




10 Gigabit Ethernet Technology Overview

Intel[®] PRO
Network Connections



Contents

Executive Summary	3
The 10 Gigabit Ethernet Standard	3
The 10 Gigabit Ethernet in the Marketplace	4
Applications for 10 Gigabit Ethernet	4
10 Gigabit Ethernet as a Fabric Interconnect	4
10 Gigabit Ethernet in Local Area Networks	5
10 Gigabit Ethernet in Metropolitan and Storage Applications	6
10 Gigabit Ethernet in Wide Area Networks	7
Using Fiber in 10 Gigabit Ethernet	7
The Physical-Media-Dependent Devices (PMDs)	7
Fiber	8
The Future of 10 Gigabit Ethernet	8
Will There be Copper?	8
Conclusion	8
Glossary	9
Definitions	9
Terms	9
Appendix A – Standards Activities	10

Executive Summary

From its origin more than 25 years ago, Ethernet has evolved to meet the increasing demands of packet-based networks. Due to its proven low implementation cost, reliability, and relative simplicity of installation and maintenance, Ethernet's popularity has grown to the point that nearly all traffic on the Internet originates or terminates with an Ethernet connection. Further, as the demand for ever-faster network speeds has increased, Ethernet has been adapted to handle these higher speeds, as well as the surges in volume demand that accompany them.

The IEEE 802.3ae* 2002 (10 Gigabit Ethernet standard) is different in some respects from earlier Ethernet standards in that it will only function over optical fiber, and only operates in full-duplex mode (collision-detection protocols are unnecessary). Ethernet can now progress to 10 gigabits per second while retaining its critical Ethernet properties, such as the packet format, and the current capabilities are easily transferable to the new standard.

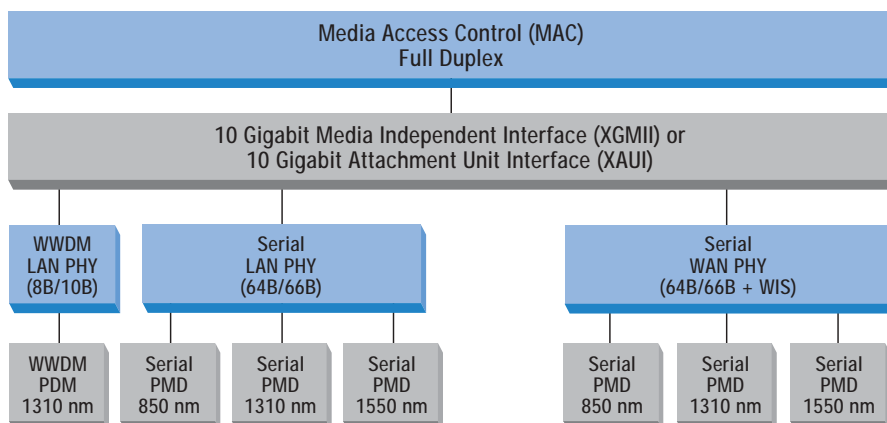
The 10 Gigabit Ethernet Standard

The 10 Gigabit Ethernet standard extends the IEEE 802.3ae* standard protocols to a wire speed of 10 Gbps and expands the Ethernet application space to include WAN-compatible links. The 10 Gigabit Ethernet standard provides a significant increase in bandwidth while maintaining maximum compatibility with the installed base of 802.3 standard interfaces, protects previous investment in research and development, and retains the existing principles of network operation and management.

Under the Open Systems Interconnection (OSI) model, Ethernet is fundamentally a Layer 1 and 2 protocol. 10 Gigabit Ethernet retains key Ethernet architecture, including the Media Access Control (MAC) protocol, the Ethernet frame format, and the minimum and maximum frame size. Just as Gigabit Ethernet, both 1000BASE-X and 1000BASE-T, followed the standard Ethernet model, 10 Gigabit Ethernet continues the evolution of Ethernet in speed and distance, while retaining the same Ethernet architecture used in other Ethernet specifications, except for one key ingredient. Since 10 Gigabit Ethernet is a full-duplex only technology, it does not need the carrier-sensing multiple-access with collision detection (CSMA/CD) protocol used in other Ethernet technologies. In every other respect, 10 Gigabit Ethernet matches the original Ethernet model.

