

Enabling Enterprise 10 Gigabit Ethernet Deployment with Long Reach Multimode Optics

Abstract


Multimode fiber (MMF) has been widely adopted and deployed as the medium to transport high-speed data in enterprise networks. As data rates increased to the gigabit range, the traditional copper medium supported only limited transmission distances. Multimode fibers emerged as the medium of choice capable to address the shortcomings of copper. With MMF, IT managers can deploy 1-Gbps links over a few hundred meters, interconnect buildings, and build scalable unified campus networks.

Nevertheless, as the bandwidth demand increases with the growing number of end user applications, network architects and IT managers are facing a new type of challenge: the need to scale their network to 10 Gbps. A significant portion of the MMF deployed for 1-Gbps links are legacy types such as Fiber Distributed Data Interface (FDDI)-grade and OM1 fibers, which were not readily suited for a smooth transition to 10 Gbps. The signal propagating along multimode fibers gets distorted. This distortion phenomenon, also known as modal dispersion, is emphasized with data rates increasing and with transmissions over older fiber types. Therefore, one major limitation of MMF resides in the diversity in fiber types currently existing in the installed base. There is a wide difference in their performance characteristics, and this makes it difficult, in many instances, to manage an upgrade to 10 Gbps without changing the fiber plant.

This resulted in the need for a robust and cost-effective 10 Gigabit Ethernet optical transceiver solution that can help upgrade campus and building backbone links regardless of the fiber type already installed. The IEEE 802.3.aq defined an interface that will cost-effectively allow an upgrade path to 10 Gbps without a change to the existing fiber plant. In this paper, we describe the performance and the capabilities enabled by this new device standardized as the 10GBASE-LRM (Long Reach Multimode).

State of the Art of Enterprise Network Connectivity

Gigabit Ethernet optical interfaces are currently the dominant type in enterprise networks. A transition to 10 Gigabit Ethernet is starting to take place as bandwidth demand increases and transceiver pricing decreases. The chart in Figure 1 describes the current and forecast mix of Gigabit Ethernet and 10 Gigabit Ethernet transceiver shipments. In 2011, the numbers will nearly be equal. The demand is promoted by high-volume applications such as data server connections, wiring closet backhaul, rapid deployment of video services, and the availability of cost-effective technologies such as 10GBASE-LRM.

Figure 1. Distribution of Gigabit Ethernet vs. 10 Gigabit Ethernet Interfaces

Source: Cisco internal analysis.

Multimode Fiber Types: An Overview

The main challenge in upgrading 1-Gbps links to 10-Gbps links is the variety of fiber types already deployed in the installed base. The main property of a multimode fiber is its capability to transmit a certain amount of information over a certain distance. This property is known as “modal bandwidth” and is expressed in MHz*km.

In general, an inverse linear relationship between the data rate and the maximum distance achievable is considered at a certain modal bandwidth. As an example, this means that an increase of the data rate from Gigabit Ethernet to 10 Gigabit Ethernet (which represents a factor of ten) will result in a tenfold reduction of achievable distance on the same fiber and under the same launch conditions. However, some complex nonlinear parameters need to be taken into account to calculate the exact reach possible at a certain data rate, and standard bodies such as IEEE define minimum transmission distance requirements for various combinations of data rate and fiber modal bandwidth. Please refer to Table 1 for a list of various fiber types defined by major international standard bodies.


Table 1. Various Types of Multimode Optical Fibers

Standard Body	Document	Notes
TIA—Telecommunications Industry Association	492AAAA	62.5 μm fibers with 160/500 MHz.km OFL BW
	492AAAB	50 μm fibers with 500/500 MHz.km OFL BW
	492AAAC	Laser-optimized 50 μm fibers with 2000 MHz.km EMB at 850 nm
IEC—International Electrotechnical Commission	60793-2-10	A1a.1 fiber—50 μm fibers with a range of OFL BW
		A1a.2 fiber—Laser-optimized 50 μm fibers with 2000 MHz.km EMB at 850 nm
		A1b fiber—62.5 μm fibers with a range of OFL BW
ISO—International Standards Organization	11801	OM1 fiber—200/500 MHz.km OFL BW (in practice OM1 fibers are 62.5 μm fibers)
		OM2 fiber—500/500 MHz.km OFL BW (in practice OM2 fibers are 50 μm fibers)
		OM3 fiber—Laser-optimized 50 μm fibers with 2000 MHz.km EMB at 850 nm

