Fiber Choices, New Standards Impact Network Performance

David Mazzarese
Technical Manager, Optical Fiber Systems Engineering

For network managers, installers and consultants, all eyes are on the data center. In businesses, educational and health facilities, and government organizations, this critical facility is at the hub of an explosion in bandwidth demand. The drivers behind this growth include the tremendous popularity of video and other high-bandwidth content on the Internet, the growing interest in videoconferencing, greater demand for data storage and record keeping, and the rise in supercomputing applications.

This trend is expected to continue, in part because of government data warehousing legislation and recommendations for the medical and financial industries, along with the need for redundancy to protect against catastrophic loss. As a result, data centers and storage area networks (SANs) are expected to see further upgrades to higher networking speeds of 40 and 100 Gigabits per second (Gb/s), depending on the application.

Optical fiber is the transmission medium of choice for these networks, due to its low loss and high bandwidth, small size, and low power consumption and heat generation. In this article, we'll review the fiber choices available to the data center user and discuss the evolution of standards that will determine which solutions are being defined by the industry as the most effective and cost-efficient.

Architectures and Protocols

Today's enterprise networks are increasingly taking advantage of 10 Gb/s-capable multimode fiber for backbone cabling in order to support 1 Gb/s-capable copper or fiber horizontal links. Traditional Hierarchal Star architecture is still used predominantly, but there are increased deployments of Fiber-to-the-Enclosure (FTTE) architecture that extend the high-performance capability of optical fiber much closer to the work station.

In the data center, where much of the information traveling over the local area network (LAN) is processed and stored, systems are becoming predominately optical fiber in order to keep up with the amount of information that needs to be managed.

Data centers typically consist of a SAN and a bank of servers that control the information traveling over the network. Data centers are connected to the LAN though an intranet and to the World Wide Web through the Internet. With more data being processed both internally and externally, the data center needs to be able to handle ever-increasing data rates.

Switches and servers in the data center typically use Ethernet as their communications protocol. Currently, 10 Gb/s (or "10G") is the fastest Ethernet speed that has been standardized (IEEE 802.3ae for fiber, published in 2002, and IEEE 802.3an for copper, published in 2006). However, IEEE is already working on standards for the next Ethernet speeds, 40G and 100G. These IEEE standards identify transceiver port types, and the requirements and characteristics of the physical layer.

In the SAN portion of the data center, Fibre Channel is the predominant protocol used. Heavily focused on optical fiber, Fibre Channel uses "Base2" speeds, doubling with each new generation (2GFC, 4GFC, 8GFC, etc.). Current efforts are focusing on 16GFC for the next Base2 speed.

Fibre Channel also uses a "Base10" protocol for inter-switch links and core connections. 10GFC was published on the heels of 10G and includes virtually the same 850-nm serial VCSEL

Evolution of Ethernet Speeds

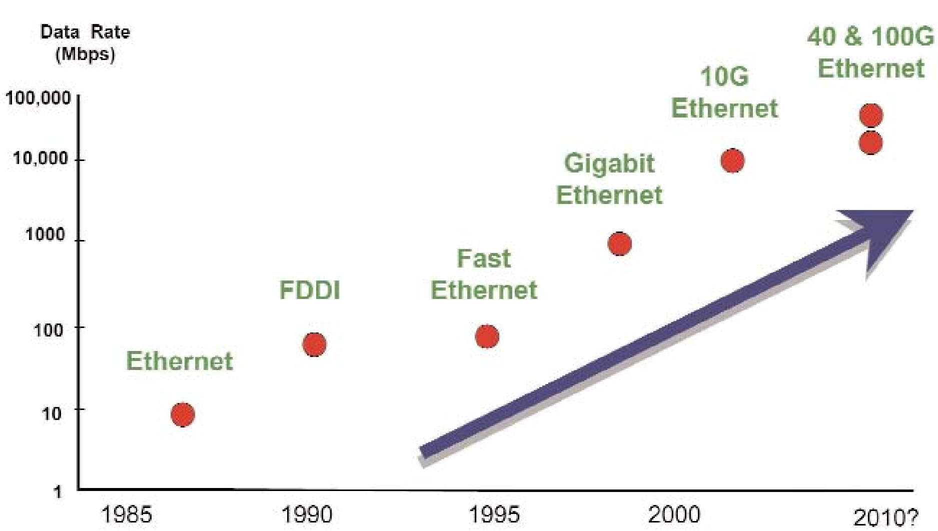

solution for 10 Gb/s up to 300 meters on 50 micron (µm) laser-optimized multimode fiber (also called OM3 fiber). Looking ahead, Fibre Channel is working on 20GFC with sights set on 40GFC as the next Base10 speed.