At the physical layer (Layer 1), an Ethernet physical layer device (PHY) connects the optical or copper media to the MAC layer through a connectivity technology (Figure 1). Ethernet architecture further divides the physical layer into three sublayers: Physical Medium Dependent (PMD), Physical Medium Attachment (PMA), and Physical Coding Sublayer (PCS). PMDs provide the physical connection and signaling to the medium; optical transceivers, for example, are PMDs. The PCS consists of coding (e.g., 64B/66B) and a serializer or multiplexor. The IEEE 802.3ae* standard defines two PHY types: the LAN PHY and the WAN PHY. They provide the same functionality, except the WAN PHY has an extended feature set in the PCS that enables connectivity with SONET STS-192c/SHD VC-4-64c networks.

Figure 1. The architectural components of the 802.3ae* standard.



10 Gigabit Ethernet in the Marketplace

Ethernet technology is currently the most deployed technology for high-performance LAN environments. Enterprises around the world have invested cabling, equipment, processes, and training in Ethernet. In addition, the ubiquity of Ethernet keeps its costs low, and with each deployment of next-generation Ethernet technology, deployment costs have trended downward. In networks today, the increase in worldwide network traffic is driving service providers, enterprise network managers and architects to look to faster network technologies to solve increased bandwidth demands. 10 Gigabit Ethernet has ten times the performance over Gigabit Ethernet today. With the addition of 10 Gigabit Ethernet to the Ethernet technology family, a LAN now can reach further distances and support even more bandwidth-hungry applications. 10 Gigabit Ethernet also meets several criteria for efficient and effective high-speed network performance, which makes it a natural choice for expanding, extending, and upgrading existing Ethernet networks:

- A customer's existing Ethernet infrastructure is easily interoperable with 10 Gigabit Ethernet. The new technology provides lower cost of ownership including both acquisition and support costs versus current alternative technologies.
- Using processes, protocols, and management tools already deployed in the management infrastructure, 10 Gigabit Ethernet draws on familiar management tools and a common skills base.
- Flexibility in network design with server, switch, and router connections.
- Multiple vendor sourcing of standards-based products provides proven interoperability.

As 10 Gigabit Ethernet enters the market and equipment vendors deliver 10 Gigabit Ethernet network devices, the next step for enterprise and service provider networks is the combination of multi-gigabit bandwidth with intelligent services, which leads to scaled, intelligent, multi-gigabit networks with backbone and server connections ranging up to 10 Gbps. Convergence of voice and data networks running over Ethernet becomes a very real option. And, as TCP/IP incorporates enhanced services and features, such as packetized voice and video, the underlying Ethernet can also carry these services without modification.

The 10 Gigabit Ethernet standard not only increases the speed of Ethernet to 10 Gbps, but also extends its interconnectivity and its operating distance up to 40 km. Like Gigabit Ethernet, the 10 Gigabit Ethernet standard (IEEE 802.3ae*) supports both single-

mode and multimode fiber mediums. However, in 10 Gigabit, the distance for single-mode (SM) fiber has expanded from 5 km in Gigabit Ethernet to 40 km in 10 Gigabit Ethernet. The advantage of reaching new distances gives companies who manage their own LAN environments the option to extend their data center to a more cost-effective location up to 40 km away from their campuses. This also allows them to support multiple campus locations within the 40 km distance.

As we have seen with previous versions of Ethernet, the cost for 10 Gbps communications has the potential to drop significantly with the development of 10 Gigabit Ethernet-based technologies. Compared to 10 Gbps telecommunications lasers, the 10 Gigabit Ethernet technology, as defined in the IEEE 802.3ae*, will be capable of using lower cost, non-cooled optics, and vertical cavity surface emitting lasers (VCSEL), which can lower PMD device costs. In addition, an aggressive merchant chip market that provides highly integrated silicon solutions supports the industry.

Applications for 10 Gigabit Ethernet

Vendors and users generally agree that Ethernet is inexpensive, well understood, widely deployed and backwards compatible in today's LAN networks. Today, a packet can leave a server on a short-haul optic Gigabit Ethernet port, move cross-country via a DWDM (dense-wave division multiplexing) network, and find its way down to a PC attached to a Gigabit copper port, all without any re-framing or protocol conversion. Ethernet is literally everywhere, and 10 Gigabit Ethernet maintains this seamless migration in functionality for any application in which Ethernet can be applied.