Another important parameter of an optical fiber is the size of its core, which is the cylindrical section made of glass that serves as the waveguide for optical information propagation. In enterprise networks, there are two main types of MMF: one with a core diameter of 62.5 μm , and the other one with a core diameter of 50 μm . When an optical signal is transmitted over multimode fibers, it spreads into multiple modal components as it enters the medium and travels along the path. These components need to be recombined in a timely manner at the other end of the transmission link in order to be recognized and processed by the detection subsystem converting the optical signal into an electrical signal. The larger the size of the core, the easier it is for modes to be spread, and the harder it is for them to be recombined.

Figure 2 describes examples of multimode and single-mode fibers, as well as the typical channel impulse responses showing how the pulse spreads depending on the fiber type.


Figure 2. Pulse Spreading Severity on Various Types of Optical Fibers



Fiber types known as OM1 or FDDI were the first ones to populate campus and building backbones, connecting servers to hubs and early-generation switches. These fibers have a 62.5 μm core with a typical modal bandwidth of 160 to 200MHz*km. As glass processing techniques improved, newer types of 50 μm MMF with higher modal bandwidth were easier to manufacture. Nevertheless, these new types of better performing fiber did not immediately replace the installed base, but continued to expand the enterprise network with the latest technology available. As a result, enterprise networks are currently populated with a mix of all types of MMF. Traditional OM1 and FDDI-grade fiber types still represent more than 50 percent of the new deployments and will be accounting for more than 40 percent of the installed base by 2010. This defines the urge for identifying a 10 Gigabit Ethernet transceiver solution that will be compatible with the existing infrastructure.


Figures 3 and 4 depict the current and forecast distributions of fiber types in campus and building backbones.

Figure 3. Installed Base: Worldwide Campus Backbones



Source: In Premises Optical Fiber Installed Base Analysis to 2010, Alan Flatman, April 2007.

Figure 4. Installed Base: Worldwide Building Backbones



Source: In Premises Optical Fiber Installed Base Analysis to 2010, Alan Flatman, April 2007.

Upgrading from Gigabit Ethernet to 10 Gigabit Ethernet: A Brief Comparison Between Existing Solutions

This section details the existing solutions for 10 Gbps over MMF and introduces the newly defined interface known as 10GBASE-LRM.

In 2002, and again in 2005, updates to the IEEE 802.3 standard were published with the introduction of 10GBASE-SR and 10GBASE-LX4 interfaces. These optical transceivers represent the current options that support 10-Gbps links over multimode fibers.

The 10GBASE-LX4 is a longwave parallel interface with built-in wavelength-division multiplexing (WDM) multiplexer and demultiplexer filters and a set of four distributed feedback (DFB) lasers and four receiving PIN diodes. This enables the transmission of four different fixed wavelengths simultaneously over the same fiber strand. This interface is designed to transmit data over up to 300 meters on multimode fibers. In contrast, the 10GBASE-SR is a very low-cost vertical-cavity surface-emitting laser (VCSEL)-based shortwave serial interface specifically designed to transmit data over strands of MMF varying from 26 meters to 300 meters. In both cases, the performance of the device and the reach achieved are dependant on the fiber medium type.

