

Gigabit Ethernet

[Vijay Moorthy, \(moorthy@cis.ohio-state.edu\)](mailto:moorthy@cis.ohio-state.edu)

Ethernet is the world's most pervasive networking technology. Gigabit Ethernet is the latest version of Ethernet. It offers 1000 Mbps (1 Gbps) raw bandwidth, that is 100 times faster than the original Ethernet, yet is compatible with existing Ethernets, as it uses the same CSMA/CD and MAC protocols. When Gigabit Ethernet enters the market it will compete directly with ATM. This paper presents a survey of Gigabit Ethernet technology.

[Other Reports on Recent Advances in Networking](#)

[Back to Raj Jain's Home Page](#)

Table of Contents

1. [Introduction](#)
 - 1.1 [History of Ethernet](#)
 - 1.2 [Gigabit Ethernet Alliance](#)
 2. [Physical Layer](#)
 - 2.1 [1000Base-X](#)
 - 2.2 [1000Base-T](#)
 3. [MAC Layer](#)
 - 3.1 [Carrier Extension](#)
 - 3.2 [Packet Bursting](#)
 4. [GMII \(Gigabit Media Independent Interface \)](#)
 - 4.1 [PCS \(Physical Coding Sublayer\)](#)
 - 4.2 [PMA \(Physical Medium Attachment\)](#)
 - 4.3 [PMD \(Physical Medium Dependent\)](#)
 5. [Buffered Distributor](#)
 6. [Topologies](#)
 - 6.1 [Upgrading server-switch connections](#)
 - 6.2 [Upgrading switch-switch connections](#)
 - 6.3 [Upgrading a Fast Ethernet backbone](#)
 - 6.4 [Upgrading a Shared FDDI Backbone](#)
 - 6.5 [Upgrading High Performance Workstations](#)
 7. [ATM vs. Gigabit Ethernet](#)
 8. [Summary](#)
 9. [Bibliography and Links](#)
-

1. Introduction

Ethernet is the world's most pervasive networking technology, since the 1970's. It is estimated that in 1996, 82% of all networking equipment shipped was Ethernet. In 1995, the Fast Ethernet Standard was approved by the IEEE. Fast Ethernet provided 10 times higher bandwidth, and other new features such as full-duplex operation, and auto-negotiation. This established Ethernet as a scalable technology. Now, with the emerging Gigabit Ethernet standard, it is expected to scale even further.

The Fast Ethernet standard was pushed by an industry consortium called the Fast Ethernet Alliance. A similar alliance, called the Gigabit Ethernet Alliance was formed by 11 companies in May 1996, soon after IEEE announced the formation of the 802.3z Gigabit Ethernet Standards project. At last count, there were over 95 companies in the alliance from the networking, computer and integrated circuit industries.

A draft 802.3z standard was issued by IEEE in July 1997. The last technical changes are expected to be resolved by September. The standard is expected to be adopted by March 1998.

The new Gigabit Ethernet standards will be fully compatible with existing Ethernet installations. It will retain Carrier Sense Multiple Access/ Collision Detection (CSMA/CD) as the access method. It will support full-duplex as well as half duplex modes of operation. Initially, single-mode and multi mode fiber and short-haul coaxial cable will be supported. Standards for twisted pair cables are expected by 1999. The standard uses physical signalling technology used in Fiber Channel to support Gigabit rates over optical fibers.

Initially, Gigabit Ethernet is expected to be deployed as a backbone in existing networks. It can be used to aggregate traffic between clients and "server farms", and for connecting Fast Ethernet switches. It can also be used for connecting workstations and servers for high - bandwidth applications such as medical imaging or CAD.

1.1 History of Ethernet

Today, Ethernet is synonymous with the IEEE 802.3 standard for a "1-persistent CSMA/CD LAN". The 802.3 standard has an interesting history. The beginning, is generally considered to be the University of Hawaii ALOHA network. This system is the ancestor of all shared media networks. The original Ethernet, developed by Xerox was based on the ALOHA system. It was a 2.94 Mbps CSMA/CD system and was used to connect over 100 personal workstations on a 1 Km cable. It was so successful, that Xerox, DEC and Intel came up with a 10 Mbps standard. The IEEE 802.3 standard was based on the 10 Mbps Ethernet.

CSMA/CD refers to the protocol used by stations sharing the medium, to arbitrate use of the medium. A sender has to "listen" to the medium. If no one else is transmitting, then the sender may transmit. If two senders start transmitting at the same time, then a *collision* is said to have occurred. Transmitting stations, therefore, have to listen to the medium for collisions while transmitting, and retransmit a packet after some time, if a collision occurs.

