WHITE PAPER

November 1996

Compaq Computer Corporation

Communications Products Group

Networking Products Division

CONTENTS

COMPAQ

Gigabit Ethernet Technology

The increasing complexity of desktop computing and bandwidth-hungry applications such as multimedia, medical imaging, CAD/CAM, and pre-press processing are demanding unprecedented LAN speeds. Out of all the available high-speed LAN technologies, Fast Ethernet, based on the popular 10Base-T Ethernet, has become the leading choice in meeting those demands by offering greater bandwidth and improved client/server response times. Now, however, the increased use of Fast Ethernet connections at the server and desktop is creating a need for even higher-speed network technology at the backbone.

An IEEE (802.3z) task force is developing a new Ethernet standard called "Gigabit Ethernet" that runs at 1000 Mb/s. Gigabit Ethernet employs the same CSMA/CD protocol, frame format, and frame size as its predecessors, plus interoperability and backward compatibility with the installed Ethernet. Gigabit Ethernet will also support existing applications, Network Operating Systems (NOS), and network management.

Gigabit Ethernet will initially be deployed as a backbone interconnection between 10/100Base-T switches and as a connection to high-performance servers.

In November 1995, Compaq Computer Corporation recognized the need for Gigabit Ethernet technology and proposed the basic architecture to the IEEE 802 committee. In early 1996, the IEEE formed the Gigabit Ethernet Task Force (802.3z) whose commission was to draft a Gigabit Ethernet standard. Compaq has contributed significant technical input and resources to expedite the development of this standard.

In April, 1996, the Gigabit Ethernet Alliance was formed by Compaq Computer Corporation and a few other leading networking and computer companies to promote the Gigabit Ethernet technology and to ensure multi-vendor interoperability. Since then the Alliance membership has grown to more than 72 companies.

This paper presents an overview of Gigabit Ethernet technology and provides the current status of the IEEE 802.3z standard.

NOTICE

The information in this publication is subject to change without notice.

THE COMPETITIVE INFORMATION CONTAINED IN THIS PUBLICATION IS BASED ON DATA AVAILABLE FROM PUBLIC SOURCES AT THE TIME OF PUBLICATION. COMPAQ COMPUTER CORPORATION SHALL NOT BE RESPONSIBLE FOR ANY INACCURACIES, ERRORS, OR OMISSIONS IN INFORMATION CONTAINED HEREIN, INCLUDING, BUT NOT LIMITED TO, INFORMATION OBTAINED FROM THIRD PARTY SOURCES, SUCH AS PUBLICATIONS OF OTHER COMPANIES, THE PRESS, OR COMPETITIVE DATA ORGANIZATIONS.

THIS PUBLICATION IS MADE AVAILABLE ON AN "AS IS" BASIS AND COMPAQ SPECIFICALLY DISCLAIMS ALL ASSOCIATED WARRANTIES, WHETHER EXPRESS OR IMPLIED. IN NO EVENT WILL COMPAQ BE LIABLE FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE USE OF OR RELIANCE ON THE MATERIAL CONTAINED IN THIS PUBLICATION, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. COMPAQ RESERVES THE RIGHT TO MAKE IMPROVEMENTS OR CHANGES TO THIS PUBLICATION AND THE PRODUCTS AND SERVICES HEREIN DESCRIBED AT ANY TIME, WITHOUT NOTICE OR OBLIGATION.

This publication does not constitute an endorsement of the product or products that were tested. The configuration or configurations tested or described may or may not be the only available solution. This test is not a determination of product quality or correctness, nor does it ensure compliance with any federal, state or local requirements. Compaq does not warrant products other than its own and then, only as strictly stated in Compaq product warranties.

Gigabit Ethernet Technology

November 1996
Document Number 130A/1196

GIGABIT ETHERNET OVERVIEW

Gigabit Ethernet represents the next generation of the Ethernet standard. Operating at 1000 Mb/s, Gigabit Ethernet ensures maximum compatibility with the installed base of 10/100 Mb/s Ethernet by adapting the existing CSMA/CD MAC protocol and Management Information Base (MIB). Gigabit Ethernet supports full-duplex operation for point-to-point connections (switched networks) and half-duplex operation for shared connections. Initial products will operate only over optical fiber. Future products will provide copper solutions.

