

Gigabit Ethernet

accelerating

the standard

for speed

Whitepaper

**Gigabit
Ethernet
Alliance**

1Gb

Copyright 1998
Gigabit Ethernet Alliance

Introduction

Modern organizations depend upon their local-area networks (LANs) to provide connectivity for a growing number of complex, mission-critical desktop computing applications. As the volume of network traffic increases, however, the bandwidth offered by a typical 10 Mbps Ethernet LAN quickly becomes inadequate to maintain acceptable performance for a growing number of desktop/server computing environments. These traffic jams are fueling the need for higher-speed networks.

Among the high-speed LAN technologies available today, Fast Ethernet, or 100BASE-T, has become the leading choice. Building on the near-universal acceptance of 10BASE-T Ethernet, Fast Ethernet technology provides a smooth, non-disruptive evolution to 100 Mbps performance. The growing use of 100BASE-T connections to servers and desktops, however, is creating a clear need for an even higher-speed network technology at the backbone and server level. Ideally, this technology should also provide a smooth upgrade path, be cost effective and not require retraining.

The most appropriate solution is Gigabit Ethernet. Gigabit Ethernet will provide 1 Gbps bandwidth for campus networks with the simplicity of Ethernet at lower cost than other technologies of comparable speed. It offers a natural upgrade path for current Ethernet installations, leveraging existing end stations, management tools and training.

Gigabit
1000 Ethernet
1000 MBPS
Highly Reliable
LOW COST

Gigabit Ethernet employs the same Carrier Sense Multiple Access with Collision Detection (CSMA/CD) protocol, same frame format and same frame size as its predecessors. For the vast majority of network users, this means their existing network investment can be extended to gigabit speeds at reasonable initial cost without the need to re-educate their support staffs and users, and without the need to invest in additional protocol stacks or middleware. The result is low cost of ownership for users.

Because of these attributes, as well as support for full-duplex operation, Gigabit Ethernet is an ideal backbone interconnect technology for use between 10/100BASE-T switches, as a connection to high-performance servers and as an upgrade path for future high-end desktop computers requiring more bandwidth than 100BASE-T can offer.

This paper presents an overview of Gigabit Ethernet, including Ethernet's current position in the industry, applications driving the need for more bandwidth, Gigabit Ethernet as a strategic evolution for networks, Gigabit Ethernet technology fundamentals and network migration scenarios.

Ethernet: The Dominant Network Technology

Ethernet technology is ubiquitous. More than 85 percent of all installed network connections were Ethernet by the end of 1997 according to IDC. This represents over 118 million interconnected PCs, workstations and servers. The remaining network connections are a combination of Token Ring, Fiber Distributed Data Interface (FDDI), Asynchronous Transfer Mode (ATM) and other protocols. All popular operating systems and applications are Ethernet-compatible, as are upper-layer protocol stacks such as Transmission Control Protocol/Internet Protocol (TCP/IP), IPX, NetBEUI and DECnet.

1998 will be a milestone year for Ethernet network equipment as the technology captures 86 percent of shipments. IDC projects Ethernet network interface card (NIC) shipments will exceed 48 million units, and Ethernet hub shipments will exceed 48 million ports. In contrast, ATM, FDDI/CDDI and Token Ring network interface card shipments combined are expected to reach just 5 million in 1998, 10 percent of the total. ATM, FDDI/CDDI and Token Ring hub ports are expected to be 4 million, 7 percent of the total. IDC projects that Ethernet dominance will continue beyond 1999. Several factors have contributed to making Ethernet the most popular network technology in use today.

Network Reliability Highly reliable networks are critical to the success of the enterprise, so ease of installation and support are primary considerations in the choice of network technology. Since the introduction in 1986 of star-wired 10BASE-T hubs, structured wiring systems have continued to evolve and hubs and switches have become increasingly reliable. Today, Ethernet networks are rapidly approaching the reliability level associated with their telephone ancestors, and are relatively simple to understand and administer.

Availability of Management and Troubleshooting Tools Management tools for Ethernet, made possible by widespread adoption of management standards including Simple Network Management Protocol (SNMP) and its successors, allow an administrator to view the status of all desktops and network elements, including redundant elements, from a central station. Ethernet troubleshooting tools span a range of capabilities, from simple link indicator lights to sophisticated network analyzers. As a result of Ethernet's popularity, large numbers of people have been trained on its installation, maintenance and troubleshooting.

Scalability The Fast Ethernet standard, approved in 1995, established Ethernet as a scalable technology. Now, the development of Gigabit Ethernet, approved in 1998, extends the scalability of Ethernet even further. Now Ethernet scales from 10 to 100 to 1000 Mbps.

Low Cost IDC's analysis of Ethernet and Fast Ethernet indicates a rapid decrease in price per port for both technologies as shown in Table 1. There has been a rapid overall decrease in the per-port price for Fast Ethernet and Ethernet products, and for shared hub ports, the difference between their costs is narrowing. Over time, Gigabit Ethernet per port costs can be expected to experience similar price decreases.

Availability of
Tools and Management
Technology

Intelligent Hub Per-Port Average Price	1996	1998
Ethernet	\$ 53	\$ 46
Fast Ethernet	\$160	\$ 73
Ratio	3.0	1.6

Highly
Reliable

Switch Per-Port Average Price	1996	1998
Ethernet	\$421	\$212
Fast Ethernet	\$716	\$432
Ratio	1.7	1.7

IEEE 802.3z

NIC Per-Port Average Price	1996	1998
Ethernet	\$107	\$ 59
Fast Ethernet	\$137	\$ 71
Ratio	1.3	1.2

New Capabilities

Source: IDC, 1997

Table 1. Ethernet Network Equipment Price Trends

Applications Driving Network Growth

As new and existing network applications evolve to embrace high-resolution graphics, video and other rich media data types, pressure is growing at the desktop, the server, the hub and the switch for increased bandwidth. Table 2 summarizes the applications and their impact on the network.

