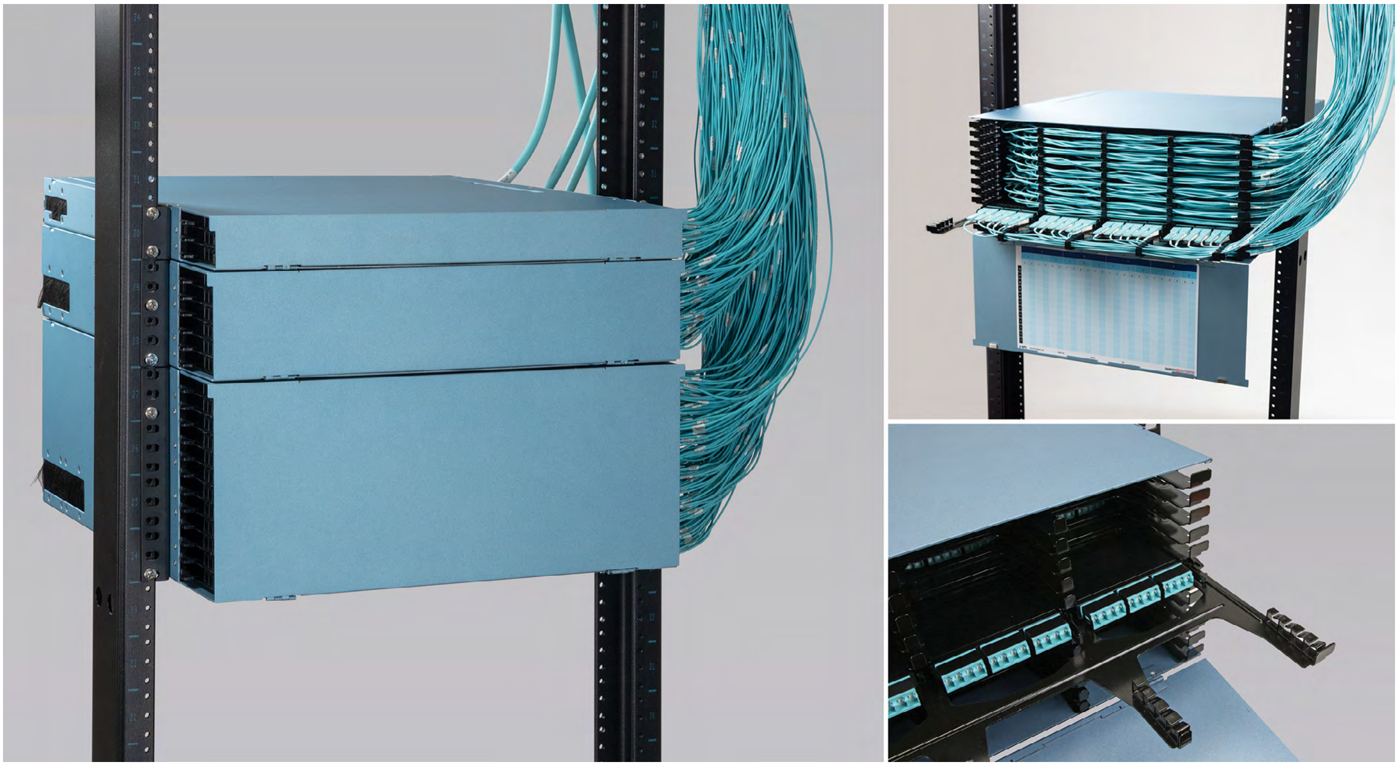

FIGURE 6: High-density connectivity platform.

economic benefit. The different module options are key in supporting a forward path. Along with the standard splice and MPO-to-LC fanout cassettes, there is need for conversion and patching solutions that enable easy migration through new technologies (i.e., 40 GbE to 100 GbE) while enabling use of the existing cabling infrastructure (Figures 6 and 7).

It is crucial to understand what the roadmap is of any given fiber management provider. The benefit of a modular solution is that new modules can be developed to support new and emerging technologies or architectures. One example is the nano connector from various companies targeting primarily the 400 GbE and beyond architectures. These reduced-size connectors provide a greater level of density than traditional LCs. Designed around a ceramic ferrule technology, they do not experience the same connection repeatability challenges often observed with MPOs.