APPLICATIONS AND TOPOLOGIES

Just as in 10 and 100 Mb/s Ethernet, Gigabit Ethernet will be switched, routed, and shared. All of today's internetworking technologies, as well as emerging technologies such as IP-specific switching and Layer 3 switching, are fully compatible with Gigabit Ethernet. (The emergence of new protocols, such as RSVP, 802.1Q, and/or 802.1p, will provide multimedia support as well as VLAN and explicit priority support for Gigabit and 100Base-T Ethernet). The initial application for Gigabit Ethernet will be in collapsed backbones between Fast Ethernet switches, repeaters, routers, and servers.

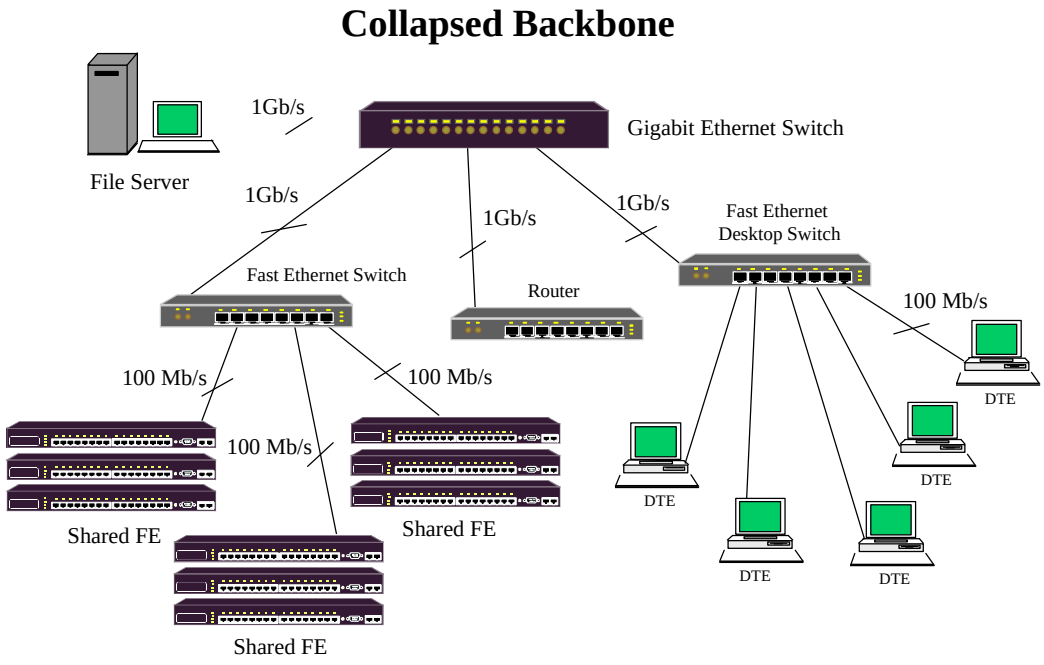

Figure 1. Gigabit Ethernet Application Example in a Collapsed Backbone

COST GOALS

The IEEE 802.3z goal is to provide Gigabit Ethernet connections at only two to three times the cost of the 100Base-FX interface. Based on Dell'Oro Group information and IEEE goals, the estimated cost per port for shared Gigabit Ethernet is $920 to $1400 and for switched Gigabit Ethernet, $1850 to $2800.

GIGABIT ETHERNET CHARACTERISTICS

A goal of the Gigabit Ethernet Task Force is that the Gigabit Ethernet standard meet the following criteria:

- Run at 1000 Mb/s at the MAC/PLS service interface
- Adopt the IEEE 802.3 Ethernet frame format
- Meet the IEEE 802 functional requirement
- Provide simple forwarding between 10 Mb/s, 100 Mb/s, and 1000 Mb/s Ethernet
- Preserve the minimum and maximum frame size of the current IEEE 802.3 standard
- Provide full- and half-duplex operation
- Support star-wired topologies
- Utilize the CSMA/CD access method with support for at least one repeater/collision domain
- Utilize ANSI Fiber Channel's FC-1 and FC-0 as a basis for work
- Provide a family of physical layer specifications that support a link distance of
    - At least 500 meters on multimode fiber
    - At least 25 meters on copper (100 meters preferred)
    - At least 3,000 meters on single mode fiber
- Support a maximum collision domain diameter of 200 meters
- Support media selected from ISO/IEC 11801
- Adopt flow control based on the IEEE 802.3x standard
- Specify an optional GMII (Gigabit Ethernet Media Independent Interface)

OSI REFERENCE MODEL

Gigabit Ethernet fits seamlessly into the OSI reference model, altering only those layers dealing with direct connection to the physical layer. This guarantees interoperability with existing Ethernet and Fast Ethernet applications. The details on each sublayer are described below.