This drove the need for a new low-cost transceiver that could have the same performance over any type of MMF. In November 2006, the IEEE 802.3aq committee finalized and published an amendment to the standard and created the 10GBASE-LRM. As a quick comparison with the two other interfaces, the 10GBASE-LRM is a longwave serial interface that includes an electronic dispersion compensation (EDC) chip on the receiving end placed immediately after the receiver optical sub-assembly (ROSA). This enables the adaptive equalization of incoming modal dispersion and thus eliminates the dependency on fiber types. 10GBASE-LRM modules can transmit data over 220 meters on any type of MMF. Table 2 gives us an overview of the minimum reaches specified by IEEE 802.3 for the various interfaces.

Table 2. Operating Ranges of 10GBASE-LX4

Fiber Type	Modal bandwidth at 1300 nm (Minimum Overfilled Launch) (MHz km)	Minimum Range (Meters)
62.5 μm MMF	500 (OM1)	2 to 300
50 μm MMF	400	2 to 240
50 μm MMF	500 (OM2/OM3)	2 to 300
10 μm SMF	n/a	2 to 10,000

Table 3. Operating Ranges of 10GBASE-SR

Fiber Type	Modal bandwidth at 850 nm (Minimum Overfilled Launch) (MHz km)	Operating Range (Meters)
62.5 μm MMF	160	2 to 26
	200	2 to 33
50 μm MMF	400	2 to 66
	500	2 to 82
	2000	2 to 300

Table 4. Operating Ranges of 10GBASE-LRM

Multimode Fiber Type ¹	ISO/IEC 11801: 2002 Fiber Type	Operating Range (Meters)	Maximum Channel Insertion Loss (dB) ²
62.5 μm 160/500 ³		0.5 to 220	1.9
62.5 μm 200/500	OM1	0.5 to 220	1.9
50 μm 500/500	OM2	0.5 to 220	1.9
50 μm 400/400		0.5 to 100	1.7
50 μm 1500/500 ⁴	OM3	0.5 to 220	1.9

Source: IEEE 802.3-2005 and IEEE 802.3aq.

10GBASE-LRM Standard

10GBASE-LRM is poised to become the solution of choice to upgrade 1-Gbps links to 10 Gbps for reaches of up to 220 meters in campus and building backbones. Its standardization was finalized and published by IEEE in November 2006. In the process of producing this standard, the IEEE LRM Task Force took a statistical approach, modeling and testing a number of fibers to define transmitter and receiver parameters that would result in error-free transmission over 99 percent of 220-meter-long fibers deployed in the field.

Two New Parameters Defined by IEEE

As a result of IEEE LRM Task Force studies, two new parameters specifically relating the modeled fiber transmission channel are introduced to help ensure performance compliance on all fiber types: transmitter waveform dispersion penalty (TWDP) and comprehensive stressed receiver sensitivity (CSRS).


At a glance, TWDP is a measure of the deterministic dispersion penalty (that is, directly derived from intersymbol interference [ISI]) caused by a particular transmitter with reference to the emulated multimode fibers and a well-characterized receiver. Measuring TWDP involves capturing a transmitter waveform and processing it using a code to calculate the penalty of that waveform on a reference equalizer. CSRS, on the receiver end, is a measure of the robustness of the receiving subsystem (which includes the equalization chip, which will be described in more details further in this document). To test LRM receivers, IEEE defined a set of three different stressed optical signals known as precursor pulse, postcursor pulse, and split symmetric pulse. Theoretical predictions and calculations conducted by the study group show that these stressed pulses would be responsible for, respectively, 4.1dB, 4.2dB, and 3.9dB TWDP at the receiver, making the postcursor pulse the worst-case stressor and 4.2dB the maximum penalty observed in CSRS testing. It was also identified that 220 meters should be the minimum reach achieved on any type of fiber with a modal bandwidth of 500MHz*km in the 1300nm wavelength window. Knowing the inverse linear relationship between the data rate and the achievable distance, the effective modal bandwidth (EMBW) of worst-case fibers for a reach of 220 meters is equal to 2.272GHz. This is how IEEE defined a chart (Figure 5) showing the permitted area for a 10GBASE-LRM-compliant link.