The original 802.3 standard was published in 1985. Originally two types of coaxial cables were used called *Thick Ethernet* and *Thin Ethernet*. Later unshielded copper twisted pair (UTP), used for telephones, was added.

In 1980, when Xerox, DEC and Intel published the DIX Ethernet standard, 10 Mbps was a lot of bandwidth. Since then, as computing technology improved, network bandwidth requirements also increased. In 1995, IEEE adopted the 802.3u Fast Ethernet standard. Fast Ethernet is a 100 Mbps Ethernet standard. Fast Ethernet established Ethernet scalability. With Fast Ethernet came full-duplex Ethernet. Until, now, all Ethernets worked in half-duplex mode, that is, if there were only two station on a segment, both could not transmit simultaneously. With full-duplex operation, this was now possible.

The next step in the evolution of Ethernet is Gigabit Ethernet. The standard is being developed by the IEEE 802.3z committee.

1.2 The Gigabit Ethernet Alliance (GEA)

In March 1996, the IEEE 802.3 committee approved the 802.3z Gigabit Ethernet Standardization project. At that time as many as 54 companies expressed their intent to participate in the standardization project. The Gigabit Ethernet Alliance was formed in May 1996 by 11 companies: 3Com Corp., Bay Networks Inc., Cisco Systems Inc., Compaq Computer Corp., Granite Systems Inc., Intel Corporation, LSI Logic, Packet Engines Inc., Sun Microsystems Computer Company, UB Networks and VLSI Technology.

The Alliance represents a multi-vendor effort to provide open and inter-operable Gigabit Ethernet products. The objectives of the alliance are:

- supporting extension of existing Ethernet and Fast Ethernet technology in response to demand for higher network

bandwidth.

- developing technical proposals for the inclusion in the standard
- establishment of inter-operability test procedures and processes

Currently membership of the alliance is over 95 companies. This indicates that the emerging standard will be backed by the industry. The alliance is pushing for speedy approval of the standard. So far, the standardization is proceeding without any delays, and is expected to be approved by March 1998.

[Contents.](#)

2. Physical Layer

The Physical Layer of Gigabit Ethernet uses a mixture of proven technologies from the original Ethernet and the ANSI X3T11 Fibre Channel Specification. Gigabit Ethernet is finally expected to support 4 physical media types . These will be defined in 802.3z (1000Base-X) and 802.3ab (1000Base-T).

2.1 1000Base-X

The 1000Base-X standard is based on the Fibre Channel Physical Layer. Fibre Channel is an interconnection technology for connecting workstations, supercomputers, storage devices and peripherals. Fibre Channel has a 4 layer architecture. The lowest two layers FC-0 (Interface and media) and FC-1 (Encode/Decode) are used in Gigabit Ethernet. Since Fibre Channel is a proven technology, re-using it will greatly reduce the Gigabit Ethernet standard development time.

Three types of media are include in the 1000Base-X standard :

- **1000Base-SX** 850 nm laser on multi mode fiber.
- **1000Base-LX** 1300 nm laser on single mode and multi mode fiber.
- **1000Base-CX** Short haul copper "twinax" STP (Shielded Twisted Pair) cable

The cabling distances to be supported are given in Table 1 :

Table 1. Cabling Types and Distances

Cable Type	Distance
Single-mode Fiber (9 micron)	3000 m using 1300 nm laser (LX)
Multi mode Fiber (62.5 micron)	300 m using 850 nm laser (SX)
	550 m using 1300 nm laser (LX)
Multi mode Fiber (50 micron)	550 m using 850nm laser (SX)
	550 m using 1300 nm laser (LX)
Short-haul Copper	25 m

Source : Sun Microsystems (Sun and Gigabit Ethernet White Paper)

2.2 1000Base-T

1000Base-T is a standard for Gigabit Ethernet over long haul copper UTP. The standards committee's goals are to allow up to 25-100 m over 4 pairs of Category 5 UTP. This standard is being developed by the 802.3ab task force and is expected to be completed by early 1999.

[Back to Table of Contents](#)

3. MAC Layer

The MAC Layer of Gigabit Ethernet uses the same CSMA/CD protocol as Ethernet. The maximum length of a cable segment used to connect stations is limited by the CSMA/CD protocol. If two stations simultaneously detect an idle medium and start transmitting, a collision occurs.

Ethernet has a minimum frame size of 64 bytes. The reason for having a minimum size frame is to prevent a station from completing the transmission of a frame before the first bit has reached the far end of the cable, where it may collide with another frame. Therefore, the minimum time to detect a collision is the time it takes for the signal to propagate from one end of the cable to the other. This minimum time is called the *Slot Time*. (A more useful metric is *Slot Size*, the number of bytes that can be transmitted in one Slot Time. In Ethernet, the slot size is 64 bytes, the minimum frame length.)