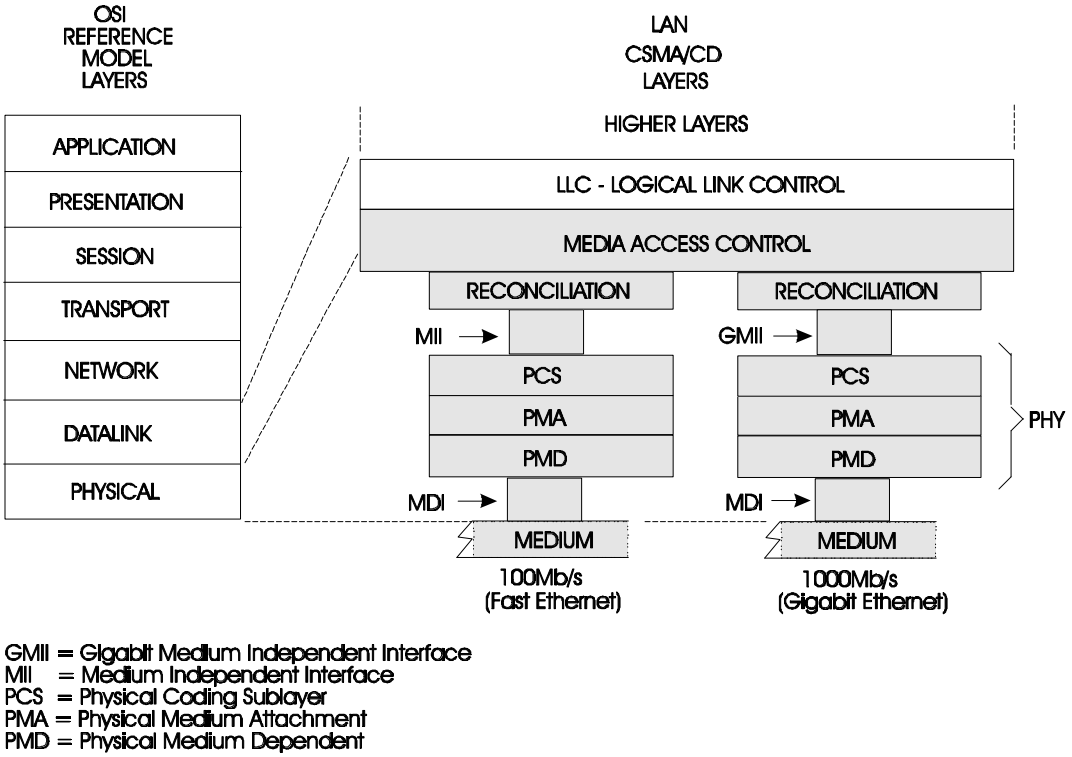

Figure 2: Fast Ethernet and Gigabit Ethernet Relationship to the ISO Open System Interconnection (OSI) Reference Model

Logical Link Control

The Logical Link Control layer is unchanged for the Gigabit Ethernet standard. All upper layer protocols (including TCP/IP, SPX/IPX, etc.) can run over Gigabit Ethernet.

Media Access Controller (MAC)

Gigabit Ethernet MAC is a scaled-up version of the Fast Ethernet MAC with an effective data rate of 1000 Mb/s. The MAC supports both full-duplex and half-duplex operation. In full-duplex mode, the MAC uses the frame-based flow control (as defined in IEEE 802.3x). In half-duplex mode, the MAC supports the CSMA/CD access method.

Reconciliation Sublayer and Gigabit Media Independent Interface (GMII)

The reconciliation sublayer and Gigabit Media Independent Interface interconnect the MAC sublayer and physical layer (PHY) entities. This includes an 8-bit data bus operating at 125 MHz plus control signals such as transmit/receive clocks, carrier indicators, and error conditions.

Physical Layer (PHY)

The Gigabit Ethernet standard will support a family of physical layer implementations for various media as indicated below:

- 62.5 μm and 50 μm multimode fiber (MMF)
- Single mode fiber (SMF)
- Short copper links (25 meters)
- Horizontal copper (100 meters)

The physical layers for optical and short copper links are based on the ANSI fiber channel FC-1 and FC-0 standard. The physical layer for long copper links is not yet defined.

Physical Coding Sublayer (PCS)

The physical coding sublayer contains the 8B/10B encoder/decoder (adopted from the FC-1 fiber channel specification) for use with optical or short copper links. The 8B/10B encoding scheme transmits eight bits as a 10-bit code group. The transmission code is DC balanced to support the electrical requirements of the receiving units. 8B/10B encoding ensures that enough transitions are present in the serial bit stream to improve clock recovery and error detection at the receiver. The PCS for horizontal copper is not yet defined.