It is important that designers and installers understand that some of the new nano connector solutions utilizing adapters, certain patch panels, fanouts and other non-direct LC and MPO connectivity may not comply with certain industry standards. For example,

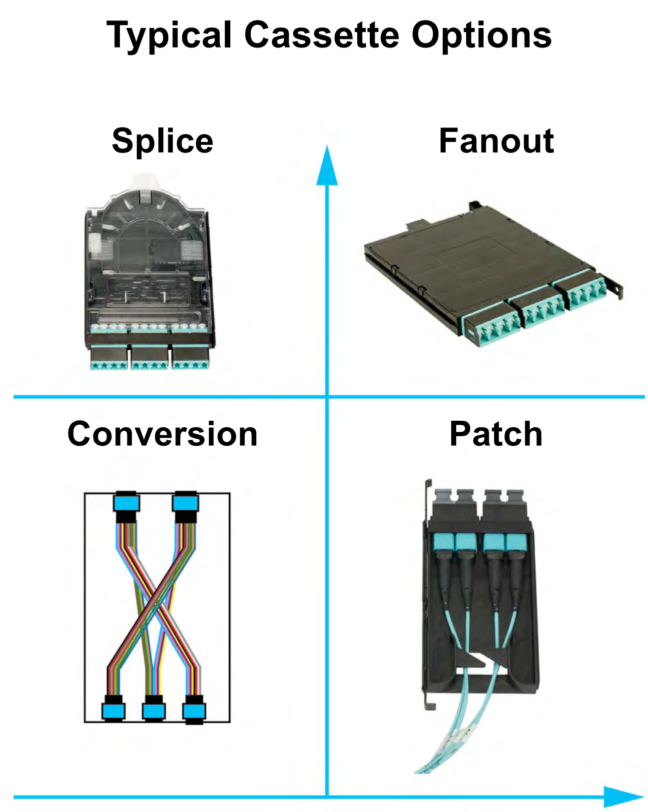

FIGURE 7: Cassette options.

ANSI/TIA-942-B specifies the LC connector for one or two optical fibers and MPO for more than two fibers. ISO/IEC 11801-5 and CENELEC EN 50173-5 specify the same optical fiber connectors at the equipment outlet and external network interface (LC and MPO). Unlike TIA-942-B, they permit the use of other connectors that meet the performance requirements of the standards in other areas, such as in main, intermediate and zone distributors as well as the local distribution point.

Another example is the use of course wavelength division multiplexing (CWDM) or dense wavelength division multiplexing (DWDM) within the data center. Both have been used in the backhaul networks for years and are often used in delivering connectivity to the data center from the carrier's network. Course wavelength division multiplexing and DWDM allow for data to be communicated through multiple wavelengths on the same fiber, and then optical filtering is used to break out into individual channels. The connectivity solution roadmap needs to keep these in mind and provide long-term support.

Fusion Splicing and Ribbon Fiber

Delivering the promise of tomorrow requires millions of low loss connections. A single hyperscale data center can easily contain over 100,000 servers and 500 to 1 million connectors. This does not account for the number of splices that could be in the order of tens of thousands. If pre-terminated trunks are not being used, then the number of splices can easily expand by an order of magnitude. Dealing with tens to hundreds of thousands of splices is no easy challenge. This level of splice density can be equated to time-spent, making true the old adage that "time equals money."

Ribbon fiber, especially a flexible gel-free ribbon product, can greatly help with this challenge. Mass fusion splicing of no-gel ribbon has been shown to reduce the overall splice time of a job by up to 87 percent (Figure 8). However, if these splices are not high quality and low loss, then any time benefit is lost by having to rework the connection. With loss budgets approaching 2 dB between any two active devices within the data center, the need to achieve splice losses of a hundredth of dB is critical. Today's leading splicers can achieve an average optical fiber splice loss of 0.03 dB. Typically, the challenge faced by low-end splicers is their ability to maintain a consistent fusion arc, especially across thousands of splices.

Having a quality splicer alone does not guarantee success. Technicians and installers also need to follow good workmanship techniques. It begins with careful preparation of the optical fiber. Ribbon stripping and cleaning the fiber is extremely important. Next, achieving a good cleave is critical for a low-loss connection. Following a good technique requires a sharp cleaver blade. Ever try cutting a tomato with a dull knife? Picture that happening to the fiber if using a dull or worn cleaver blade. Except in this case, the technician or installer is cutting a thousand tomatoes. The fiber splicing process can get messy very quickly. Keeping track of the blade rotation is also crucial for consistent low-loss results. Field data discloses how the splices at the start of a job achieve a near-zero loss, and by the end of the job they can approach 0.1 dB. As previously noted, in a connection path where the total loss budget is nearing 2 dB, this becomes a significant factor. It is for this reason that some splicer manufacturers have integrated intelligence into their splicer and cleaver families, allowing them to communicate via Bluetooth, for example. Manufacturers have also incorporated a feature that auto-advances the cleaver blade after a certain number of splices or when the loss level begins to move above a pre-set threshold.