The maximum cable length permitted in Ethernet is 2.5 km (with a maximum of four repeaters on any path). As the bit rate increases, the sender transmits the frame faster. As a result, if the same frames sizes and cable lengths are maintained, then a station may transmit a frame too fast and not detect a collision at the other end of the cable. So, one of two things has to be done : (i) Keep the maximum cable length and increase the slot time (and therefore, minimum frame size) OR (ii) keep the slot time same and decrease the maximum cable length OR both. In Fast Ethernet, the maximum cable length is reduced to only 100 meters, leaving the minimum frame size and slot time intact.

Gigabit Ethernet maintains the minimum and maximum frame sizes of Ethernet. Since, Gigabit Ethernet is 10 times faster than Fast Ethernet, to maintain the same slot size, maximum cable length would have to be reduced to about 10 meters, which is not very useful. Instead, Gigabit Ethernet uses a bigger slot size of 512 bytes. To maintain compatibility with Ethernet, the minimum frame size is not increased, but the "carrier event" is extended. If the frame is shorter than 512 bytes, then it is padded with extension symbols. These are special symbols, which cannot occur in the payload. This process is called *Carrier Extension*.

3.1 Carrier Extension

Gigabit Ethernet should be inter-operable with existing 802.3 networks. Carrier Extension is a way of maintaining 802.3 minimum and maximum frame sizes with meaningful cabling distances.

For carrier extended frames, the non-data extension symbols are included in the "collision window", that is, the entire extended frame is considered for collision and dropped. However, the Frame Check Sequence (FCS) is calculated only on the original (without extension symbols) frame. The extension symbols are removed before the FCS is checked by the receiver. So the LLC (Logical Link Control) layer is not even aware of the carrier extension. Fig. 1 shows the ethernet frame format when Carrier Extension is used.

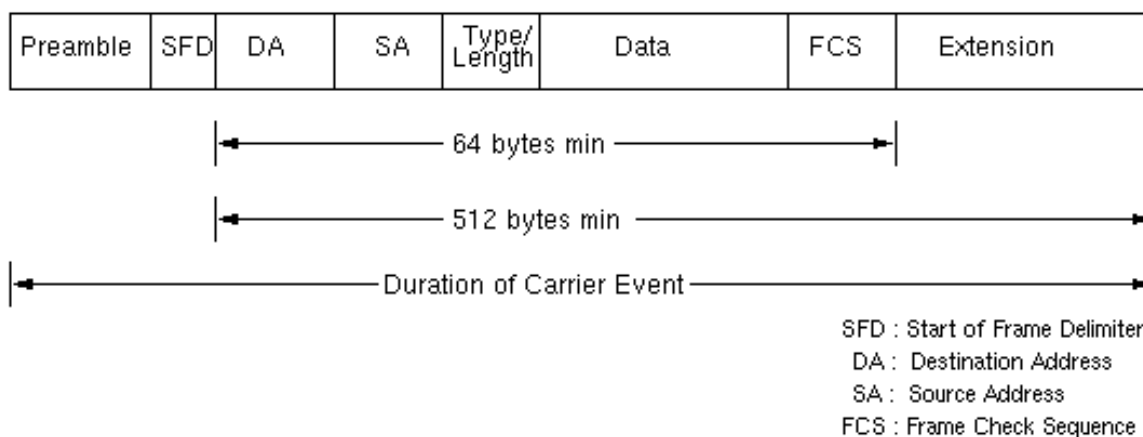


Fig 1. Ethernet Frame Format with Carrier Extension

3.2 Packet Bursting

Carrier Extension is a simple solution, but it wastes bandwidth. Up to 448 padding bytes may be sent for small packets. This results in low throughput. In fact, for a large number of small packets, the throughput is only marginally better than Fast Ethernet.

Packet Bursting is an extension of Carrier Extension. Packet Bursting is "Carrier Extension plus a burst of packets". When a station has a number of packets to transmit, the first packet is padded to the slot time if necessary using carrier extension. Subsequent packets are transmitted back to back, with the minimum Inter-packet gap (IPG) until a burst timer (of 1500 bytes) expires. Packet Bursting substantially increases the throughput. Fig. 2. shows how Packet Bursting works.

