The 10 Gigabit Ethernet

G. Kiranmai[1], R. V. Krishnaiah[2]

[1]M. Tech (Scholar) DRK College of Engineering & Technology, Quthbullapur, India

[2]PhD, M. Tech (CSE), M. Tech (EIE) B. Tech (ECE)

Abstract: ***Since its inception at Xerox Corporation in the early 1970s, Ethernet has been the dominant networking protocol. Of all current networking protocols, 10 Gigabit Ethernet builds on top of the Ethernet protocol, but increases speed tenfold over Fast Ethernet to 10000 Mbps, or 10 gigabit per second (Gbps). This protocol, which was standardized in august 2002, promises to be a dominant player in high-speed local area network backbones and server connectivity. Since10 Gigabit Ethernet significantly leverages on Ethernet, customers will be able to leverage their existi0ng knowledge base to manage and maintain gigabit networks.***

Keywords: Gigabit Ethernet, 10G Ethernet

1. Introduction

The 10 Gigabit Ethernet standards extend the IEEE 802.3ae* standard protocols to a wire speed of 10 Gbps and expand the Ethernet application space to include WAN-compatible links. The 10 Gigabit Ethernet standards 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 sub-layers: Physical Medium Dependent (PMD), Physical Medium Attachment (PMA), and Physical Coding Sub-layer (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.10 Gigabit Ethernet Technology Overview.

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.

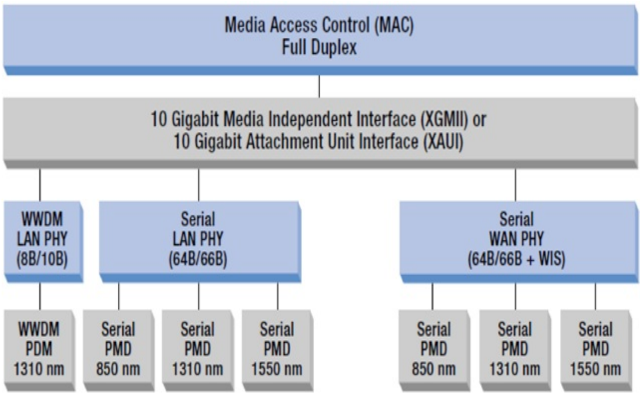

Figure 1: The architectural components of the 802.3ae* standard

2. 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 is required in network design with server, switch and router connections. Multiple vendors sourcing of standards-based products provide 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 standards not only increase the speed of Ethernet to 10 Gbps, but also extend 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.

3. 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.

3.1 10 Gigabit Ethernet as a Fabric Interconnect

Fabric interconnects, whether they are for server area networks or storage area networks, has 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*, Wulfkit* 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 Fiber 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 interconnects both server and storage-area networks for several reasons.

3.1.1 10 Gigabit Ethernet Offers the Necessary Bandwidth

In fact, InfiniBand and Fiber Channel will also begin mass deployments of 10 Gigabit technologies, indicating a convergence on 10 Gigabit throughputs. 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.

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.

3.2 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).