10 Gigabit Ethernet as a Fabric Interconnect

Fabric interconnects, whether they are for server area networks or storage area networks, have traditionally been the domain of dedicated, often proprietary, networks with relatively small user bases when compared to Ethernet. These server area networks include InfiniBand*, Servernet*, Myranet*, Wulffit* and Quadrics* technologies, and offer excellent bandwidth and latency performance for very short-haul (generally less than 20 m) networks. However, with the exception of InfiniBand, these are proprietary networks that can be difficult to deploy and maintain due to the small number of experienced IT professionals familiar with the technology. The small volumes also result in higher costs for server adapters and switches. And, as with any proprietary solution, they are not interoperable with other technologies without the appropriate routers and switches.

In storage area networks, the lack of standards and a slew of interoperability problems plagued the early Fibre Channel deployments. However, these technologies also suffer similar problems as those seen by proprietary server area networks in that they are considered difficult to deploy due to lack of a skilled IT pool, are relatively expensive at the adapter and switch port, are still not directly interoperable with other network technologies without expensive routers or switching devices, and generally focus on short-haul deployments.

10 Gigabit Ethernet is in a position to replace these proprietary technologies as a next-generation interconnect for both server and storage-area networks for several reasons:

- **10 Gigabit Ethernet Offers the Necessary Bandwidth.** In fact, InfiniBand and Fibre Channel will also begin mass deployments of 10 Gigabit technologies, indicating a convergence on 10 Gigabit throughput.
- **Cost-Saving Server Consolidation.** 10 Gigabit Ethernet grants a single server the bandwidth needed to replace several servers that were doing different jobs. Centralization of management is also a major benefit of server consolidation. With a single powerful server, IT managers can monitor, manage, and tune servers and application resources from a single console, which saves time and maximizes IT resources. According to IDC, companies realize a seven-to-one savings in management when processes and servers are consolidated.¹

¹ IDC, Worldwide Server Consolidation Forecast, September 2002

- **Planned Growth of 10 Gigabit Network Features.** For the first time ever, Ethernet can be a low-latency network due to RDMA (Remote Direct Memory Access) support, which is critical in the server-to-server communication typically associated with clustering and server area networks.

In addition, the expected universal deployment of TOE (TCP/IP Offload Engine) technology in 10 Gigabit Ethernet adapters may make it extremely efficient on host systems with expected CPU utilization well below anything seen on today's systems deploying Gigabit Ethernet. Due to the wide adoption rate of Ethernet, TOE technology will become extremely cost efficient compared to the lower volume, niche alternatives.

10 Gigabit Ethernet in Local Area Networks

Ethernet technology is already the most deployed technology for high-performance LAN environments. With the extension of 10 Gigabit Ethernet into the family of Ethernet technologies, LANs can provide better support the rising number of bandwidth-hungry applications and reach greater distances. Similar to Gigabit Ethernet technology, the 10 Gigabit standard supports both single-mode and multimode fiber media.

With links up to 40 km, 10 Gigabit Ethernet allows companies that manage their own LAN environments the ability to strategically choose the location of their data center and server farms – up to 40 km away from their campuses. This enables them to support multiple campus locations within that 40 km range (Figure 2). Within data centers, switch-to-switch

Figure 2. 10 Gigabit Ethernet use in expanded LAN environments.

