

advancing information transport systems

BICSI news

PRESIDENT'S MESSAGE	3
EXECUTIVE DIRECTOR MESSAGE	4
BICSI UPDATE	38-40
COURSE SCHEDULE	44-45
STANDARDS REPORT	46

Volume 29, Number 5

The Evolution of Structured Cabling Standards >> 18

Optical Fiber in the Data Center >> 24

The Greening of the Data Center >> 30

Improve Your Business Writing >> 35



Technology Infrastructure for K-12 Schools, Part 2

Telecommunications Spaces, Pathways and Outlets

>> 6

Bicsi[®]

OPTICAL FIBER in the DATA CENTER

Specification of lower loss cables and connectors and optical fiber

rated for longer distances offer design flexibility. **BY DAVID MAZZARESE**

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 recordkeeping, 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 generation of heat. In this article, optical fiber choices available to the data center user will be reviewed and

the evolution of standards that will determine which solutions are being defined by the industry as the most effective and cost-efficient will be discussed.

Network Architectures and Protocols

Today's enterprise networks are increasingly taking advantage of 10 Gb/s-capable multimode optical fiber for backbone cabling in order to support 1 Gb/s-capable copper or optical fiber horizontal links. Traditional hierarchical 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 workstation.

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 through 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 optical fiber, published in 2002, and IEEE 802.3an for copper, published in 2006). However, the Institute of Electrical and Electronics Engineers Inc.® (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 vertical cavity surface emitting laser (VCSEL) solution for

Fibre Channel Connects Storage to Servers



Fibre Channel SANs are hidden from users behind servers

April 2008

IEC SC86A WG1

BE - 23

Courtesy of Terry Cobb, Fibre Channel liaison to the IEC

10 Gb/s up to 300 meters (m [984 feet (ft)]) on 50 μm laser-optimized multimode optical fiber (also called OM3 optical fiber). Looking ahead, Fibre Channel is working on 20GFC and already has sights set on 40GFC as the next Base10 speed.

Multimode Optical 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 m (33 ft) or so. 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 singlemode and multimode optical fiber. Singlemode optical fiber has very high bandwidth that can be transmitted long distances, but the optoelectronics required to do so are quite a bit more expensive than multimode (on the order of 25–30 percent higher). Even if you only need to go a few hundred meters, as with data centers, you still need the more expensive optical fibers if you were to use singlemode.

There are two types of multimode optical fiber—62.5 μm and 50 μm , 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 40 and 100G example, optical fiber is needed to transmit greater than 10 m (33 ft). Fortunately, distances of 100–200 m (328–656 ft) or more are expected to be achievable using existing, standards-based OM3 multimode optical fiber (also known as laser-optimized 50 μm) and soon-to-be-standardized OM4 multimode

optical fiber (extended-length laser-optimized 50 μm).

Why is optical fiber more expensive for singlemode than multimode? Two factors come into play—the wavelength of operation and, more significantly, the size of the optical 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 singlemode optical fibers require significantly tighter alignment tolerances in order to couple, or capture, the light into its tiny (9 micron) 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 optical fiber for data center applications.

So for shorter reach premises applications like a data center, multimode optical 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 singlemode optical fiber.

Differential Mode Delay (DMD)-Controlled Optical Fiber Helps Ensure Performance

Today, approximately 70 percent of the multimode optical fiber installed in the data center is OM3 or OM4 optical fiber. These laser optimized optical 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 optical fibers are intended for high performance applications, they can still support 1 Gb/s operation, and

their 50-micron core size couples sufficient power from light emitting diode (LED) sources to support legacy protocols like Token Ring, fiber distributed data interface (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 optical fiber, providing 10 Gb/s transmission with low cost 850 nm serial applications for distances up to 300 m (984 ft). For longer distances (e.g., large building backbones, medium-length campus backbones) and more sensitive power budget applications (e.g., data center equipment interconnects), OM4 optical fiber with specifications that are significantly tighter than the current standards for OM3 are often deployed.

OM4 optical fiber, which can support 10 Gb/s Ethernet, Fibre Channel, and OIF applications to 550 m (1804 ft) 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 optical fibers is a manufacturing process that produces an optical fiber with almost no DMD and 4700 MHz-km of effective modal bandwidth (EMB), more than double the IEEE® requirement for 10 Gb/s 300 m (984 ft) 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 optical 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 optical 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 optical fiber again is the better choice for transmission media compared to singlemode. First, the larger core of a multimode optical fiber makes it easier to align their cores at a connection point, making them less sensitive to connection loss. Second, using a higher bandwidth optical fiber such as OM4 over a distance less than what it is rated for (typically 550 m [1804 ft] 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 (e.g., cable routing, termination), link degradation from moves, adds, or changes (MACs), or from aging of electronics.

Power Consumption and Cooling Considerations

One of the greatest challenges with today’s data centers is minimizing costs associated with power consumption and cooling. The more power that is consumed, the higher the cost and more heat generated. This requires more cooling, which adds even more cost. The comparatively low power requirements of optical networks give them a big advantage over copper.

For example, a 10G BASE-T transceiver in a copper system uses about 6 watts (W) of power. The comparable 10G BASE-SR optical transceiver uses less than 1 W to transmit the same signal. Thus, each optical connection saves about 5 W of power. Data centers vary in size, but if 10,000 connections at 5 W each were considered, that is 50 kW less power—significant savings from using less power-hungry optical technology.

Furthermore, 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. Typical air conditioning has an energy efficiency rating of around 10, meaning that it takes 10 W of power to remove one W of heat. Removing the 50 kW of power described above would require about 500 kW of energy! The energy costs of cooling are 10 times the energy of operating the transceivers themselves. That is a total of 550 extra kW of power needed to operate a 10,000 port, copper-based data center.

Looking Ahead to Higher Speeds

As mentioned previously, 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 applications, network aggregation, and high performance computing).

For shorter reach data center and equipment interconnects, IEEE 802.3ba is focusing on a physical medium dependent (PMD) solution that takes advantage of parallel optical fiber technology (which

is already being used in current platforms such as InfiniBand), thereby helping to keep costs as low as possible. Parallel optical fibers entails simultaneous transmission of one 10 Gb/s signal on each of 4 or 10 optical 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 optical 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 (328 ft) using OM3 optical fiber.

For some data centers and other applications that may require support beyond 100 m (328 ft), an ad hoc group within IEEE® is studying how this could be accomplished in a cost-effective manner. It might make use of better performing transceivers, or of an OM4 grade of optical 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 optical fiber, to the next. Skew can be affected by differences in the physical lengths of each optical fiber within the cable, and by any difference in speed that the light signal travels down one optical fiber compared to adjacent optical 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 optical fiber applications are already well established.

The new standard for 40 and 100 Gb/s transmission will include procedures for compensating for skew, ensuring that industry-recognized cable designs such as loose tube, tight buffer, and ribbon cable all can be accommodated, and the full performance range of current, industry-standard OM3 multimode optical 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 optical fiber, and as the systems migrate to higher speeds of 40-100 Gb/s, laser optimized 50/125 multimode optical fiber can provide the lowest cost and most reliable solution as compared to copper cable or singlemode optical fiber.

Once settled on the optical 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 briefly mentioned earlier in this article, 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 optical fiber to reduce intersymbol 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 an optical fiber rated for a longer distance than what it will be used.
- Do not assume that all products that meet a particular standard are equal; it is 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 optical 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. ■



David Mazzaresse

David Mazzaresse is technical manager, optical fiber systems engineering, for OFS in Sturbridge, Mass.

For more information, contact David at dmazzaresse@ofsoptics.com or at +1 413.436.5052.