¹ Each fiber type is identified by its core diameter followed by a pair of OFL bandwidth values separated by "/". The OFL bandwidth are in MHz km and are for 850 nm and 1300 nm respectively.

² Channel insertion loss includes cable attenuation and an allocation of 1.5 dB for connectors.

³ 160/500, 62.5 μm fiber is commonly referred to as "FDDI-grade" fiber.


⁴ The OM3 fiber specification includes the 850 nm laser launch bandwidth in addition to the OFL bandwidths.

Figure 5. TWDP vs. EMBW Permitted Area for 10GBASE-LRM Fiber Links


These new parameters, among others defined by IEEE, are designed to guarantee repeatable interoperability and link performance for 10GBASE-LRM. Thanks to their definition, the difference in modal bandwidth properties of multimode fibers is now well assimilated. Nevertheless, other dynamic random parameters such as temperature variations, movements of fiber cables, polarization state variations, and component aging can affect transmission performance by causing unpredictable fluctuations to signal distortion.

Electronic Dispersion Compensation: Adapting Continuously to Varying Conditions

Multiple factors cause gradual or rapid variations in transmission performance and add complexity to the dispersion phenomenon. As a reminder, the term modal dispersion in transmissions over MMF refers to the phenomenon where different components of a signal travel through the fiber medium at different speeds and therefore arrive at the receiving end at different times. This can be better understood with the alternate term "pulse spreading." A data stream consists of a rapid succession of pulses. Modal dispersion causes the pulses to spread and arrive at the receiving end with a wide shape. The detection subsystem needs to be robust enough to distinguish which components belong to which pulse, because received pulses overlap. Figure 6 shows the case of one pulse and two pulses transmitted over a fiber strand.

Figure 6. Overlapping of Spread Pulses in Data Transmissions over MMF

If we consider the random variations contributing to dispersion, there is a need for a robust and adaptive solution to compensate for these factors. Theoretically, this effect can only be removed in the optical domain and not after a conversion of the signal in the electrical domain, because phase information is lost during detection. However, the effect of dispersion can at least be mitigated in the electrical domain, after a photodetector such as a PIN diode has converted the optical signal in an electrical signal. These equalization techniques, known as electronic dispersion compensation (EDC), are in many ways more cost-effective than true optical dispersion compensation methods. Electronic filters are placed right behind the detection system, which traditionally is composed of a photodetector and a transimpedance amplifier. EDC is usually implemented with a set of transversal filters helping the time-delayed portions of the incoming electronic signal get recombined and amplified with suitable levels. The output of these filters is inherently the weighted sum of a number of time-delayed inputs. We can see a generic set of transversal filters in Figure 7.

Figure 7. Typical Transversal Filtering Scheme in an EDC Chip

The optimum settings of the weights in the filtering scheme depend upon the nature of the dispersion, which again varies greatly from one link to another depending on the fiber, the launch conditions, the link length, and environmental changes. Therefore, the most important aspect of an EDC solution is its ability to automatically adjust the filter weights according to the characteristics of the received signal. This enhanced and unique feature reveals EDC as a continuously adaptive solution, responding quickly to changes in the characteristics of the received signal.

This implementation of EDC is called continuous time filtering (CTF) and depicts the simplest, most cost-effective, and lowest power solution. It is usually appropriate for applications in which dispersion is not excessive. For more complex cases where the dispersion amount results in ISI exceeding one unit interval (UI) of the data rate, other sophisticated equalization algorithms are implemented. Indeed, one UI represents the time allocated to a generated pulse and equals the