Multimode Fiber Offers Better Performance, Lower Costs

Several transmission media are available for use in the data center. These include various performance grades, or "categories," of copper cabling, and different types and performance grades of optical fiber.

Copper cabling has long been considered the least expensive option for data center applications, but its performance is limited in terms of transmission capacity and reach. For example, looking ahead at 40G and 100G transmission speeds, it is expected that copper will only be able to handle these speeds for very short distances, on the order of 10 meters or so. And it is too early to tell what makeup or type of copper cable will be necessary for these speeds. Historically, as transmission speeds increase, copper-based systems become more complex and costly.

On the optical fiber side, users have a choice between single-mode and multimode fiber. Single-mode fiber has very high bandwidth that can be transmitted long distances, but the opto-electronics required to do so are quite a bit more expensive than multimode (on the order of 25 - 30% higher). Even if you only need to go a few hundred meters, as with data centers, you would still need the more expensive optics if you used single-mode.

There are two types of multimode fiber - 62.5 µm and 50 µm, so named because of their core sizes - and various performance grades, listed here in increasing order of reach and performance capability: OM1 62.5 µm, and OM2, OM3, and soon to be OM4 50 µm.

Again using the 40G and 100G example, you will need optical fiber to transmit greater than 10 meters. Fortunately, distances of 100 - 200 meters or more are expected to be achievable using existing, standards-based OM3 multimode fiber (also known as laser-optimized 50 µm) and soon-to-be-standardized OM4 multimode fiber (extended-length laser-optimized 50 µm).

Why are the optics more expensive for single-mode than multimode? Two factors come into play - the wavelength of operation and, more significantly, the size of the fiber cores where the light is carried. The material used for the laser to achieve long wavelength (1310 nm, 1550 nm) transmission is more expensive than that for 850 nm short wavelength lasers. But more importantly, the transceivers used with single-mode fibers require significantly tighter alignment tolerances in order to couple the light into its tiny (9 µm) core. Not only is high-precision transceiver packaging required, but also tighter tolerance connectors and careful cable installation and termination practices are necessary. All this adds considerable cost as compared with multimode fiber for data center applications.

So for shorter reach premises applications like a data center, multimode fiber can easily provide the needed bandwidth (supporting up to 10 Gb/s or more serially and 40 and 100 Gb/s in parallel arrays) well into the future at much lower expense than single-mode fiber.

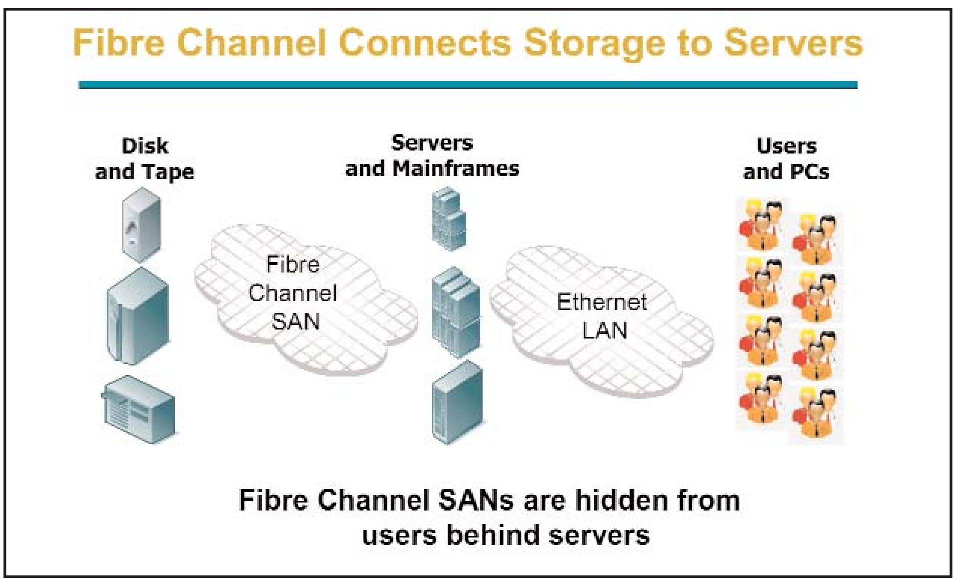

Courtesy of Terry Cobb, FibreChannel liaison to IEC

DMD-Controlled Fiber Helps Ensure Performance

Today, approximately 70 percent of the multimode fiber installed in the data center is OM3 or OM4 fiber. These laser optimized fibers, designed for 850 nm transmission using VCSELs as a light source, all feature a Differential Mode Delay (DMD)-controlled core that helps ensure 10 Gb/s support with low-cost 850 nm serial applications up to their rated distances. Even though these fibers are intended for high performance applications, they can still support 1 Gb/s operation, and their 50 µmcore size couples sufficient power from LED sources to support legacy protocols like Token Ring, FDDI, Ethernet and Fast Ethernet, and slower Fibre Channel speeds for virtually all in-building networks.