Application	Data Types/Size	Network Traffic Implication	Network Need
Scientific Modeling, Engineering	<ul style="list-style-type: none"> Data files 100's of megabytes to gigabytes 	<ul style="list-style-type: none"> Large files increase bandwidth required 	<ul style="list-style-type: none"> Higher bandwidth for desktops, servers and backbone
Publications, Medical Data Transfer	<ul style="list-style-type: none"> Data files 100's of megabytes to gigabytes 	<ul style="list-style-type: none"> Large files increase bandwidth required 	<ul style="list-style-type: none"> Higher bandwidth for desktops, servers and backbone
Internet/Intranet	<ul style="list-style-type: none"> Data files now Audio now Video is emerging High transaction rate Large files, 1 MB to 100 MB 	<ul style="list-style-type: none"> Large files increase bandwidth required Low transmission latency High volume of data streams 	<ul style="list-style-type: none"> Higher bandwidth for servers and backbone Low latency
Data Warehousing, Network Backup	<ul style="list-style-type: none"> Data files Gigabytes to terabytes 	<ul style="list-style-type: none"> Large files increase bandwidth required Transmitted during fixed time period 	<ul style="list-style-type: none"> Higher bandwidth for servers and backbone Low latency
Desktop Video Conferencing, Interactive Whiteboarding	<ul style="list-style-type: none"> Constant data stream 1.5 to 3.5 Mbps at the desktop 	<ul style="list-style-type: none"> Class of service reservation High volume of data streams 	<ul style="list-style-type: none"> Higher bandwidth for servers and backbones Low latency Predictable latency

Table 2. Summary of Applications Driving Network Growth

Many of these applications require the transmission of large files over the network. Scientific applications demand ultra-high bandwidth networks to communicate 3-D visualizations of complex objects ranging from molecules to aircraft. Magazines, brochures and other complex, full-color publications prepared on desktop computers are transmitted directly to digital-input printing facilities. Many medical facilities are transmitting complex images over LAN and WAN links, enabling the sharing of expensive equipment and specialized medical expertise. Engineers are using electronic and mechanical design automation tools to work interactively in distributed development teams, sharing files in the hundreds of gigabytes.

Many companies are now employing Internet technologies to build private intranets, enabling users in an organization to go beyond electronic mail and access critical data through familiar Web browsers, opening the door to a new generation of multimedia client/server applications. While intranet traffic is currently composed primarily of text, graphics and images, this is expected to expand in the near future to include more bandwidth-intensive audio, video and voice.

Data warehousing has become popular as a way of making enterprise data available to decision makers for reporting and analysis without sacrificing the performance, security or integrity of production systems. These warehouses may comprise gigabytes or terabytes of data distributed over hundreds of platforms and accessed by thousands of users, and must be updated regularly to provide users near-real-time data for critical business reports and analyses.

Network backup of servers and storage systems is common in many industries which require enterprise information to be archived. Such backups usually occur during off hours and require large amounts of bandwidth during a fixed amount of time (4 to 8 hours). The backup involves gigabytes or terabytes of data distributed over hundreds of servers and storage systems throughout an enterprise.

As these applications proliferate and demand ever greater shares of bandwidth at the desktop, and as the total number of network users continues to grow, organizations will need to migrate critical portions of their networks to higher-bandwidth technologies.

Gigabit Ethernet: The Strategic Alternative for LANs and Intranets

The accelerating growth of LAN traffic is pushing network administrators to look to higher-speed network technologies to solve the bandwidth crunch. These administrators—who typically have either Ethernet or FDDI backbones today—have several alternatives to choose from. Although each network faces different issues, Gigabit Ethernet meets several key criteria for choosing a high-speed network:

- Easy, straightforward migration to higher performance levels without disruption
- Low cost of ownership—including both purchase cost and support cost
- Capability to support new applications and data types
- Network design flexibility

Easy Migration to Higher Performance

One of the most important questions network administrators face is how to get higher bandwidth without disrupting the existing network. Gigabit Ethernet follows the same form, fit and function as its 10 Mbps and 100 Mbps Ethernet precursors, allowing a straightforward, incremental migration to higher-speed networking. All three Ethernet speeds use the same IEEE 802.3 frame format, full-duplex operation and flow control methods. In half-duplex mode, Gigabit Ethernet employs the same fundamental CSMA/CD access method to resolve contention for the shared media. And, Gigabit Ethernet uses the same management objects defined by the IEEE 802.3 group. Gigabit Ethernet is Ethernet, only faster.

Ethernet Frame Format It is simple to connect existing lower-speed Ethernet devices to Gigabit Ethernet devices using LAN switches or routers to adapt one physical line speed to the other. Gigabit Ethernet uses the same variable-length (64- to 1514-byte packets) IEEE 802.3 frame format found in Ethernet and Fast Ethernet (Figure 1). Because the frame format and size are the same for all Ethernet technologies, no other network changes are necessary. This evolutionary upgrade path allows Gigabit Ethernet to be seamlessly integrated into existing Ethernet and Fast Ethernet networks.

Bytes	8	6	6	2	0 - 1500
	Preamble	Destination address	Source address	Length of data field	Protocol header, data and pad

Figure 1. IEEE 802.3 Frame

In contrast, other high speed technologies use fundamentally different frame formats. High-speed ATM, for example, implements a fixed-length data cell. When connecting Ethernet and Fast Ethernet to ATM, the switch or router must translate each ATM cell to an Ethernet frame, and vice versa.

Full-and Half-Duplex Operation As defined by the IEEE 802.3x specification, two nodes connected via a full-duplex, switched path can simultaneously send and receive packets. Gigabit Ethernet follows this standard to communicate in full-duplex mode. Almost all Gigabit Ethernet products on the market in 1998 are full duplex.