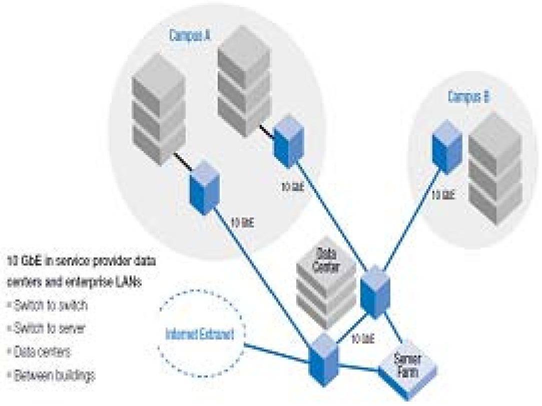

Figure 2: 10 Gigabit Ethernet use in expanded LAN environments

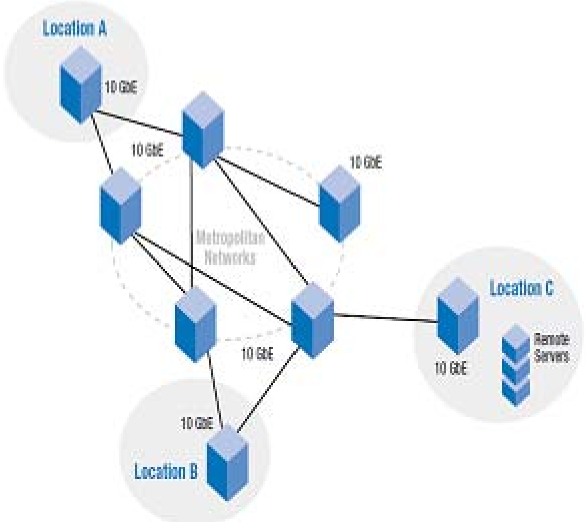

Figure 3: Example of 10 Gigabit Ethernet use in a MAN

Within data centers, switch-to-switch 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, and 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.

3.3 10Gigabit 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 computer 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 is:

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

3.4. 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.

4. Using Fiber in 10 Gigabit Ethernet

4.1 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 (WWDM) 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.

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. 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.

Table 2: The multimode optical fiber options, as defined in the IEEE 802.3ae* standard Fiber

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

Table 1: 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 WWDM*	*1310 nm Serial*	*1550 nm Serial*
10GBASE-SR		x		x			
10GBASE-DW		x	x	x			
10GBASE-LX4	x				x		
10GBASE-LR		x				x	
10GBASE-LW		x	x			x	
10GBASE-ER		x					x
10GBASE-ER		X	x				x

5. The Future of 10 Gigabit Ethernet

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 cables 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.

6. 10 Gigabit Ethernet in Servers

6.1 Benefits of 10 Gigabit Ethernet

Server Consolidation & Virtualization In a drive to reduce operating costs, corporations are consolidating their servers into fewer but larger data centers. This physical consolidation has the effect of concentrating network traffic. Because most servers already have Gigabit Ethernet (GbE) connections, aggregating the traffic from a server farm requires a higher-bandwidth connection than can be achieved with GbE links. One of the first applications of 10 Gigabit Ethernet inXthe data center is in removing this bottleneck by connecting servers to a GbE switch with10GbE uplink ports. Taking consolidation one step further, blade servers combine multiple servers into a single chassis.

Each server blade connects to the system backplane using a GbE connection. An integrated Ethernet switch connects the blades and provides the connections to the local-area network (LAN). With dozens of CPUs generating network traffic within the chassis, GbE connections to the LAN would be massively oversubscribed. By using 10GbE connections to the LAN, blade servers can scale to larger numbers of processors without the network connection becoming a major bottleneck.

Although blade servers are another step in physical consolidation, most applications still reside on a dedicated server. This model of having one application per server has caused a proliferation of servers in corporate networks. It also makes it difficult to scale and balance performance across servers. Virtualization—decoupling the operating system image from the physical system—promises to eliminate these scalability issues. Virtualization should also allow powerful multiprocessor systems to be more fully utilized as CPU resources are balanced across multiple applications. The combination of consolidation and virtualization will drive server performance to new levels. But as system performance scales, GbE connections to the LAN become a bottleneck. Trunking multiple GbE links is an interim solution, but trunking introduces processing overhead and requires multiple cables (or fibers). As a result, trunking more than four GbE links is uncommon. With 10 times the bandwidth of GbE, 10GbE is the clear answer to increasing LAN connection speeds.

7. 10G Ethernet in Servers

7.1 Benefits and Challenges

LAN, SAN, and Cluster Convergence:

Today, large data centers use specialized networks for storage and for server clusters. Fiber Channel (FC) is the technology of choice for storage-area networks (SAN). FC networks typically operate at 1Gbps or 2Gbps, while 4Gbps products are just becoming available. Because FC SANs are primarily deployed in large data centers, volumes remain small compared with Ethernet-based LANs. IP storage, which uses Ethernet as a storage network, is emerging as an alternative to FC. The key protocol in migrating storage traffic to the LAN is iSCSI. The iSCSI protocol is already supported in shipping versions of Windows, Linux, and Unix operating systems, and iSCSI host-bus adapters (HBA) are available from major vendors.