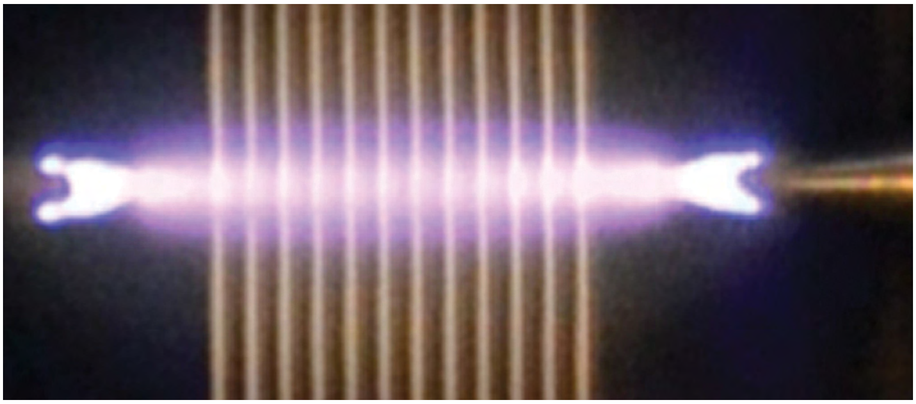

FIGURE 8: Mass fusion splicing of 12-fiber ribbon.

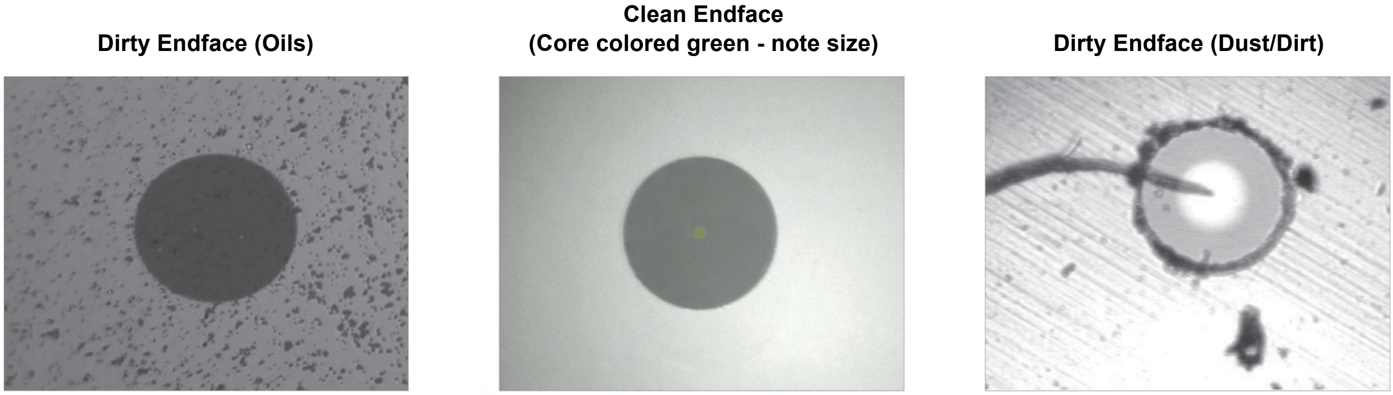

FIGURE 9: Comparison of most common dirt and debris found on connector endfaces.

Connector Cleaning and Best Practices

It can be argued that the modern data center relies on a connector or rather a million connectors. They might be factory installed or they might be spliced on. Either way, the technician or installer should worry about whether the connectors are clean and undamaged. It is reasonable to expect that factory-installed connectors are undamaged but expecting them to be clean is a different situation. As one industry expert said, “Dust caps are aptly named… they hold dust.” If the connector is not clean, then it is not going to deliver the performance required in today’s high-speed networks. Have you ever tried driving 100 mph with a dirty windshield? It does not end well. Trying to drive 100 Gb through a dirty connector is no different. It is estimated that over 80 percent of network turn-up failures are due to dirty or damaged connectors.

What is worse is that even if the dirt or damage did not impact the connection on Day 1, perhaps due to 10 GbE transmission, there is the possibility of what may happen on Day 1001 when there is an upgrade to 25 GbE transmission. If the world were a perfect place with perfect conditions, product, and staff, technicians and installers would not need to clean, inspect, and test. Even in a tightly environmentally controlled data center, there are plenty of small particles, such as dust and pollen (Figure 9) that can block the light on an 8 to 10 μm optical core endface (Figure 10).