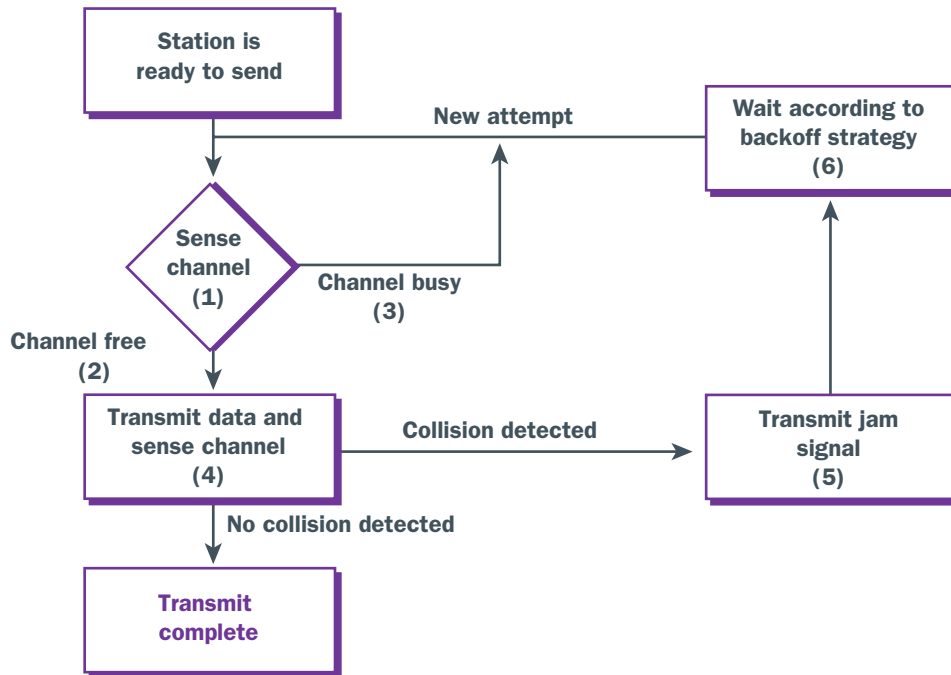


Figure 2. CSMA/CD Flow Chart

Gigabit Ethernet also employs standard Ethernet flow control methods to avoid congestion and overloading.

When operating in half-duplex mode, Gigabit Ethernet adopts the same fundamental CSMA/CD access method to resolve contention for the shared media. The CSMA/CD method is illustrated in Figure 2. The Gigabit Ethernet CSMA/CD method was enhanced in order to maintain a 200-meter collision diameter at gigabit speeds. Without this enhancement, minimum-sized Ethernet packets could complete transmission before the transmitting station senses a collision, thereby violating the CSMA/CD method.

To resolve this issue, both the minimum CSMA/CD carrier time and the Ethernet slot time have been extended from their present value of 64 bytes to a new value of 512 bytes. (Note that the minimum packet length of 64 bytes has not been affected.) Packets smaller than 512 bytes have been augmented with a new carrier extension field following the CRC field. Packets longer than 512 bytes have not been extended. These changes, which can impact small-packet performance, have been offset by incorporating a new feature, called packet bursting, into the CSMA/CD algorithm. Packet bursting will allow servers, switches and other devices to send bursts of small packets in order to fully utilize available bandwidth.

Devices that operate in full-duplex mode (switches and buffered distributors) are not subject to the carrier extension, slot time extension or packet bursting changes. Full-duplex devices will continue to use the regular Ethernet 96-bit interframe gap (IFG) and 64-byte minimum packet size.

Management Objects As in the transition from Ethernet to Fast Ethernet, the fundamental management objects familiar to most network managers are carried forward with Gigabit Ethernet. For example, SNMP defines a standard method to collect device-level Ethernet information. SNMP uses management information base (MIB) structures to record key statistics such as collision count, packets transmitted or received, error rates and other device-level information. Additional information is collected by remote monitoring (RMON) agents to aggregate the statistics for presentation via a network management application. Because Gigabit Ethernet uses standard Ethernet frames, the same MIBs and RMON agents can be utilized to provide network management at gigabit speeds.

Low Cost of Ownership

Cost of ownership is an important factor in evaluating any new networking technology. The overall cost of ownership includes not only the purchase price of equipment, but also the cost of training, maintenance and troubleshooting.

Competition and economies of scale have driven the purchase price of Ethernet connections down significantly. Though Fast Ethernet products have been shipping only since 1994, even these products have experienced significant price declines over the past two years. Gigabit Ethernet will follow the same price trends as Fast Ethernet. Products on the market today provide cost-effective connections for gigabit transmission rates. The IEEE's goal was to provide Gigabit Ethernet connections at two to three times the cost of a 100BASE-FX interface. As volume builds, reduced line width IC processes are implemented and low-cost opto-electronic devices are developed, the cost of Gigabit Ethernet interfaces will decline.

Switched Gigabit Ethernet connections are lower in cost than 622 Mbps ATM interfaces (assuming identical physical media interfaces), because of the relative simplicity of Ethernet and higher shipment volumes. Gigabit Ethernet repeater interfaces will be significantly lower in cost than 622 Mbps ATM connections, providing users with cost-effective alternatives for data center network backbone and server connections. Table 3 illustrates current prices for Ethernet, Fast Ethernet, FDDI and 622 Mbps ATM multimode and the target range for Gigabit Ethernet based on the IEEE 802.3z goals and Dell'Oro Group estimates.

Technology	Equipment Type	1996 Equipment Price/Port	1999 Equipment Price/Port	Change %
Shared Fast Ethernet	Hub	\$137	\$85	-39%
Switched Fast Ethernet	Switch	\$785	\$390	-50%
Shared FDDI	Concentrator	\$835	\$650	-22%
Switched FDDI	Switch	\$4000	\$1860	-54%
ATM 622 Mbps Estimate (multimode fiber)	Switch	\$6600	\$4800	-27%
Shared Gigabit Ethernet IEEE goal (multimode fiber)	Hub	N.A.	\$470 to \$700 ** (2x to 3x Fast Ethernet MM)	
Switched Gigabit Ethernet IEEE goal (multimode fiber)	Switch	N.A.	\$1070 to \$1610 ** (2x to 3x Fast Ethernet MM)	

Source: Dell'Oro Group

** Estimates based on Dell'Oro Group info and IEEE goals

Table 3. Network Backbone Connection Prices

Over time, advances in silicon, including 0.35-micron CMOS ASIC technology, will provide even greater performance gains and cost reduction opportunities that will result in a new, even more cost-effective generation of Ethernet technology. Analysis indicates that 0.35-micron processes will achieve 1250 Mbps operation and economically fit one million gates on a single die. This is more than enough to fit a complete Ethernet switch, including management, a significant amount of buffer memory, and an embedded 32-bit controller, on a single die—with obvious cost advantages.

Finally, because the installed base of users is already familiar with Ethernet technology, maintenance and troubleshooting tools, the support costs associated with Gigabit Ethernet will be far lower than other technologies. Gigabit Ethernet requires only incremental training of personnel and incremental purchase of maintenance and troubleshooting tools. In addition, deployment of Gigabit Ethernet will be faster than alternative technologies. Once upgraded with training and tools, network support staff will be able to confidently install, troubleshoot and support Gigabit Ethernet installations

Support for New Applications and Data Types

The emergence of intranet applications portends a migration to new data types, including video and voice. In the past it was thought that video might require a different networking technology designed specifically for multimedia. But today it is possible to mix data and video over Ethernet through a combination of the following:

- Increased bandwidth provided by Fast Ethernet and Gigabit Ethernet, enhanced by LAN switching
- The emergence of new protocols, such as Resource Reservation Protocol (RSVP), that provide bandwidth reservation
- The emergence of new standards such as 802.1Q and 802.1p which will provide virtual LAN (VLAN) and explicit priority information for packets in the network
- The widespread use of advanced video compression such as MPEG-2

These technologies and protocols combine to make Gigabit Ethernet an extremely attractive solution for the delivery of video and multimedia traffic, as illustrated in Table 4.