Physical Medium Attachment (PMA)

The physical medium attachment performs 10-bit serialize/deserialize functions (analogous to the ANSI 10-bit SERDES chip). It receives 10-bit encoded data at 125 MHz from the PCS and delivers serialized data to the PMD block. It also receives serialized data from the PMD and delivers deserialized 10-bit data to the PCS.

Physical Medium Dependent (PMD)

The physical medium dependent sublayer provides the media transceivers and connectors for various media. Both 780 and 1300 nm wavelength optical drivers are supported. Laser drivers must confirm to Class 1 laser safety operation, which guarantees no eye damage.

Copper solutions include short copper link (25 meter) over Category 5 Unshielded Twisted Pair (Cat. 5 UTP) cable and Twinax cables and Horizontal Copper (100 meters) over (Cat. 5 UTP) cable.

FUNCTIONAL ELEMENTS

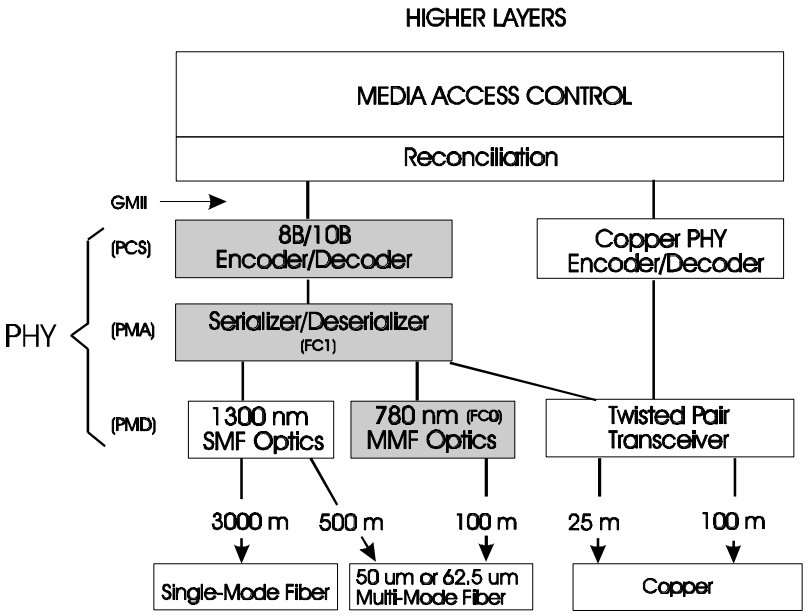

Figure 3: Functional Elements of Gigabit Ethernet

Notes:
Shaded areas indicate functions that are leveraged from the current fiber channel standard. Short-length copper implementation is identical to optical fiber implementation except copper drivers are used in place of optical drivers.

Early products for Gigabit Ethernet will use the available fiber channel FC-0 (fiber driver/receiver) and FC-1 (serializer/de-serializer) for the physical layer implementation. These devices need to operate at 1.25 GHz (Gigabit Ethernet signaling rate) rather than 1.06 GHz (FC rate).

CABLING TYPES AND DISTANCES

Cable Type	Optical Driver	Maximum Distance	Relative Cost
Multimode Fiber (50 or 62.5 µm)	780 nm	100 meters	Moderate
Multimode Fiber (50 or 62.5 µm)	1300 nm	500 meters	High
Single Mode Fiber	1300 nm	3 kilometers	High
Short-length Copper	N/A	25 meters	Low
Horizontal Copper	N/A	100 meters	Moderate

Table 1. Proposed Cabling Types and Distances

Current proposals for copper solutions include CAT5 UTP and Twinax cables. Distances indicated represent maximum length for a point-to-point connection. For a shared (CSMA/CD) network, network diameter should not exceed 200 meters regardless of media type.

SWITCHED NETWORK

Operating in full-duplex mode, a Gigabit Ethernet switch can provide an aggregate 2 Gb/s bandwidth to each node on the network. Congestion at the switch is controlled through frame-based flow control (IEEE 802.3x standard). A congested node on the switch can send a flow control message (a 64-byte packet with unique ID type) to the attached station to stop sending packets for a specified period of time. The attached station stops further transmission until the specified period of time has passed or until it receives a new packet with a delay time of zero. Flow control packets can operate between two devices only on a point-to-point connection. It cannot operate between two devices separated through a switch. The maximum cable length that Gigabit Ethernet can use in a switched network is limited by the media type and driver capability (up to 3,000 meters on single mode fiber) and is not constrained to the 200-meter network diameter limitation associated with a CSMA/CD network.