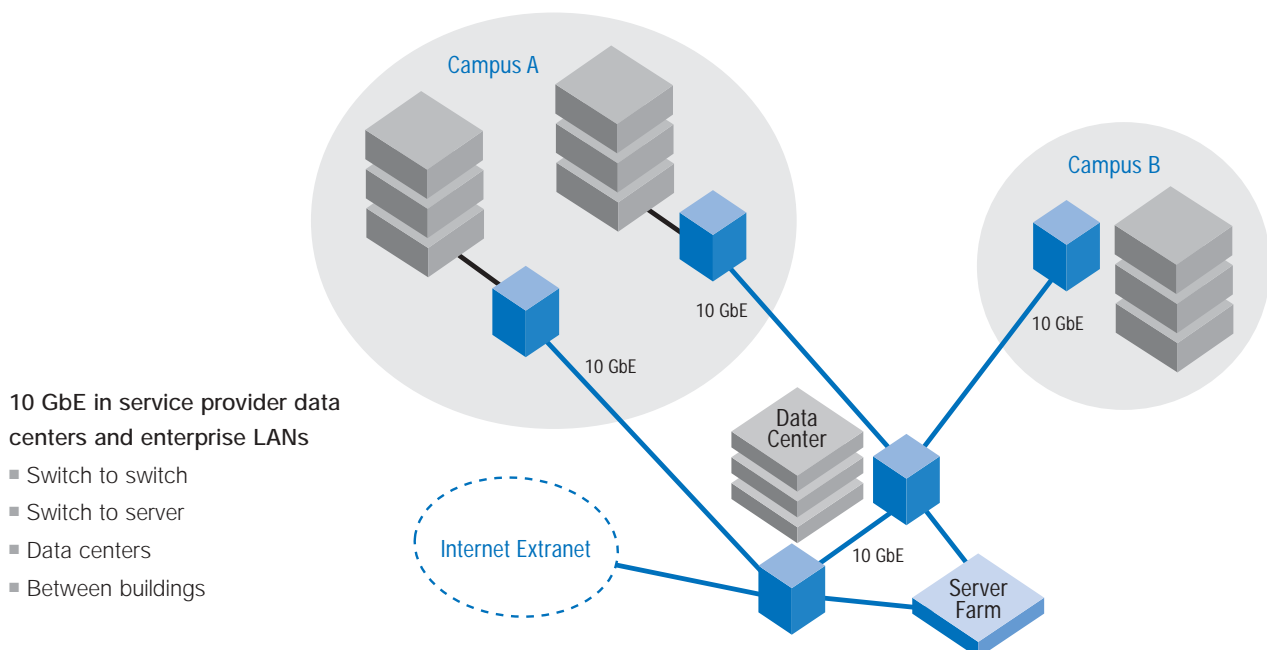
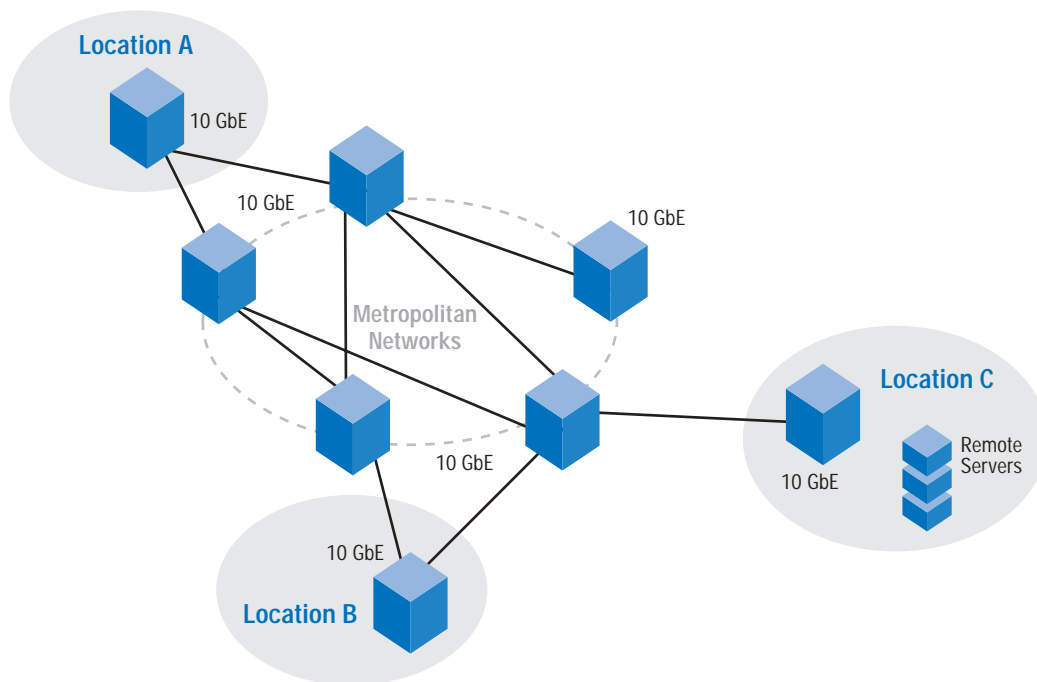


Figure 3. Example of 10 Gigabit Ethernet use in a MAN.



applications, as well as switch-to-server applications, can be deployed over a more cost-effective, short-haul, multi-mode fiber medium to create 10 Gigabit Ethernet backbones that support the continuous growth of bandwidth-hungry applications.