Capabilities	Gigabit Ethernet	Fast Ethernet	ATM	FDDI
IP Compatibility	Yes	Yes	Requires RFC 1557 or IP over LANE today; I-PNNI and/or MPOA in the future	Yes
Ethernet Packets	Yes	Yes	Requires LANE	Yes, though 802.1h translation bridging
Handle Multimedia	Yes	Yes	Yes, but application needs substantial changes	Yes
Quality of Service	Yes, with RSVP and/or 802.1p	Yes with RSVP and/or 802.1p	Yes with SVCs or RSVP with complex mapping from IETF (work in progress)	Yes, with RSVP and/or 802.1p
VLANs with 802.1Q/p	Yes	Yes	Requires mapping LANE and/or SVCs to 802.1Q	Yes

Table 4. High-Speed Network Capabilities

Flexible Internetworking and Network Design

Network administrators today face a myriad of internetworking choices and network design options. They are combining routed and switched networks, and building intranets of increasing scale. Ethernet networks are shared (using repeaters) and switched based on bandwidth and cost requirements. The choice of a high-speed network, however, should not restrict the choice of internetworking or network topology.

Gigabit Ethernet can be switched, routed and shared. All of today's internetworking technologies, as well as such technologies such as IP-specific switching and Layer 3 switching, are fully compatible with Gigabit Ethernet, just as they are with Ethernet and Fast Ethernet. Gigabit Ethernet is available in a full duplex repeater (with the accompanying low cost per port) as well as on LAN switches and routers.

Gigabit Ethernet Technology

The simple migration and support offered by Ethernet, combined with the scalability and flexibility to handle new applications and data types, makes Gigabit Ethernet the strategic choice for high-speed, high-bandwidth networking.

Gigabit Ethernet is an extension to the highly successful 10 Mbps and 100 Mbps IEEE 802.3 Ethernet standards. Offering a raw data bandwidth of 1000 Mbps, Gigabit Ethernet maintains full compatibility with the huge installed base of Ethernet nodes.

The Gigabit Ethernet Standard—IEEE 802.3z

To recap the recent history of the Gigabit Ethernet standards process, in July, 1996, after months of initial feasibility studies, the IEEE 802.3 working group created the 802.3z Gigabit Ethernet task force. The key objectives of the 802.3z Gigabit Ethernet task force were to develop a Gigabit Ethernet standard that does the following:


- Allows half- and full-duplex operation at speeds of 1000 Mbps
- Uses the 802.3 Ethernet frame format
- Uses the CSMA/CD access method with support for one repeater per collision domain
- Addresses backward compatibility with 10BASE-T and 100BASE-T technologies

The task force identified three specific objectives for link distances: a multimode fiber-optic link with a maximum length of 550 meters; a single-mode fiber-optic link with a maximum length of 3 kilometers (later extended to 5 kilometers); and a copper based link with a maximum length of at least 25 meters. The IEEE is also actively investigating technology that would support link distances of at least 100 meters over Category 5 unshielded twisted pair (UTP) wiring. This standards work will be done by 1999. In addition, the task force decided to include a specification for an optional Gigabit Media Independent Interface (GMII) in the scope of its work.

One of the primary goals of the Gigabit Ethernet Alliance was to accelerate the standards activity for Gigabit Ethernet. Fast Ethernet took approximately 13 months to go from first draft to final approval, and Gigabit Ethernet required was approximately the same amount of time (Figure 3). The goal of the IEEE 802.3z Gigabit Ethernet task force was to complete the Gigabit Ethernet standard by 1998, although pre-standard products appeared in 1997.

Standards Timeline

Fast Ethernet



Gigabit Ethernet 802.3z Milestones




Figure 3. Fast Ethernet and Gigabit Ethernet Standards Timelines

Distance Specifications for Fiber Optic Media

Table 5 shows the Gigabit Ethernet distance specifications for fiber optic media. For the 62.5 micron diameter 160 MHz*km multimode (MM) fiber often called FDDI-grade fiber, the distance is specified at 220 meters. As is evident from the table, as the bandwidth of the fiber increases, the minimum range for MM fiber increases up to 550 meters. The longwave length transceiver (1000BASE-LX) reaches 550 meters for all media types. For single mode fiber with 1000BASE-LX, the distance is specified at 5 km.

Standard	Fiber Type	Diameter (microns)	Modal Bandwidth (MHz* km)	Minimum Range Meters
1000BASE-SX	MM	62.5	160	2 to 220*
	MM	62.5	200	2 to 275**
	MM	50	400	2 to 500
	MM	50	500	2 to 550***
1000BASE-LX	MM	62.5	500	2 to 550
	MM	50	400	2 to 550
	MM	50	500	2 to 550
	SM	9	N/A	2 to 5000

Table 5. Gigabit Ethernet Distances for Fiber Optic Media

Notes:

- * The TIA 568 building wiring standard specifies 160/500 MHz*km multimode fiber
- ** The international ISO/IEC 11801 building wiring standard specifies 200/500 MHz*km multimode fiber
- *** The ANSI Fibre Channel specification specifies 500/500 MHz*km 50 micron multimode fiber and 500/500 fiber has been proposed for inclusion in ISO/IEC 11801.