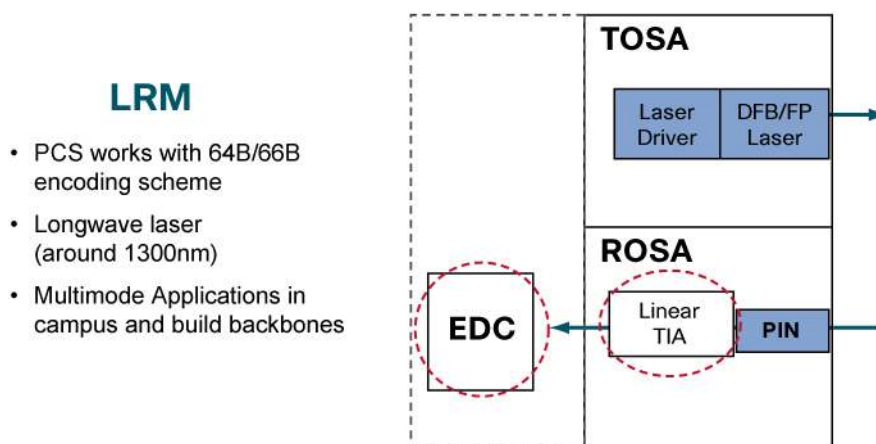
inverse of the signal bandwidth. For 10-Gbps transmissions, one unit interval equals about 100ps. When there is only a single UI of interference, compensation involves determining whether an adjacent symbol has spread into the current symbol and then adding or subtracting the symbol. When more than one UI of interference is present, a symbol can spread and distort several adjacent symbols, making compensation more complex. For those cases, a first solution is known as feed-forward/decision-feedback equalization (FFE/DFE) and uses multitap algorithms, where FFE removes distortion before a symbol's primary energy point or precursor area, and DFE compensates for interference following a symbol's primary energy point or postcursor area. Another more sophisticated solution is known as maximum likelihood sequence estimator (MLSE). MLSE equalization architectures incorporate a Viterbi decoder and require a digital signal processor (DSP) approach to filtering. The DSP is often a source of high power consumption, making MLSE-based approaches implemented only in applications having little room for compromise in performance.

As a summary, traditional receivers without EDC can recover an optical signal only if the dispersion is less than approximately one-half UI over the length of fiber. The IEEE 802.3aq standard, however, supports runs as long as 220m over OM1 and FDDI-grade fiber types and specifies that the receiver must be able to handle more than four UI of dispersion. Without EDC, we cannot possibly meet this requirement.

Summary: Deploying 10GBASE-LRM in Enterprise Networks

In this paper, we have seen the major benefits offered by the LRM solution and the EDC feature. Figure 8 shows the block diagram highlighting the simplicity and robustness of the design.

Figure 8. 10GBASE-LRM Block Diagram and Primary Features



IT managers and network architects now have a robust, cost-effective alternative to upgrade their campus and building backbone links from 1 Gbps to 10 Gbps without a change to the existing fiber plant and can address the ever-growing demand for bandwidth. 10GBASE-LRM devices are capable of transmitting and receiving 10-Gbps traffic over 220-meter strands of legacy FDDI-grade and OM1 fibers, the fiber types that offer the highest level of technical challenges. Potentially, transmission distances could be as much as 300 meters or even longer with applications over OM3 and future OM3+ fibers. Even though these reaches are currently not guaranteed by the IEEE standard, additional amendments could be published with the support of new high-performance MMF types. This would help address 99 percent of overall upgrade requirements in enterprise connectivity.

Cisco® stands as a primary player in the development of the product and becomes in June 2007 the first networking company offering this solution to its customers.

References

1. IEEE Std 802.3aq-2006, Amendment 2: Physical Layer and Management Parameters for 10Gb/s Operation, Type 10GBASE-LRM.
2. IEEE Std 802.3-2005.
3. "Testing and Interoperability of 10GBASE-LRM Optical Interfaces", IEEE Communications Magazine, March 2007.

For additional information about Cisco transceiver modules, refer to <http://www.cisco.com/en/US/products/ps6574/index.html>.

About the Author

Houfar Azgomi

Technical Marketing Engineer, Transceiver Module Group, Cisco



Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Europe Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

©2007 Cisco Systems, Inc. All rights reserved. CCVP, the Cisco logo, and the Cisco Square Bridge logo are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networking Academy, Network Registrar, Packet, PIX, ProConnect, ScriptShare, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0705R)