Server clusters have been connected using proprietary protocols, such as Myrinet, and more recently using InfiniBand. Although InfiniBand is an industry standard, it serves an arrow market and has not achieved high volumes. InfinBand also has limited cable reach, making it suitable for use only within a data center or a small HPC cluster. On the plus side, InfiniBand provides 10Gbps connections with little delay (or latency), which is critical to parallel-processing applications. Next-generation InfiniBand products will scale to 30Gbps, keeping InfiniBand ahead of Ethernet in peak performance. There are also industry efforts to redirect the remote direct-memory access (RDMA) technology developed for InifiniBand for use over lower-cost Ethernet. Commonly known as I WARP, the protocols for RDMA-over-IP promise to reduce the latency of 10GbE to a level similar to that of InfiniBand.

An application-layer interface, iWARP will allow cluster applications to be ported to 10GbE networks. Although FC and InfiniBand serve their niches well, installing, maintaining, and managing three networks is clearly inefficient. By leveraging new technologies such as iSCSI and iWARP, 10GbE promises to consolidate the LAN, SAN, and cluster into a single network. This convergence of three fabrics into one will be especially important for blade servers due to their physical constraints. The use of lower-cost Ethernet technology should also broaden the market for SAN and cluster capabilities. For these reasons, the convergence of the LAN, SAN, and cluster networks using 10GbE has broad industry support from chip and system vendors alike. Despite this convergence at 10Gbps, we expect the installed base of FC and InifiniBand to continue to be used, but most new 10Gbps installation will eventually move to 10GbE.

7.2 Technology Issues in 10GbE Adoption

Before 10GbE gains broad adoption, some technology issues must be addressed. For servers, the biggest issue is protocol processing. As any IT manager knows, a network connection consumes a portion of a server's processing power (measured as CPU utilization). In a single-CPU server, a GbE link can consume close to half of the server's processing cycles. If a conventional network-interface card (NIC) architecture was 10G Ethernet in Servers: Benefits and Challenges One major source of processing overhead is the TCP/IP stack.

As a result, there have been ongoing efforts to offload some TCP processing from the system CPU onto the NIC hardware. Some of these TCP-offload functions are available today in 10GbE NIC implementations, while others are still under development. Current products support TCP checksum and large-send (or segmentation) offloads, which yield sizable reductions in CPU utilization. These offloads are supported natively in current versions of Windows and through modified stacks supplied by NIC vendors or system OEMs for Linux and UNIX. Another way to reduce TCP overhead is by using special large packets known as jumbo frames. Jumbo frames are useful in storage applications that transfer large blocks of data.

Although NIC vendors can offload all TCP processing by bypassing operating system network stacks, end users have been reluctant to use such products unless a major server vendor directly supports them. This support is not an issue for UNIX systems, because the server vendor qualifies specific NIC hardware and a modified TCP/IP stack for use on their systems. To address this problem in Windows servers, Microsoft is developing technology known as TCP Chimney. When deployed, TCP Chimney will enable full TCP offload by NIC hardware. TCP Chimney is due for release in the Scalable Networking Pack for Windows Server 2003 during 1H05. Linux environments are problematic; NIC vendors can easily implement modified TCP/IP stacks, but rigorous qualification requirements are lacking as compared with Windows and UNIX systems

Another source of processing overhead is data copying. In a conventional networking stack, received packets are stored into operating-system memory and later copied to application memory. This copying consumes CPU cycles and also introduces a delay. For parallel-processing applications that use small buffers, such as distributed databases, data copying is a major performance drain. iWARP enables data to be written directly into application memory, eliminating the extra copy operation.