The IEC defines a clean connector as having no scratches and no defects (e.g., dirt, pollen) in the core. There cannot even be any scratches greater than 3 μm or defects greater than 5 μm in the cladding. It is important to remember that any piece of debris anywhere on

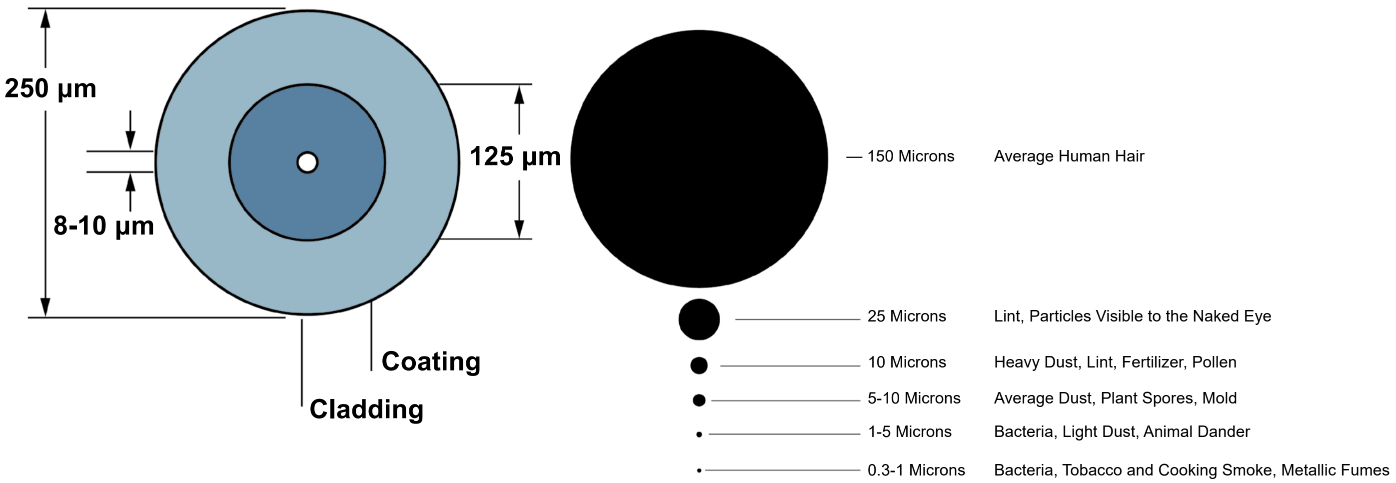

FIGURE 10: Fiber-endface cross section.

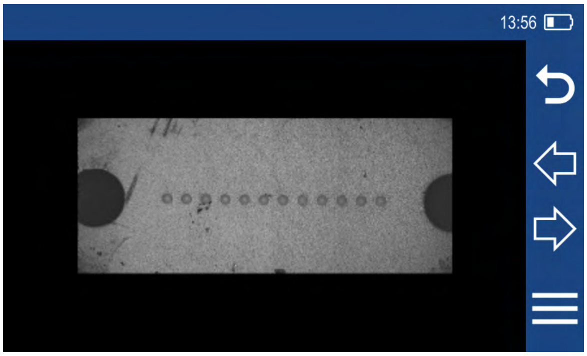

FIGURE 11: Image of MPO endface with dirt and debris.

the endface can easily migrate to the core during connectorization. What is worse is that the debris now contaminates both endfaces. It is similar to spreading a cold. Moreover, the problem actually grows exponentially more difficult when dealing with MPOs (Figure 11).

MPO connectors already have an inherent mating challenge for consistent low-loss interconnection. It is for this reason that some large hyperscale and data center providers have tried to move away from them as much as possible. Typically, this is difficult due to the tradeoff in density. For any connector type, it comes down to using the right tools and techniques to guarantee performance. Using the appropriate cleaning tools for the connector in question is highly important. As data rates continue to increase and loss budgets continue to decrease, it is important to inspect each connector to evaluate and confirm, either manually or through automatic pass/fail algorithms offered in many digital scopes, that there is no damage or debris. A recommended approach is to follow the flow as illustrated in Figure 12.