With 10 Gigabit backbones, companies can easily support Gigabit Ethernet connectivity in workstations and desktops with reduced network congestion, enabling greater implementation of bandwidth-intensive applications, such as streaming video, medical imaging, centralized applications, and high-end graphics. 10 Gigabit Ethernet also improves network latency, due to the speed of the link and over-provisioning bandwidth, to compensate for the bursty nature of data in enterprise applications.

The bandwidth that 10 Gigabit backbones provide also enables the next generation of network applications. It can help make telemedicine, telecommuting, distance learning and interactive, digital videoconferencing everyday realities instead of remote future possibilities. And the fun stuff too, like HDTV, video-on-demand, or extreme Internet gaming.

10 Gigabit Ethernet enables enterprises to reduce network congestion, increase use of bandwidth-intensive applications, and make more strategic decisions about the location of their key networking assets by extending their LAN up to 40 km.

10 Gigabit Ethernet in Metropolitan and Storage Applications

Gigabit Ethernet is already being deployed as a backbone technology for dark fiber metropolitan networks. With appropriate 10 Gigabit Ethernet interfaces, optical transceivers and single-mode fiber, network and Internet service providers will be able to build links reaching 40 km or more (Figure 3), encircling metropolitan areas with city-wide networks.

10 Gigabit Ethernet now enables cost-effective, high-speed infrastructure for both network attached storage (NAS) and storage area networks (SAN). Prior to the introduction of 10 Gigabit Ethernet, some industry observers maintained that Ethernet lacked sufficient horsepower to get the job done.

10 Gigabit Ethernet can now offer equivalent or superior data carrying capacity at latencies similar to many other storage networking technologies, including Fiber Channel, Ultra160 or 320 SCSI, ATM OC-3, OC-12, and OC-192, and HIPPI (High-Performance Parallel Interface). Gigabit Ethernet storage servers, tape libraries, and compute servers are already available; 10 Gigabit Ethernet end-point devices will soon appear on the market as well.

There are numerous applications for Gigabit Ethernet today, such as back-up and database mining. Some of the applications that will take advantage of 10 Gigabit Ethernet are:

- Business continuance/disaster recovery
- Remote back-up
- Storage on demand
- Streaming media

10 Gigabit Ethernet in Wide Area Networks

10 Gigabit Ethernet enables ISPs and NSPs to create very high-speed links at a very low cost from co-located, carrier-class switches and routers to the optical equipment directly attached to the SONET/SDH cloud. 10 Gigabit Ethernet, with the WAN PHY, also allows the construction of WANs that connect geographically dispersed LANs between campuses or points of presence (POPs) over existing SONET/SDH/TDM networks. 10 Gigabit Ethernet links between a service provider's switch and a DWDM device or LTE (line termination equipment) might in fact be very short – less than 300 meters.

Using Fiber in 10 Gigabit Ethernet

The Physical-Media-Dependent Devices (PMDs)

The IEEE 802.3ae* standard provides a physical layer that supports specific link distances for fiber-optic media. To meet the distance objectives, four PMDs (physical-media-dependent devices) were selected (Table A). The task force selected

- A 1310 nanometer serial PMD to support single-mode fiber a maximum distance of 10 km
- A 1550 nanometer serial PMD to support single-mode fiber a maximum distance of 40 km
- An 850 nanometer serial PMD to support multimode fiber a maximum distance of 300 meters
- A 1310 nanometer wide-wave division multiplexing (WDM) PMD to support a maximum distance of 10 km on single-mode fiber, as well as a maximum distance of 300 meters on multimode fiber

Table A. PMDs that have been selected to meet the 803.2ae* standard's distance objectives.

Device	8B/10B PCS	64B/66B PCS	WIS	850 nm Serial	1310 nm WDM	1310 nm Serial	1550 nm Serial
10GBASE-SR		■		■			
10GBASE-SW		■	■	■			
10GBASE-LX4	■				■		
10GBASE-LR		■				■	
10GBASE-LW		■	■			■	
10GBASE-ER		■					■
10GBASE-ER		■	■				■

Table B. The multimode optical fiber options, as defined in the IEEE 802.3ae* standard.