Lasers and Multimode Fiber

The challenges associated with using lasers on multimode fiber have become more apparent as the operating speed has been increased. In addition, the IEEE 802.3z task force was particularly attentive to the characteristics of the installed base of network cabling, be it copper or fiber optic cabling. The group conducted numerous experimental and field tests to insure that the underlying signaling technology would work on the vast majority of installed cabling. Thus, the IEEE applied far greater scrutiny to the operating characteristics of lasers on multimode fiber than has ever been applied before. This is the first time that laser launch has been extensively tested over long multimode fiber. These rigorous tests discovered a jitter component caused by a phenomenon known as differential mode delay (DMD). (For more details, see the Gigabit Ethernet Alliance web site: www.gigabit-ethernet.org)

This jitter effect regarding multimode fiber transmissions was resolved. For 1000BASE-SX, the solution was achieved by qualifying the launch of the laser transmitter, introducing conformance tests for stressed receiver sensitivity and stressed receiver jitter, and reallocating the jitter budget. In addition to these refined transceiver conformance tests, the link distance for the lowest bandwidth multimode fiber was specified at 220 meters. Other fiber types can go further, see Table 5. With 1000BASE-LX transceivers over multimode fiber, external patch cords are used to mitigate DMD. Existing technologies that use the combination of short wavelength lasers and multimode fiber, such as Fibre Channel (FC), have not seen DMD effects because of the short distances used in FC applications.

Full- and Half-Duplex over Fiber Today, UTP in the Future

Gigabit Ethernet supports new full-duplex operating modes for switch-to-switch and switch-to-end-station connections, and half-duplex operating modes for shared connections using repeaters and the CSMA/CD access method. Initially operating over optical fiber, Gigabit Ethernet will also be able to use Category 5 UTP cabling. Figure 4 illustrates the functional elements of Gigabit Ethernet.

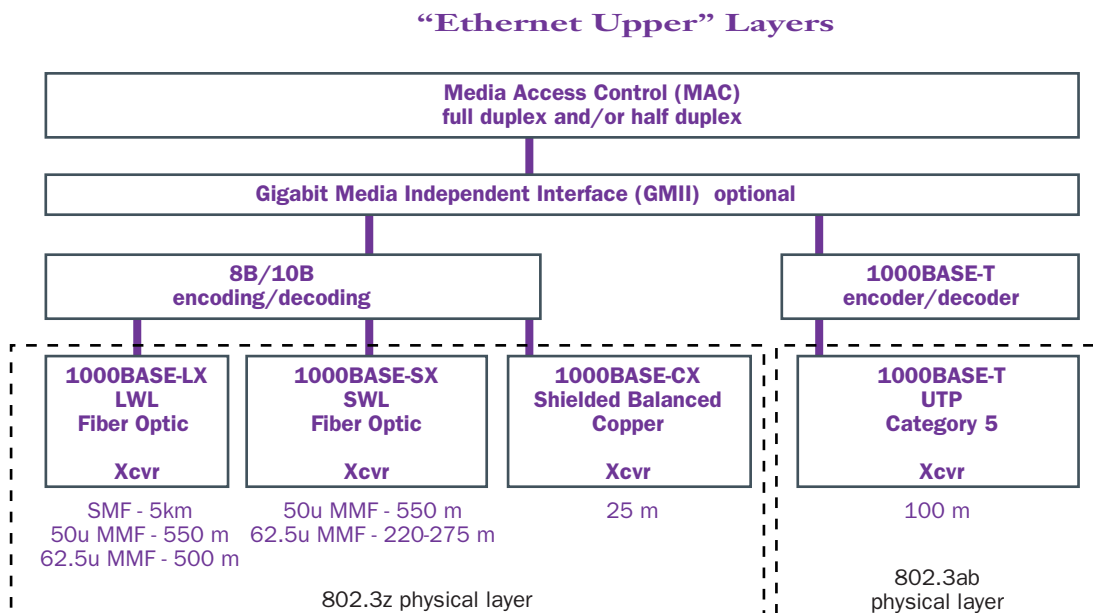


Figure 4. Functional Elements of Gigabit Ethernet Technology

Much of the effort of the IEEE 802.3z task force was devoted to the definition of transceiver or PHY standards for Gigabit Ethernet. Like other standards based on the International Standards Organization (ISO) model, Gigabit Ethernet implements functionality adhering to a physical layer standard. In general, the PHY or physical layer is responsible for defining the mechanical, electrical and procedural characteristics for establishing, maintaining and deactivating the physical link between network devices. For Gigabit Ethernet communications, several physical layer standards are emerging from the IEEE 802.3z effort.

Two PHYs provide Gigabit transmission over fiber-optic cabling. 1000BASE-SX is targeted at lowest cost multimode fiber runs in horizontal and shorter backbone applications. 1000BASE-LX is targeted at longer multimode building fiber backbones and single-mode campus backbones. For multimode fiber, these standards define gigabit transmission over maximum distances ranging from 220 to 550 meters. Single-mode fiber, which is covered by the long-wavelength standard, is defined to cover distances of 5 kilometers.

There are also two standards efforts for Gigabit Ethernet transmission over copper cabling. The first copper link standard is defined by the 802.3z task force and is referred to as 1000BASE-CX. This standard supports interconnection of equipment clusters where the physical interface is short-haul copper. It supports a switching closet or computer room as a short jumper interconnection for 25 meter distances. This standard uses the Fibre Channel-based 8B/10B coding at the serial line rate of 1.25 Gbps, and runs over 150-ohm balanced, shielded, specialty cabling assemblies. IBM Type I cabling is not recommended. This copper physical layer standard has the advantage that it can be generated quickly and is inexpensive to implement. The short copper link standard was completed in the same time frame as the fiber links.

The second copper link standard is intended for use in horizontal copper cabling applications. In March, 1997, a Project Authorization Request (PAR) was approved by the IEEE Standards Board, enabling the creation of a separate but related committee referred to as the 802.3ab Task Force. This new group is chartered with the development of a 1000BASE-T physical layer standard providing 1 Gbps Ethernet signal transmission over four pairs of category 5 UTP cable, covering cabling distances of up to 100 meters or networks with a diameter of 200 meters. This standard will outline communications used for horizontal copper runs on a floor within a building using structured generic cabling, taking advantage of the existing UTP cable already deployed. This effort requires new technology and new coding schemes in order to meet the potentially difficult and demanding parameters set by the previous Ethernet and Fast Ethernet standards. The 1000 Base-T standards process should be complete by the first half of 1999.

The IEEE 802.3z standards activity is planning ahead for the expected advances in silicon technology and digital signal processing that will eventually enable Gigabit Ethernet to operate over UTP cabling. To accommodate this, a logical interface will be specified between the MAC and PHY layers that will decouple the Fibre Channel 8B/10B encoding, allowing other encoding schemes that more readily support the use of cost-effective UTP cabling. As Figure 4 shows, 1000BASE-T does not use the Fibre Channel 8B/10B encoding scheme.