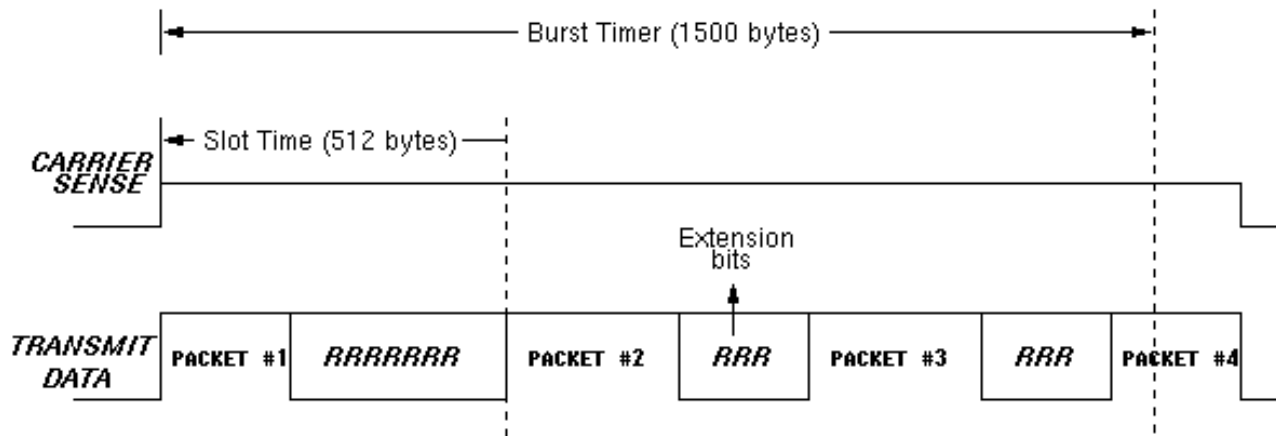


Fig. 2. Packet Bursting

[Back to Table of Contents](#)

4. GMII (Gigabit Media Independent Interface)

The various layers of the Gigabit Ethernet protocol architecture are shown in Fig. 3. The GMII is the interface between the MAC layer and the Physical layer. It allows any physical layer to be used with the MAC layer. It is an extension of the MII (Media Independent Interface) used in Fast Ethernet. It uses the same management interface as MII. It supports 10, 100 and 1000 Mbps data rates. It provides separate 8-bit wide receive and transmit data paths, so it can support both full-duplex as well as half-duplex operation.

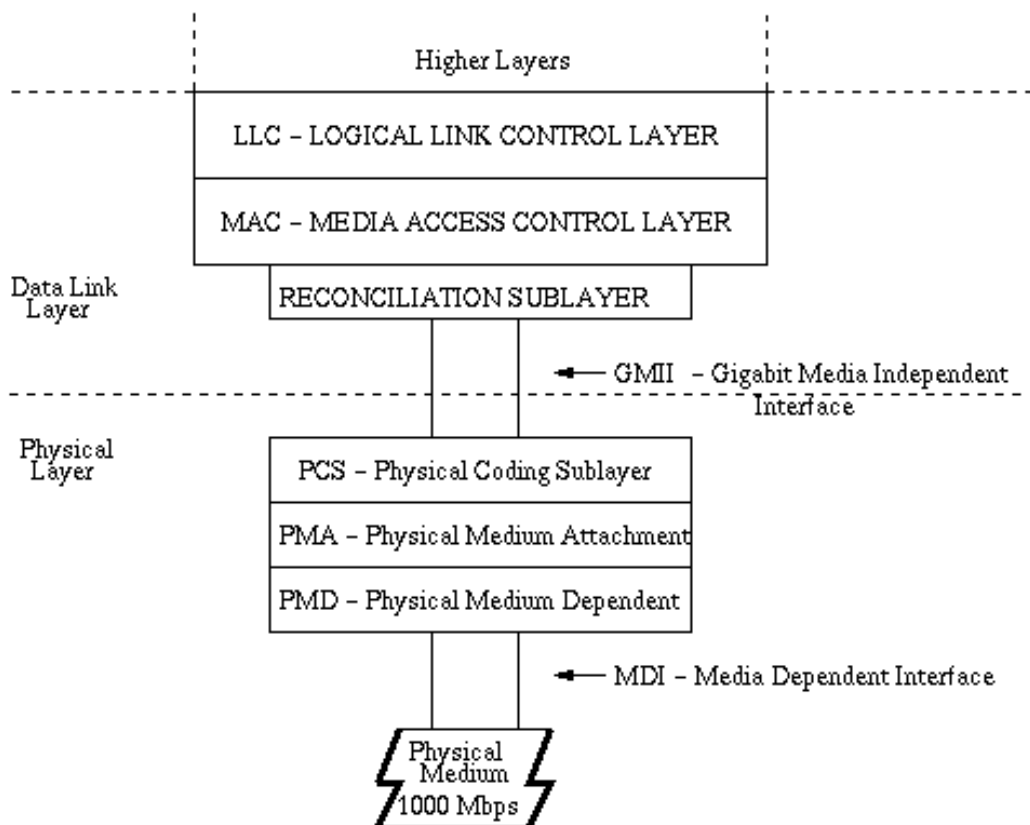


Fig.3 : Gigabit Ethernet Protocol Architecture

The GMII provides 2 media status signals : one indicates presence of the carrier, and the other indicates absence of collision. The Reconciliation Sublayer (RS) maps these signals to Physical Signalling (PLS) primitives understood by the existing MAC sublayer. With the GMII, it is possible to connect various media types such as shielded and unshielded twisted pair, and single-mode and multi mode optical fibre, while using the same MAC controller.