Multimode Fiber	62.5 MMF			50 MMF	
	160	200	400	500	2000
850 nm Serial	26m	33m	66m	82m	300m
1310 nm LX4	300m @ 500MHz*km		240m	300m	

Fiber

There are two types of optical fiber, multimode and single-mode fiber, that are currently used in data networking and telecommunications applications. The 10 Gigabit Ethernet technology, as defined in the IEEE 802.3ae* standard, supports both optical fiber types. However, the distances supported vary based on the type of fiber and wavelength (nm) is implemented in the application. In single-mode fiber applications, the IEEE 802.3ae standard supports 10 km with 1310 nm optical transmissions and 40 km with 1550 nm optical transmissions. With multimode optical fiber, the distances are not as easily defined due to the variety of fiber types and the way each type is defined. Multimode fiber is commonly defined by the core and cladding diameters. For example, fiber with a core of 62.5 microns and a cladding diameter of 125 microns is referred to as 62.5/125 micron fiber. The acceptance of multimode fiber in networks today dates back to the inclusion of 62.5/125 micron fiber into the Fiber Distribution Data Interface (FDDI) standard in the 1980s. The other portion that influences distance capabilities in multimode fiber is the fiber information carrying capacity (measured in MHz-km), which determines the distance and bit rate at which a system can operate (i.e., 1 Gbps or 10 Gbps). The distance a signal can run greatly decreases as transmission speed increases (Table B). When implementing multimode fiber for 10 Gigabit Ethernet applications, understanding the distance capabilities is a critical piece to the 10 Gigabit Ethernet solutions.

The Future of 10 Gigabit Ethernet:

Will There be Copper?

IEEE 802.3* has recently formed two new study groups to investigate 10 Gigabit Ethernet over copper cabling. The 10GBASE-CX4 study group is developing a standard for transmitting and receiving XAUI signals via a 4-pair twinax cable, commonly referred to as a 4x InfiniBand cable. The goal of the study group is to provide a standard for a low-cost inter-rack and rack-to-rack solution. It is expected to take about one year to develop a standard. The 10GBASE-T study group is developing a standard for the transmission and reception of 10 Gigabit Ethernet via a Category 5 or better unshielded twisted pair (UTP) copper cable up to 100 m. This effort is expected to take much longer than the 10GBASE-CX4 effort and current estimates are that the effort will complete sometime in late 2005 or early 2006.

Conclusion

Ethernet has withstood the test of time to become the most widely adopted networking technology in the world. With the rising dependency on networks and the increasing number of bandwidth-intensive applications, service providers seek higher capacity networking solutions that simplify and reduce the total cost of network connectivity, thus permitting profitable service differentiation, while maintaining very high levels of reliability. The 10 Gigabit Ethernet IEEE 802.3ae* 10 Gigabit Ethernet standard is proving to be a solid solution to network challenges.

10 Gigabit Ethernet is the natural evolution of the well-established IEEE 802.3* standard in speed and distance. In addition to increasing the line speed for enterprise networks, it extends Ethernet's proven value set and economics to metropolitan and wide area networks by providing

- Potentially lowest total cost-of-ownership (infrastructure/operational/human capital)
- Straight-forward migration to higher performance levels
- Proven multi-vendor and installed-base interoperability (Plug and Play)
- Familiar network management feature set

An Ethernet-optimized infrastructure is taking place in the metropolitan area and many metropolitan areas are currently the focus of intense network development intending to deliver optical Ethernet services. 10 Gigabit Ethernet is on the roadmap of most switch, router and metropolitan optical system vendors to enable:

- Cost-effective, Gigabit-level connections between customer access gear and service provider POPs in native Ethernet format.
- Simple, high-speed, low-cost access to the metropolitan optical infrastructure.
- Metropolitan-based campus interconnection over dark fiber, targeting distances of 10 to 40 km.
- End-to-end optical networks with common management systems.

Glossary

Definitions

802.3ae – The IEEE standard for 10 Gigabit Ethernet

802.3ab – The IEEE standard for UTP Gigabit Ethernet (1000BASE-T)

802.3z – The IEEE standard for Gigabit Ethernet (1000BASE-X)