Fibre Channel Components

During the development of the standard, the IEEE 802.3 committee drew heavily on the physical layer technology and expertise developed by ANSI as part of the X3.230 Fibre Channel project. As a result, many manufacturers have been able to re-use physical-layer components originally designed for Fibre Channel in their new Gigabit Ethernet designs. This reuse of existing, proven technologies and methods has minimized the development time for many Gigabit Ethernet products. It has also made it possible for the vendors of optical components and high-speed 8B/10B serializers to manufacture one set of parts that may be used with either standard, thus aggregating the benefits of large volume production across both marketplaces.

Gigabit Ethernet Products

Since Gigabit Ethernet is Ethernet, the types of Gigabit Ethernet products are quite straightforward: Layer 2 switches, Layer 3 switches (or routing switches), uplink/downlink modules, NICs, Gigabit Ethernet router interfaces, and the buffered distributors. There are pure multiport Gigabit Ethernet switches with high performance backplanes, as well as devices that have both Gigabit Ethernet and Fast Ethernet ports in the same box. Gigabit Ethernet uplinks have appeared as modular upgrades for fixed-configuration Fast Ethernet devices or modular, chassis-based hubs to provide a high-speed connection to the network. Vendors of high-performance routers can be expected to deliver Gigabit Ethernet interfaces as well.

Some Gigabit Ethernet vendors have developed a new device called a full duplex repeater or buffered distributor. The full duplex repeater is a full-duplex, multiport, hub-like device that interconnects two or more 802.3 links operating at 1 Gbps or faster. Like an 802.3 repeater, it is a non-address-filtering device. The buffered distributor forwards all incoming packets to all connected links except the originating link, providing a shared bandwidth domain comparable to a 802.3 collision domain. Unlike an 802.3 repeater, the buffered distributor is permitted to buffer one or more incoming frames on each link before forwarding them.

As a shared bandwidth device, the buffered distributor should be distinguished from both routers and switches. While routers with Gigabit Ethernet interfaces may have backplanes that support bandwidths greater or less than gigabit rates, the ports attached to a Gigabit Ethernet buffered distributor's backplane share one gigabit of bandwidth. In contrast, the backplanes of high-performance, multiport Gigabit Ethernet switches will support multigigabit bandwidths.

Ethernet and Higher-Level Services

Gigabit Ethernet provides high-speed connectivity, but does not by itself provide a full set of services such as Quality of Service (QoS), automatic redundant fail-over, or higher-level routing services; these are added via other open standards. Gigabit Ethernet, like all Ethernet specifications, specifies the data link (layer 2) of the OSI protocol model, while TCP and IP in turn specify the transport (layer 4) and network (layer 3) portions and allow reliable communication services between applications. Issues such as QoS are not addressed in the original Gigabit Ethernet specifications, but must be addressed across several of these standards. RSVP, for instance, is defined at the network layer to work alongside IP. Layer 3 (routing) services also operate at the network layer (Table 6).

OSI Layer	OSI Name	Examples
4	Transport	TCP
3	Network	IP,RSVP
2	Data link	Ethernet (MAC), 802.1p, 802.1Q
1	Physical	10Base-T, 100Base-T, Fibre Channel

Table 6. Layers of Network Functionality

Various implementations of Gigabit Ethernet may include one or more of these standards in order to provide a more robust or functional networking connection, but the overall success of Gigabit Ethernet is not tied to any one of them. The advantage of modular standards is that any one piece may evolve and be adopted at a pace determined by market need and product quality. Note that all of the standards are just as readily paired with Fast Ethernet and 10 Mbps Ethernet, so that all levels of Ethernet performance can benefit from all the standards work.

Quality of Service on Ethernet Applications emerging in the mid to late 90's demand consistent bandwidth, latency, and jitter from network connections. Such applications include voice and video over LANs and WANs, multicast software distribution, and the like. Standards bodies have responded with new open definitions such as RSVP and the current work in the IEEE 802.1p and IEEE 802.1Q standards groups. RSVP is gaining industry acceptance as a preferred way to request and provide quality of network connections. In order to have RSVP function and deliver defined and consistent quality to an application, each network component in the chain between client and server must support RSVP and communicate appropriately. Because of the need to have so many components supported by RSVP before meaningful results can be achieved, some vendors are advancing proprietary schemes to deliver some degree of QoS. Some of these may deliver QoS benefits to users, but will require certain portions of the network to be vendor-specific implementations.

802.1p and 802.1Q facilitate quality of service today over Ethernet by providing a means for “tagging” packets with an indication of the priority or class of service desired for the packet. These tags allow applications to communicate the priority of packets to internetworking devices. RSVP support can be achieved by mapping RSVP sessions into 802.1p service classes.

Layer 3 Functionality Layer 3 involves determination of the eventual destination of a packet—beyond its MAC destination address on the packet header. By examining the IP address (buried deeper in the packet), the IP subnet can be determined, allowing broadcasts to be contained to the appropriate subnets and packets to be forwarded accurately to intermediate nodes for most efficient transit through the network.


The classic Layer 3 device is the router, which makes Layer 3 decisions by implementing complex algorithms and data structures in software. While such complicated routing tasks formerly required complex and software-intensive multiprotocol router products, vendors over the last year have shipped Layer 3 switches or routing switches that accomplish many of these tasks, while delivering arguably better price/performance than traditional routers. Narrowing the protocol supported to IP has allowed these devices to optimize tasks and accomplish more work with dedicated hardware.

Gigabit Ethernet Performance

Figure 5 shows third party tests results for the single port performance of a Gigabit Ethernet multi-port switch at various packet size loads: 64 bytes, 128 bytes, . . . , 1518 bytes. The switch was non-blocking and the total throughput of the switch is divided by 8 ports to obtain the throughput number for each port respectively. Depending on the packet size, throughput ranged from 81,274 pps up to 1.488 million pps.

The test shows that practically 100% efficiency is obtained. For example, with 64 byte packets, the throughput is over 1.488 million packets per second. Once the 12 byte interframe gap and the 8 byte preamble are included, the test shows that the Gigabit Ethernet link performs at full wire-speed of 1.0 billion bps! (1.488095 million pps * [64 + 12 + 8 bytes] * 8 bits/byte = 1,000.000 million bits per second.)

Similarly, for the largest 1518 byte packet with a measured throughput of 81,274 pps, efficiency exceeds 99%. (81,274 pps * [1518 + 12 + 8] * 8 bits/byte = 999.995 million bps.) See Table 7 for complete details.