OM3 is the most widely deployed laser-optimized multimode fiber, providing 10 Gb/s transmission with low cost 850 nm serial applications for distances up to 300 meters. For longer distances (large building backbones, medium-length campus backbones, etc.) and more sensitive power budget applications (e.g., data center equipment interconnects) OM4 fiber with specifications that are significantly tighter than the current standards for OM3 are often deployed.

OM4 fiber, which can support 10 Gb/s Ethernet, Fibre Channel, and OIF applications to 550 meters using the same low-cost 850 nm VCSELs, is expected to become standardized in the industry through work currently being conducted in several international bodies including the TIA TR42.12 and IEC SC86A WG1. The key to the performance of these fibers is a manufacturing process that produces a fiber with almost no DMD and 4,700 MHz-km of effective modal bandwidth (EMB), more than double the IEEE requirement for 10 Gb/s 300 meter support.

Cassette-Based Solutions

To better manage the growth and increasing number of ports in a data center, pre-terminated multi-fiber trunk cables and multi-point optical (MPO) connectors are being used. For example, this will allow 12 fibers to be terminated with one mated pair of connectors. These "pre-term" assemblies provide ease of installation, space savings, and greatly simplify the connectivity portion of the network.

In one commonly used architecture, 12 fiber cables with MPO connectors are run between cassettes that then fan out to individual optical fiber ports. This architecture simplifies installation but could result in more connections than usual in a given optical link between the transmitter and receiver. Furthermore, multi-fiber MPO connectors typically exhibit higher connection, or insertion, loss than single-fiber connectors.

In these cases, multimode fiber again is the better choice for transmission media compared to single-mode. First, the larger core of a multimode fiber makes it easier to align their cores at a connection point, making them less sensitive to connection loss. Second, using a higher bandwidth fiber such as OM4 over a distance less than what it is rated for (typically 550 meters at 10G) provides additional channel insertion loss (ChIL) margin, or "headroom," to accommodate the additional, higher loss connectors.

Finally, the additional headroom can translate to more safety margin, providing additional immunity from installation challenges (cable routing, termination, etc.), link degradation from moves, adds and changes, or from aging of electronics.

Power Consumption and Cooling Considerations

One of the greatest challenges in today's data centers is minimizing costs associated with power consumption and cooling. Higher power consumption means increased energy costs and more heat dissipation. This requires more cooling, which adds even more cost. The comparatively low power requirements of optical fiber networks give them a big advantage over copper.

Example: A 10G BASE-T transceiver in a copper system uses about 6 watts of power. The comparable 10G BASE-SR optical transceiver uses less than 1 watt to transmit the same signal. The result: each optical connection saves about 5 watts of power. Data centers vary in size, but if we assume 10,000 connections at 5 watts each, that's 50 kW less power -- significant savings from using less power-hungry optical technology.

And there's more: The power used by these transmitters is dissipated as heat, which must then be removed from the room in order to keep the electronics cool. A recent EPA study shows that each watt required to operate data center equipment will require another watt for cooling and power delivery. This means that an additional 50 kW of power is required for a 10,000-port data center.

Finally, consider cable size and its effect on airflow. A duplex 2.0 mm fiber jumper occupies approximately 8 mm^2. A single Cat 6A cable is approximately 9 mm in diameter, and occupies approximately 64 mm^2. If 40 servers are mounted in a rack and each has one or more network connections, it's easy to understand how copper cabling causes congestion. Fiber cabling greatly reduces that congestion while allowing better airflow and cooling.

Looking Ahead to Higher Speeds

As mentioned previously, the IEEE is currently developing new standards for higher speed transmission - 40 Gb/s and 100 Gb/s - in data centers and other high performance computing (HPC) applications. The IEEE 802.3ba Task Force is working to develop these 40G and 100G standards simultaneously (40 Gb/s will support the server market, while 100 Gb/s is needed for core switching and routing, network aggregation, and high performance computing).

For shorter reach data center and equipment interconnects, IEEE 802.3ba is focusing on a Physical Medium Dependant (PMD) solution that takes advantage of parallel optics technology (which is already being used in current platforms such as InfiniBand), thereby helping to keep costs as low as possible. Parallel optics entails simultaneous transmission of one 10 Gb/s signal on each of 4 or 10 fibers (for 40G and 100G, respectively). Arrayed transceivers using 4, or 10, VCSELs and detectors, as appropriate, will aggregate each 10 Gb/s signal.