SHARED NETWORK

The IEEE 802.3z task force is evaluating the following two proposals for shared network support in Gigabit Ethernet:

Carrier Extension

In this proposal, a repeater device (similar to a 10/100Base-T repeater) will be supported, and a carrier extension mechanism will be implemented to account for the problem of scaling CSMA/CD to a 1000 Mb/s speed and maintaining a 200-meter network diameter. If a received packet is less than 512 bytes, the receive carrier sense signal is extended (to 512 bytes) to ensure the detection of a delayed collision signal across the 200-meter network.

Buffered Repeater

In this proposal, CSMA/CD operation on the wire is totally eliminated. A low-end switch, called a *buffered repeater*, will be standardized to replace the repeater function. As the name implies, the buffered repeater will store the incoming packet on a port and broadcast the packet to all other ports. All ports on the buffered repeater will operate in full-duplex mode. Frame-based flow control (802.3x) will be used to handle congestion. The buffered repeater will not perform address filtering or support spanning tree management, therefore it will be less expensive than a switch. Since there is no collision overhead, network bandwidth is better utilized. Table 2 compares the costs and features of Gigabit Ethernet repeaters and switches.

	Basic Repeater	Buffered Repeater	Switch
Relative Cost	$	$	$
Full/Half duplex	HDX (CSMA/CD)	FDX	FDX
Complexity	Low	Moderate	High
Distance (Network Diameter)	200m	Length Independent	Length Independent
Network Performance	<1Gb/s	1Gb/s	>1Gb/s

Table 2. Feature Comparison Gigabit Ethernet Repeaters and Switches

STANDARDIZATION TIMELINE

The IEEE 802.3z has adopted the following timeline for completing the Gigabit Ethernet standard:

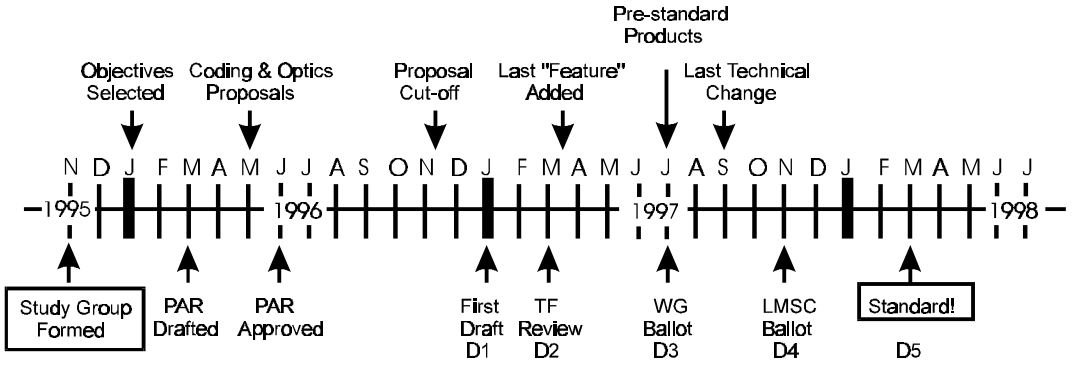

Figure 3. Standardization Timeline

GIGABIT ETHERNET ALLIANCE

The Gigabit Ethernet Alliance was formed in May, 1996, by 11 companies, including Compaq Computer Corporation, to promote Gigabit Ethernet technology. Since then, the membership has grown to more than 75 companies.

The Alliance has the following primary objectives:

- Fully support the Gigabit Ethernet standards activities being conducted in the IEEE 802.3z working group
- Contribute technical resources to facilitate convergence and consensus on technical specifications
- Provide resources to establish and demonstrate product interoperability
- Provide communications between potential suppliers and consumers of Gigabit Ethernet products. The Alliance maintains a web site with general information on this technology and the status of the standards development. The web site is located at: http://www.gigabit-ethernet.org

CONCLUSION

Over the last 10 years, the growth of client/server applications and the increasing complexity of desktop computing has significantly boosted the demand for network bandwidth.

Compaq Computer Corporation, as the leading provider of desktops and servers and the driving force in client/server computing today, is uniquely positioned to understand the bandwidth needs of the evolving network. Gigabit Ethernet is a logical extension of today's Ethernet and Fast Ethernet standards that will provide maximum performance for the client/server networks of tomorrow.

Gigabit Ethernet products are expected to be available starting in July, 1997, and the Gigabit Ethernet (IEEE 802.3z) standard is scheduled to be completed by January, 1998.