Source: Extreme Networks.

NDTL results for single port Gigabit Ethernet

Figure 5. Gigabit Ethernet Performance

Packet Size	Theoretical Throughput (pps)	Actual Throughput (pps)	bps (000s)	Efficiency
64	1,488,095	1,488,095	1,000,000	100.00%
128	844,594	844,594	999,999	100.00%
256	452,898	452,898	999,999	100.00%
512	234,962	234,962	999,998	100.00%
768	158,629	158,629	999,997	100.00%
1024	119,731	119,731	999,993	100.00%
1280	96,153	96,153	999,991	100.00%
1518	81,274	81,274	999,995	100.00%

Source: Extreme Networks, NDTL Results
 Note: bps calculation includes preamble and interframe gap.

Table 7. Gigabit Ethernet Performance

Gigabit Ethernet Migration

The initial applications for Gigabit Ethernet are for campuses or buildings requiring greater bandwidth between routers, switches, hubs, repeaters and servers. Examples include switch-to-router, switch-to-switch, switch-to-server and repeater-to-switch connections (Table 8). In its early phase, Gigabit Ethernet is not expected to be deployed widely to the desktop. In all scenarios the network operating system (NOS), applications and NIC drivers at the desktop will remain unchanged. The MIS manager can also leverage not only the existing multimode fiber, but also the current investment in network management applications and tools.

- 1 Upgrading switch-to-switch connections ➤ Obtain 1000 Mbps pipes between 100/1000 switches
- 2 Upgrading switch-to-server connections ➤ Achieve high-speed access to applications and file servers
- 3 Upgrading a switched Fast Ethernet backbone ➤ Aggregate Fast Ethernet switches with a Gigabit Ethernet switch or repeater
- 4 Upgrading a shared FDDI backbone ➤ Install FDDI switch or Ethernet-to-FDDI switches/routers with Gigabit Ethernet switches or repeaters
- 5 Upgrading high-performance workgroups ➤ Gigabit Ethernet NICs for connection to Gigabit Ethernet switches or repeaters

Table 8. The Five Most Likely Upgrade Scenarios

Upgrading Switch-to-Switch Links

A very straightforward upgrade scenario is upgrading 100 Mbps links between Fast Ethernet switches or repeaters to 1000 Mbps links between 100/1000 switches. Such high-bandwidth, switch-to-switch links would enable the 100/1000 switches to support a greater number of both switched and shared Fast Ethernet segments. Figure 6a illustrates the network prior to upgrade, while Figure 6b illustrates the network after upgrading to Gigabit Ethernet.

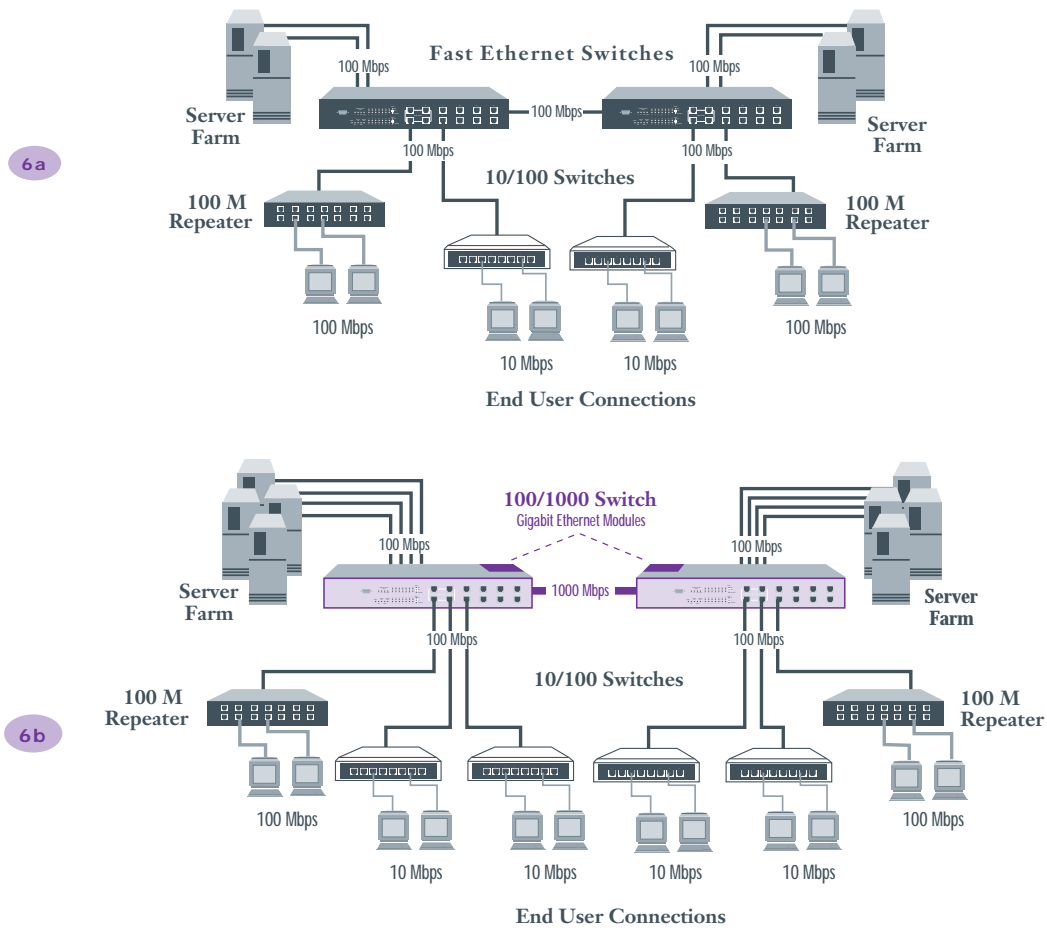
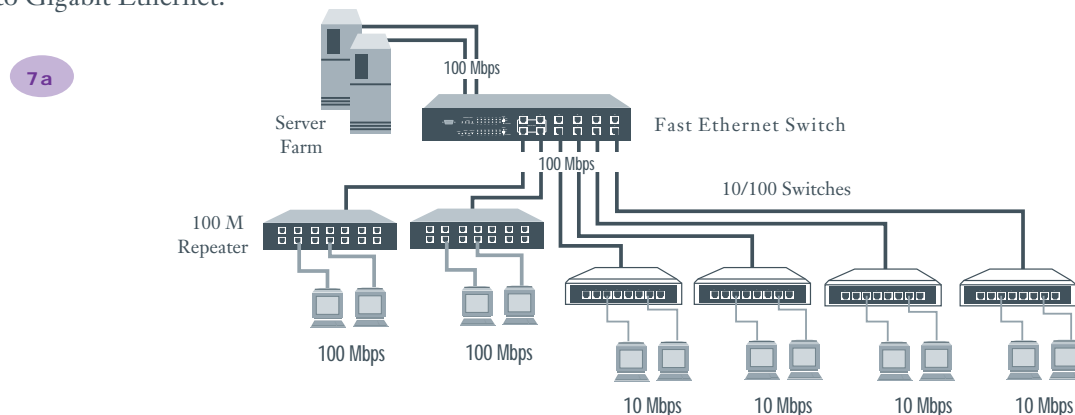