The GMII is divided into three sublayers : PCS, PMA and PMD.

4.1 PCS (Physical Coding Sublayer)

This is the GMII sublayer which provides a uniform interface to the Reconciliation layer for all physical media. It uses 8B/10B coding like Fibre Channel. In this type of coding, groups of 8 bits are represented by 10 bit "code groups". Some code groups represent 8 bit data symbols. Others are control symbols. The extension symbols used in Carrier Extension are an example of control symbols.

Carrier Sense and Collision Detect indications are generated by this sublayer. It also manages the auto-negotiation process by which the NIC (Network Interface) communicates with the network to determine the network speed (10,100 or 1000 Mbps) and mode of operation (half-duplex or full-duplex).

4.2 PMA (Physical Medium Attachment)

This sublayer provides a medium-independent means for the PCS to support various serial bit-oriented physical media. This layer serializes code groups for transmission and deserializes bits received from the medium into code groups.

4.3 PMD (Physical Medium Dependent)

This sublayer maps the physical medium to the PCS. This layer defines the physical layer signalling used for various media. The **MDI (Medium Dependent Interface)**, which is a part of PMD is the actual physical layer interface. This layer defines the actual physical attachment, such as connectors, for different media types.

[Back to Table of Contents](#)

5. Buffered Distributor

Ethernet today supports full-duplex media, physical layer as well MAC layer. However it still supports half-duplex operation to maintain compatibility. A new device has been proposed which provides hub functionality with full duplex mode of operation. It is called various names such as *Buffered Distributor*, *Full Duplex Repeater* and *Buffered Repeater*. The term "Buffered Distributor" is used for all these devices in the following discussion.

The basic principle is that CSMA/CD is used as the access method to the network and not to the link. A Buffered Distributor is a multi-port repeater with full-duplex links.

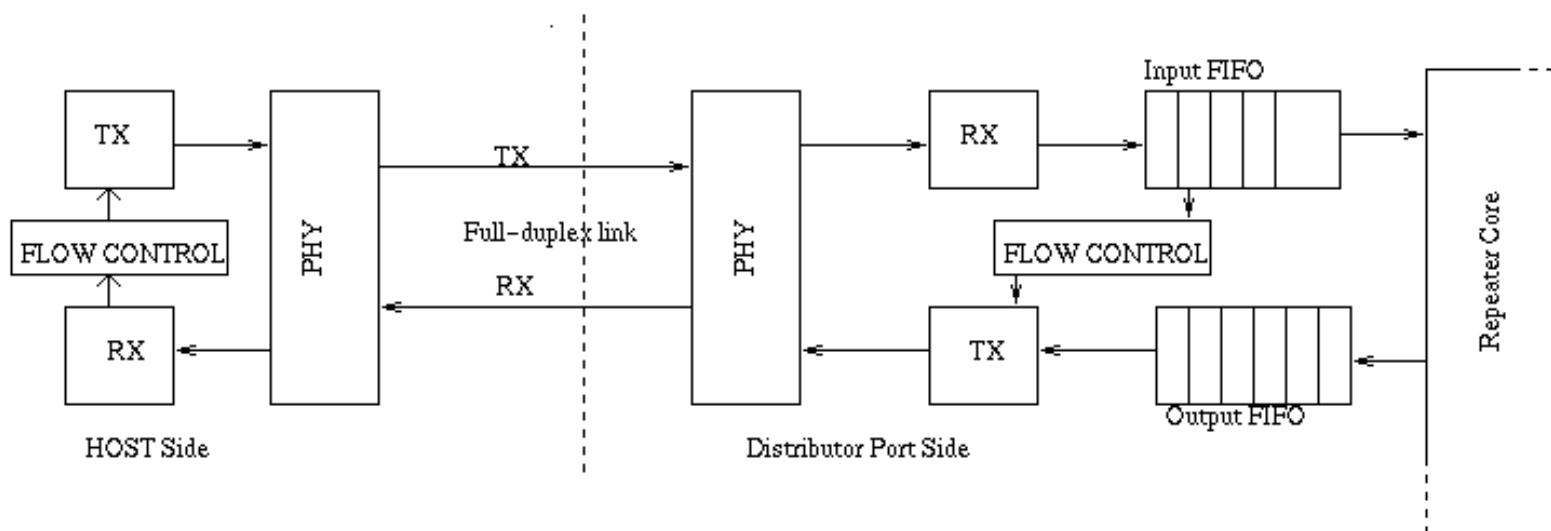


Fig. 4 : Buffered Distributor Architecture

Each port has an input FIFO queue and an output FIFO queue. A frame arriving to an input queue is forwarded to all output queues, except the one on the incoming port. Within the distributor, CSMA/CD arbitration is done to forward the frames to output queues.