CoS – Class of Service

CWDM – Coarse-Wavelength Division Multiplexing

DWDM – Dense-Wavelength Division Multiplexing

Gbps – Gigabits per second or billion bits per second

IEEE – Institute of Electrical and Electronics Engineers

IP – Internet Protocol

ISO – International Standards Organization

LAN – Local Area Network

MAC – Media Access Control

MAN – Metropolitan Area Network

Mbps – Megabits per second or million bits per second

MMF – Multimode Fiber

OC-X – Optical Carrier Level

PCS – Physical Coding Sublayer

PHY – Physical layer device

PMA – Physical Medium Attachment

PMD – Physical Medium Dependent

POP – Points of Presence

QoS – Quality of Service

RMON – Remote Monitoring

SDH – Synchronous Digital Hierarchy

SMF – Single-mode Fiber

SNMP – Simple Network Management Protocol

SONET – Synchronous Optical Network

Tbps – Terabits per second or trillion bits per second

TCP/IP – Transmission Control Protocol/Internet Protocol

TDM – Time Division Multiplexing

WAN – Wide Area Network

WDM – Wavelength Division Multiplexing

WIS – WAN Interface Sublayer

WWDM – Wide-Wavelength Division Multiplexing

Terms

Cladding – The material surrounding the core of an optical fiber. The cladding has a lower refractive index (faster speed) that is used to keep the light in the core. The cladding and core make up an optical waveguide.

Core – The central region of an optical fiber through which light is transmitted. It has a higher refractive index (slower speed) than the surrounding cladding.

Dense-Wave Division Multiplexing – Wavelengths are closely spaced, allowing more channels to be sent through one fiber. Currently, systems using 100 GHz spacing are deployed in the WAN environment. Overall wavelength range is typically between 1530 nm to 1560 nm. The minimum and maximum wavelengths are restricted by the wavelength dependent gain profile of optical amplifiers.

Media Access Control – The media access control sublayer provides a logical connection between the MAC clients of itself and its peer station. Its main responsibility is to initialize, control, and manage the connection with the peer station. The MAC layer of the 10 Gigabit protocol uses the same Ethernet address and frame formats as other speeds, and will operate in full-duplex mode. It will support a data rate of 10 Gbps using a pacing mechanism for rate adaptation when connected to a WAN-friendly PHY.

OC-192 – A speed of SONET interconnect with a payload rate of 9.584640 Gbps, primarily used in WAN environments.

Physical Coding Sublayer – Part of the PHY, the PCS sublayer is responsible for encoding the data stream from the MAC layer for transmission by the PHY layer and decoding the data stream received from the PHY layer for the MAC layer.

PHY – The physical layer device, a circuit block that includes a PMD (physical media dependent), a PMA (physical media attachment), and a PCS (physical coding sublayer).

PMD – Part of the PHY, the physical-media-dependent sublayer is responsible for signal transmission. The typical PMD functionality includes amplifier, modulation, and wave shaping. Different PMD devices may support different media.

WWDM – A technique used to effectively transmit several wavelengths (i.e., colors of light) from several laser sources through one fiber. Each laser source would be calibrated to send a unique optical wavelength (which are separated at the receiving end of the fiber).

Appendix A – Standards Activities

The Institute of Electrical and Electronics Engineers (IEEE) was founded to foster the development of standards in all fields of science and technology within the organization's scope. A key principle throughout the standards process is consensus among the participants. The IEEE-SA (Standards Association) and its Standards Board oversee the process of standards formation through two committees.

The New Standards Committee (NesCom) ensures that proposed standards fall within the IEEE's scope, that they are assigned to the correct technical committees, and that the makeup of working groups, etc., fairly represents all interested parties. It also examines Project Authorization Requests (PAR) and recommends to the IEEE-SA Standards Board whether to approve them. The second committee, the Standards Review Committee (RevCom), examines proposed new and revised standards, ensures that such proposals represent a consensus among the IEEE Sponsor balloting group members, and recommends standards to the Standards Board.

For further information on the IEEE standards process or the 10 Gigabit Ethernet technology, visit the following Web sites:

IEEE

IEEE 802 LAN/MAN Standards Committee

IEEE 802.3 CSMA/CD (ETHERNET)

IEEE P802.3ae 10Gb/s Ethernet Task Force

<http://standards.ieee.org/resources/glance.html>

10 Gigabit Ethernet Alliance (10GEA)

www.10gea.org



INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life saving, or life sustaining applications. Intel may make changes to specifications and product descriptions at any time, without notice.

* Other names and brands may be claimed as the property of others.
Copyright © 2003 Intel Corporation. All rights reserved.
Intel, the Intel logo are trademarks or registered trademarks of Intel Corporation.

Printed in USA.

0203/OC/PL/PDF

 Please Recycle

252451-001