Another hurdle for 10GbE adoption is cabling and the related physical-layer standards. The original 10GbE standard defined multiple physical layers for fiber media but did not support copper (twisted-pair) cabling. Of these initial specifications, 10GBase-SR and 10GBase-LR are the most popular today. SR supports multimode fiber up to 300m in length, but is limited to 26m over FDDI-grade multimode fiber. Fortunately, 26m is adequate for most data-center applications. Typically used in metro applications, LR supports 10km reach over single-mode fiber but does not support multimode fiber, which is found in most enterprise applications. Another of the initial specifications, LX4, provides 300m reach over FDDI-grade fiber and also supports 10km reach over single mode fiber. LX4 products are now available, but they carry a price premium over SR. More-recent efforts have focused on copper cabling and reducing the cost of 300m multimode fiber applications. As an alternative to SR in the data center, the CX4 standard enables 15m reach over InfiniBand-style copper cables. Although CX4 uses large 8-pair shielded cables, it eliminates the need for costly optical modules. Products supporting CX4 have become available during 2004. Currently under development, LRM promises 300m reach over FDDI-grade fiber at a lower cost than LX4. Finally, 10GBase-T will enable the use of twisted-pair (UTP) cabling, but this standard will not be complete until 2006.

Unlike 1000Base-T (GbE), 10GBase-T will likely require CAT6 cabling and still may not achieve 100m reach.

7.3 Performance Impact of 10GbE

For many of the reasons discussed above, the performance impact of 10GbE will vary greatly by application and implementation. First-generation 10GbE NICs, which implement partial TCP offloads and a PCI-X system interface, deliver peak performance of 6–8Gbps. At large packet sizes, these NICs consume less than 100% of a typical server CPU. Thus, a first-generation 10GbE NIC should deliver 50–100% more throughput than a conventional 4xGbE NICs with a similar level of CPU utilization.10G Ethernet in Servers: Benefits and Challenges

Second-generation 10GbE NICs with TCP-offload engines (TOE) are becoming available but still use a PCI-X system interface. These TOE NICs should achieve throughput similar to that of first-generation 10GbE NICs while lowering CPU utilization. Third generation 10GbE NIC products are likely to adopt PCI Express x8 interfaces and should achieve full 10Gbps line rate with large packets. But achieving full-duplex line-rate performance (20Gbps) is likely to consume all of the processing power of two typical server CPUs. Network traffic fluctuates greatly, however, so this situation represents a peak load rather than a sustained concern.

8. Pros and Cons of 10GbE

Hardware Cost because 10GbE is in the early stages of adoption, prices for both switch ports and NICs have been high. But competition is heating up as vendors try to stake out a position in anticipation of 10GbE ramping to high volume. During 2004, 10GbE ports for modular switches have dropped from more than $10,000 per port to as low as $2,500 per port. By mid-2005, we expect to see fixed-configuration 10GbE switches selling for less than $1,000 per port without optical modules. Some products will integrate CX4 transceivers, eliminating the need for modules altogether.

Standard form factors (such as MSA) are also helping drive competition for pluggable optical modules. XPAK and X2 are the latest module types shipping in volume; 10GbE SR module pricing to OEMs has dropped to only $350. Although end-user prices for prior-generation modules (XENPAK) have been in the $3,000 range, XPAK/X2 prices should fall below $1,000 during 2005.

9. 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.

Reference

[1] IEEE
[2] IEEE 802 LAN/MAN Standards Committee
[3] IEEE 802.3 CSMA/CD (ETHERNET)
[4] IEEE P802.3ae 10Gb/s Ethernet Task Force
[5] http://standards.ieee.org/resources/glance.ht
[6] 10 Gigabit Ethernet Alliance (10GEA)
[7] www.10gea.org

Author Profile

G. Kiranmai, M. Tech in CSE from DRK College of Engineering & Technology, Quthbullapur (Mondl), Bowrampet (village), Ranga Reddy (dist), B. Tech in CSE from Vidhya Bharathi Institute & Technology, Jangaon (Mondl), Pemberthi (village), Warangal (dist).

Dr. R. V. Krishnaiah, PhD from JNTU Ananthapur, M. Tech in CSE from JNTU Hyderabad, M. Tech in EIE from NIT (former REC) Warangal, B. Tech ECE from Bapatla Engineering College, He has published more than 50 papers in various journals. His interested includes Data mining, Software Engineering and Image Processing.