Since collisions can no longer occur on links, the distance restrictions no longer apply. The only restriction on cabling distances is the characteristics of the physical medium, and not the CSMA/CD protocol.

Since the sender can flood the FIFO, frame based flow control is used between the port and the sending station. This is defined in the 802.3x standard and already used in Ethernet switches.

The motivation behind development of the Buffered Distributor is its cost compared to a Gigabit switch and not a need to accommodate half duplex media. The Buffered Distributor provides full duplex connectivity, just like a switch, yet it is not so expensive, because it is just an extension of a repeater.

[Back to Table of Contents](#)

6. Topologies

This section discusses the various topologies in which Gigabit Ethernet may be used. Gigabit Ethernet is essentially a "campus technology", that is, for use as a backbone in a campus-wide network. It will be used between routers, switches and hubs. It can also be used to connect servers, server farms (a number of server machines bundled together), and powerful workstations.

Essentially, four types of hardware are needed to upgrade an exiting Ethernet/Fast Ethernet network to Gigabit Ethernet :

- Gigabit Ethernet Network Interface Cards (NICs)
- Aggregating switches that connect a number of Fast Ethernet segments to Gigabit Ethernet

Gigabit Ethernet

- Gigabit Ethernet switches
- Gigabit Ethernet repeaters (or Buffered Distributors)

The five most likely upgrade scenarios are given below :

6.1 Upgrading server-switch connections

Most networks have centralized file servers and compute servers. A server gets requests from a large number of clients. Therefore, it needs more bandwidth. Connecting servers to switches with Gigabit Ethernet will help achieve high speed access to servers. . This is perhaps the simplest way of taking advantage of Gigabit Ethernet.

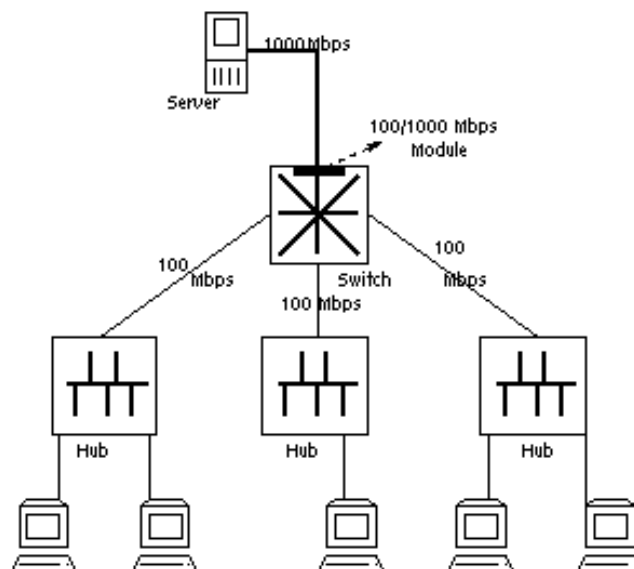


Fig. 5. Server-Switch Connection

6.2 Upgrading switch-switch connections

Another simple upgrade involves upgrading links between Fast Ethernet switches to Gigabit Ethernet links between 100/1000 Mbps switches.