To further balance cost with performance, the task force is working to leverage proven technology, media and network management practices. In fact, they will likely relax component performance specifications in some cases in order to help reduce overall cost. An example is the VCSEL light sources for shorter reach applications using multimode fiber. The 802.3ba is considering a relaxation of the spectral width of these sources from 0.45 nm (the current 10GbE requirement) to 0.65 nm. This limits the distance of such a link (due to chromatic dispersion effects) to 100 m using OM3 fiber.

For some data centers and other applications that may require support beyond 100 m, an ad hoc group within IEEE is studying how this could be accomplished in a cost-effective manner. It might use better performing transceivers, or an OM4 grade of fiber, or a combination of the two.

Handling "Delay Skew"

Any discussion of a parallel transmission approach will include a topic called "delay skew," which is being addressed by the IEEE task force. Delay skew is the difference in signal arrival time from one lane (or fiber) to the next. Skew can be affected by differences in the physical lengths of each fiber within the cable, and by any difference in speed that the light signal travels down one fiber compared to adjacent fibers.

Delay skew will not be a hindrance to parallel transmission over any cable design as it will be compensated for effectively within the transceiver circuitry. In fact, proven techniques for skew compensation in copper cabling and other parallel optics applications are already well established.

The new standard for 40 and 100 Gb/s transmission will include procedures to compensate for skew, ensuring that industry-recognized cable designs such as loose tube, tight buffer, and ribbon cable all can be accommodated, and the full range of current, industry-standard OM3 multimode fibers can be used.

Exceeding the Standards for Higher Performance

Data center designers are likely to agree that the lowest cost solution for 10 Gb/s deployment will contain a significant amount of OM3 fiber, and as the systems migrate to higher speeds of 40 to 100 Gb/s, laser optimized 50/125 multimode fiber can provide the lowest cost and most reliable solution as compared to copper cable or single-mode fiber.

Once settled on the fiber type for their data center network, the user must be sure that the optical fiber products they specify can provide the performance and reliability needed. This is especially critical in 10G applications at 850 nm, since loss budgets for these systems are lower than previous applications. As touched upon earlier, you may want your network to have extra power "headroom" to accommodate additional connections and higher loss connectors, and to improve overall reliability.

There are two ways to achieve greater power headroom (also known as power margin). First, by reducing ChIL, the end-to-end loss resulting from all connections and splices in the link, plus the attenuation of the cable itself; second, by using a higher bandwidth fiber to reduce Inter-symbol Interference (ISI), which occurs when bits of data run together.

Because network downtime can be very expensive, reliability is a key consideration for high performance networks. For greater flexibility in network design and, ultimately, greater reliability, follow these strategies:

- specify lower loss cables and connectors, which provide more power margin
- specify a fiber rated for a longer distance than it will be used for
- don't assume that all products that meet a particular standard are equal; it's possible to find higher performing products that exceed the standards.

All of this is especially true in demanding data center applications. The most cost-effective solution is OM3 fibers that have been designed and manufactured specifically for laser transmission and have performance characteristics that exceed the standards. They are available in various performance grades and should feature a DMD-controlled core that helps ensure 10 Gb/s support with low-cost 850 nm serial applications up to their rated distances.

About the Author

David Mazzarese is the Technical Manager of Fiber Systems Engineering for OFS, supporting the entire optical path. Prior to this, he spent more than 10 years leading the development of various multimode, single-mode and specialty fiber products for the company. He has also served as technical manager of perform and measurements development and engineering at OFS' Sturbridge, Mass. facility. He holds a PhD in Chemical Engineering and an MS in Electrical Engineering from the University of Massachusetts at Amherst.

For additional information, visit our website at www.ofsoptics.com/ofs-fiber or call 1-888-fiber-help. For regional assistance, contact:

North America
Telephone: 508-347-8590
Toll Free: 800-799-7732
Fax: 508-347-1211
E-mail: fibersalesnar@ofsoptics.com

Caribbean, Latin America
Telephone: 508-347-8590
Fax: 508-347-1211
E-mail: fibersalescala@ofsoptics.com

Europe, Middle East, Africa
Telephone: +45-43 48 3736
Fax: +45 4348 3444
E-mail: fibersalesemea@ofsoptics.com

Asia Pacific
Telephone: +852 2836 7102
Fax: +852 2836 7101
E-mail: fibersalesap@ofsoptics.com

Japan
Telephone: +81-3-3286-3424
Fax: +81-3-3286-3708 or 3190
E-mail: fibersalesjapan@ofsoptics.com

China
Telephone: +86 10 6505 3660
Fax: +86 10 65059515
E-mail: fibersaleschina@ofsoptics.com

ofs
A Furukawa Company