The inspect before you clean (IBYC) approach provides the greatest level of efficiency and time savings. As shown in Figure 12, it may be necessary to perform both wet and dry cleaning to remove all the debris. Technicians and installers should want a full toolkit of different cleaning products offered by various ICT companies. If after completing the flow process the endface cannot be suitably cleaned per IEC standards, then either a new cable or a splice-on field connector needs to be used.

The IEC defines a clean connector as having no scratches and no defects (e.g., dirt, pollen) in the core.

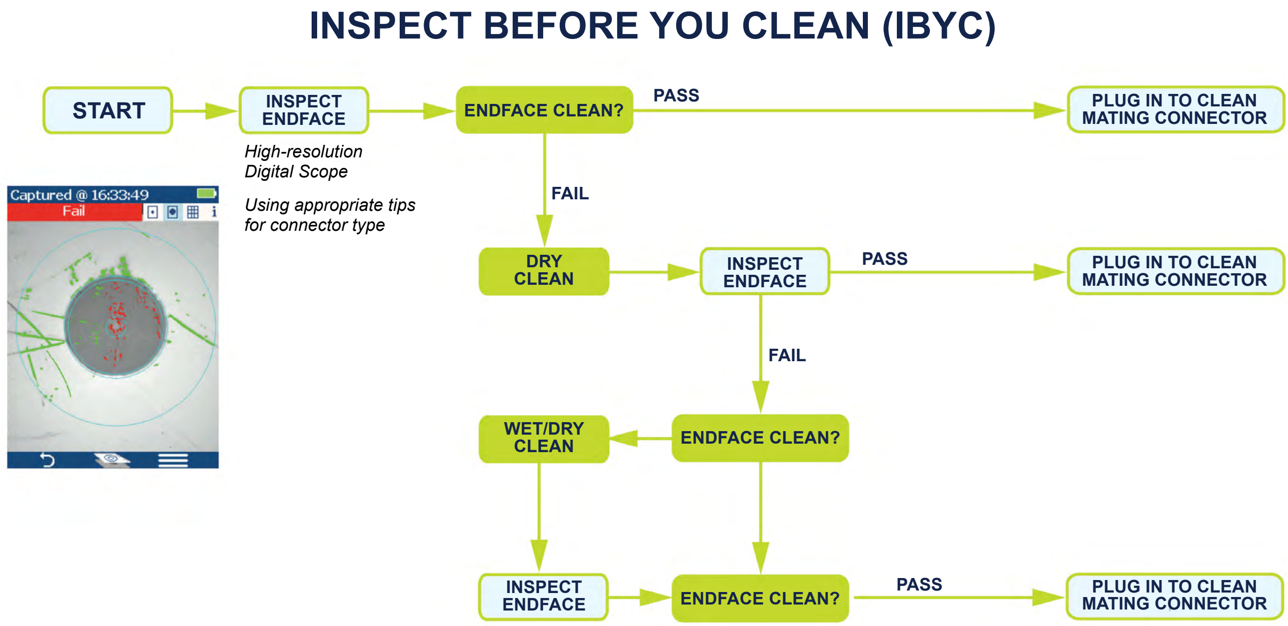

FIGURE 12: Inspect Before You Clean (IBYC) flow diagram.

For the densities required by the data center, optical fiber ribbon solutions are key because they enable easier overall cable management and the ability to perform mass fusion splicing.

CONCLUSION

Most ICT professionals agree that data is the life blood of the modern world. What is not always understood is that it all depends upon a vast infrastructure of optical fiber, optical connectors, and a myriad of panels and cabinets. High-density optical fiber cable stands as the foundation to enable today's drive for more data. Success requires implementing high-density optical cabling solutions successfully through new and innovative solutions that can handle the big problems of managing and connecting these high fiber counts. And, do not forget the tiny-but-significant challenge of keeping an 8 µm circle clean and verifying its cleanliness. Today's modern data center stands at the very center of an emerging reality. It is the very heart and the mind of tomorrow.

AUTHOR BIOGRAPHY: Seán Adam is vice president of market strategy and innovation at AFL. He holds a master's degree in electric engineering from Massachusetts Institute of Technology and has over 20 years of experience in the semiconductor and ICT industry with a focus on system-based solutions and architectures. Today, Seán is focused on developing and guiding AFL's overall growth strategy in the key markets it serves. He can be reached at Sean.Adam@aflglobal.com, via LinkedIn (www.linkedin.com/in/seanadam) or follow him on Twitter (@AFL_SeanAdam).

REFERENCES:

1. "New GSMA Study: 5G to Account for 15% of Global Mobile Industry by 2025 as 5G Network Launches Accelerate," *GSMA*, 25 February 2019.