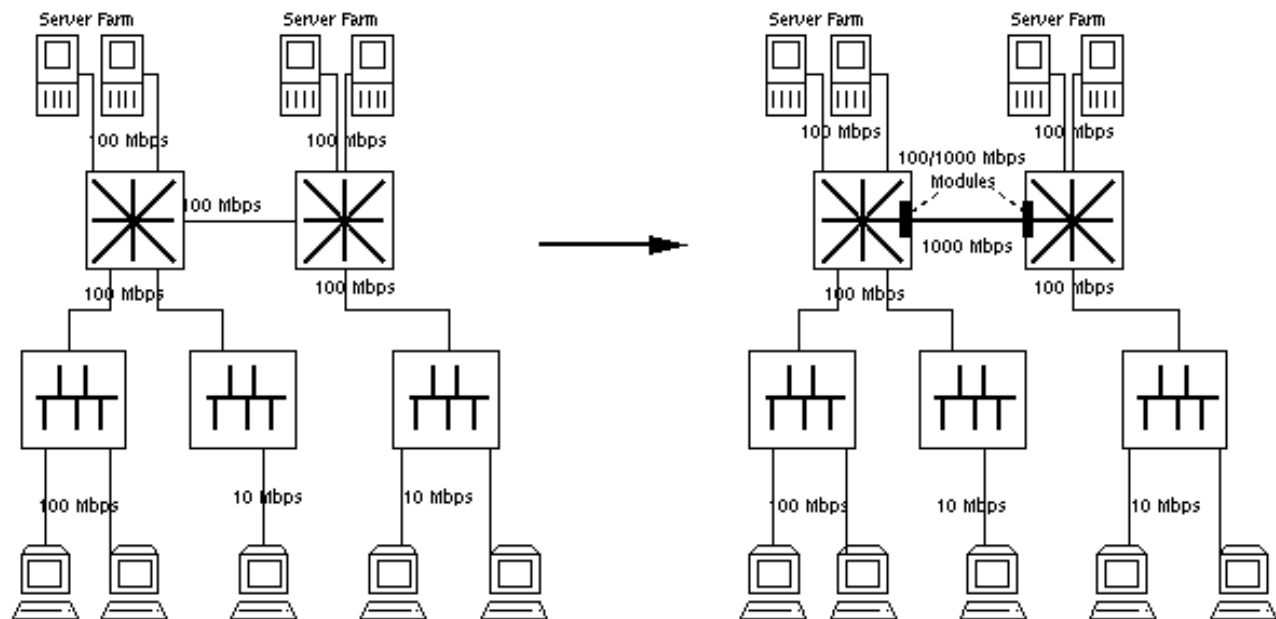


Fig. 6. Upgrading Switch-Switch connections

6.3 Upgrading a Fast Ethernet backbone

A Fast Ethernet backbone switch aggregates multiple 10/100 Mbps switches. It can be upgraded to a Gigabit Ethernet switch which supports multiple 100/1000 Mbps switches as well as routers and hubs which have Gigabit Ethernet interfaces. Once the backbone has been upgraded, high performance servers can be connected directly to the backbone. This will substantially increase throughput for applications which require high bandwidth.




Fig. 7. Upgrading the Backbone

6.4 Upgrading a Shared FDDI Backbone

Fiber Distributed Data Interface (FDDI) is a common campus or building backbone technology. An FDDI backbone can be upgraded by replacing FDDI concentrators or Ethernet-to-FDDI routers by a Gigabit Ethernet switch or repeater.

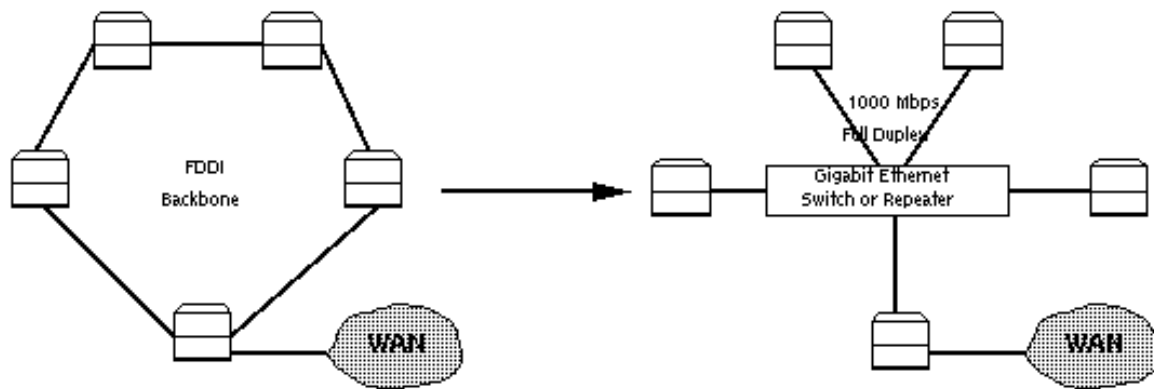


Fig. 8. Upgrading a FDDI Backbone

6.5 Upgrading High Performance Workstations

As workstations get more and more powerful, higher bandwidth network connections are required for the workstations. Current high-end PCs have buses which can pump out more than 1000 Mbps. Gigabit Ethernet can be used to connect such high speed machines.




Fig. 9. Upgrading High Performance Workstations

[Back to Table of Contents](#)

7. ATM vs. Gigabit Ethernet

When ATM (Asynchronous Transfer Mode) was introduced, it offered 155 Mbps bandwidth, which was 1.5 times faster than Fast Ethernet. ATM was ideal for new applications demanding a lot of bandwidth, especially multimedia. Demand for ATM continues to grow for LAN's as well as WAN's.

On the one hand, proponents of ATM try to emulate Ethernet networks via LANE (LAN Emulation) and IPOA (IP over ATM). On the other, proponents of Ethernet/IP try to provide ATM functionality with RSVP (Resource Reservation Protocol) and RTSP (Real-time Streaming Transport Protocol). Evidently, both technologies have their desirable features, and advantages over the other. It appears that these seemingly divergent technologies are actually converging.

ATM was touted to be the seamless and scaleable networking solution - to be used in LANs, backbones and WANs alike. However, that did not happen. And Ethernet, which was for a long time restricted to LANs alone, evolved into a scalable technology.

As Gigabit Ethernet products enter the market, both sides are gearing up for the battle. Currently, most installed workstations and personal computers do not have the capacity to use these high bandwidth networks. So, the imminent battle is for the backbones, the network connections between switches and servers in a large network.

Gigabit Ethernet seems to be ready to succeed. It is backed by the industry in the form of the Gigabit Ethernet Alliance. The standardization is currently on schedule. Pre-standard products with claims of inter-operability with standardized products have already hit the market. Many Fast Ethernet pre-standard products were inter-operable with the standard. So it is expected that most pre-standard Gigabit Ethernet products will also be compatible with the standard. This is possible because many of the companies that have come out with products are also actively participating in the standardization process.

ATM still has some advantages over Gigabit Ethernet :

- ATM is already there. So it has a head start over Gigabit Ethernet. Current products may not support gigabit speeds, but faster versions are in the pipeline.
- ATM is better suited than Ethernet for applications such as video, because ATM has QOS (Quality of Service) and different services available such as CBR (constant bit rate) which are better for such applications. Though the IETF (Internet Engineering Task Force, the standards body for internet protocols) is working on RSVP which aims to provide QOS on Ethernet, RSVP has its limitations. It is a "best effort" protocol, that is, the network may acknowledge a QOS request but not deliver it. In ATM it is possible to guarantee QOS parameters such as maximum delay in delivery.

Gigabit Ethernet has its own strengths :

- The greatest strength is that it is Ethernet. Upgrading to Gigabit Ethernet is expected to be painless. All applications that work on Ethernet will work on Gigabit Ethernet. This is not the case with ATM. Running current applications on ATM requires some amount of translation between the application and the ATM layer, which means more overhead.
- Currently, the fastest ATM products available run at 622 Mbps. At 1000 Mbps, Gigabit Ethernet is almost twice as fast.

It is not clear whether any one technology will succeed over the other. It appears that sooner or later, ATM and Ethernet will complement each other and not compete.

[Back to Table of Contents](#)

8. Summary

Gigabit Ethernet is the third generation Ethernet technology offering a speed of 1000 Mbps. It is fully compatible with existing Ethernets, and promises to offer seamless migration to higher speeds. Existing networks will be able to upgrade their performance without having to change existing wiring, protocols or applications. Gigabit Ethernet is expected to give existing high speed technologies such as ATM and FDDI a run for their money. The IEEE is working on a standard for Gigabit Ethernet, which is expected to be out by the beginning of 1998. A standard for using Gigabit Ethernet on twisted pair cable is expected by 1999.

[Back to Table of Contents](#)

9. Bibliography and Links

Annotated Bibliography

1. IEEE Draft P802.3z/D3 "Media Access Control (MAC) Parameters, Physical Layer, Repeater and Management Parameters for 1000 Mb/s operation", June 1997
This is the draft document of the 802.3z standard. It is available from IEEE on request. The [IEEE standards](#) site has

information on ordering IEEE standards and draft documents at <http://standards.ieee.org/faqs/order.html>

2. Gigabit Ethernet Alliance "Gigabit Ethernet : White Paper", Aug 1996,
<http://www.gigabit-ethernet.org/technology/whitepapers/gige/>
A whitepaper from the GEA. Gives comparison of technologies like FDDI and ATM with Gigabit Ethernet.
3. Joe Skorupa, George Prodan, "Battle of the Backbones: ATM vs. Gigabit Ethernet", Data Communications, April 1997,
<http://www.data.com/tutorials/backbones.html>
This article features arguments for and against Gigabit Ethernet vs. ATM.
4. Michelle Rae McLean, "Gigabit Speeds Mandate Fix", LAN Times, Sep 1996,
<http://www.wcmh.com/96sep/609c029a.html>
This article discusses problems in migrating to Gigabit Ethernet.

Gigabit Ethernet Related WWW Links

1. <http://www.gigabit-ethernet.org/>
Gigabit Ethernet Alliance Home Page.
2. http://www.pcwebopaedia.com/Gigabit_Ethernet.htm
Gigabit Ethernet page at PC Webopaedia. Contains Definitions and Links
3. <http://www.ots.utexas.edu:8080/ethernet/gigabit.html>
This page contains some Gigabit Ethernet information and links.

[Back to Table of Contents](#)

[Vijay Moorthy](#), Aug 14, 1997