Figure 6. Upgrading Switch-to-Switch Links

Upgrading Switch-to-Server Links

The simplest upgrade scenario is upgrading a Fast Ethernet switch to a Gigabit Ethernet switch to obtain high-speed, 1000 Mbps interconnection to a server farm of high-performance super servers with Gigabit Ethernet NICs installed. Figure 7a illustrates the network prior to upgrade, while Figure 7b illustrates the network after upgrading to Gigabit Ethernet.



7b

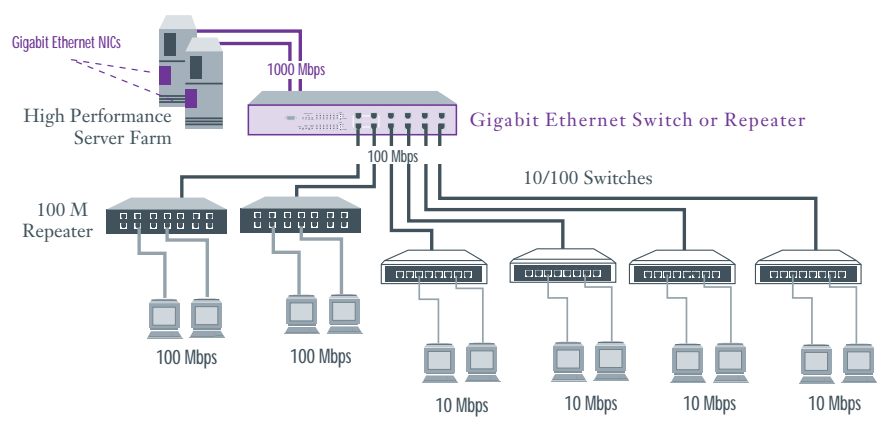
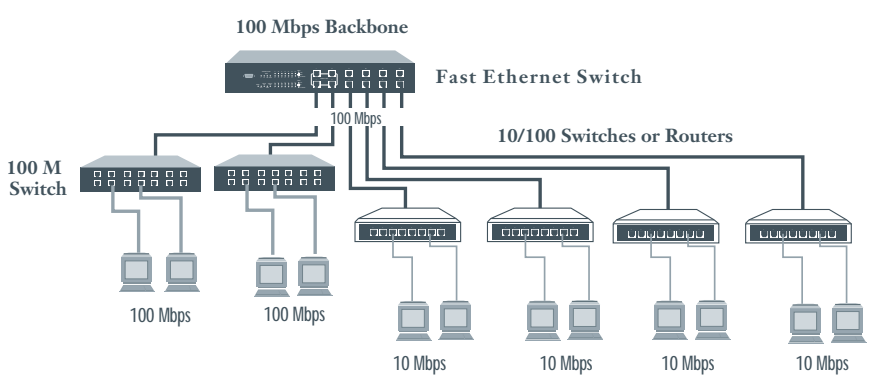


Figure 7. Upgrading Switch-to-Server Links

Upgrading a Switched Fast Ethernet Backbone

A Fast Ethernet backbone switch that aggregates multiple 10/100 switches can be upgraded to a Gigabit Ethernet switch supporting multiple 100/1000 switches as well as other devices such as routers and hubs with Gigabit Ethernet interfaces and uplinks. Gigabit repeaters and/or buffered distributors can also be installed as needed. Once the backbone is upgraded to a Gigabit Ethernet switch, high-performance server farms can be connected directly to the backbone with Gigabit Ethernet NICs, increasing throughput to the servers for users with high-bandwidth applications. Also, the network can now support a greater number of segments, more bandwidth per segment, and hence a greater number of nodes per segment. Figure 8a illustrates the network prior to upgrade, while Figure 8b illustrates the network after upgrading to Gigabit Ethernet.

8a



8b

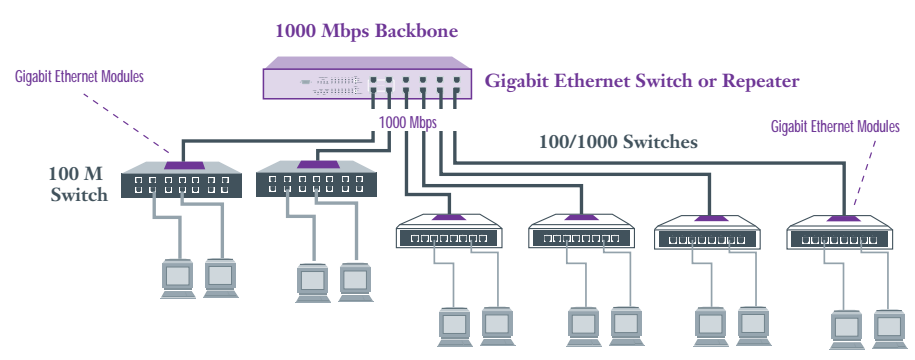


Figure 8. Upgrading a Switched Fast Ethernet Backbone

Upgrading a Shared FDDI Backbone

An FDDI campus or building backbone can be upgraded by replacing the FDDI concentrator or hub or Ethernet-to-FDDI router with a Gigabit Ethernet switch or repeater. (As an intermediate step, some users might migrate to an FDDI switch before installing a Gigabit Ethernet switch.) The only upgrade required is the installation of new Gigabit Ethernet interfaces in the routers, switches or repeaters. All the investment in fiber-optic cabling is retained, and the aggregate bandwidth is increased at least tenfold for each segment. Figure 9a illustrates the network prior to upgrade, while Figure 9b illustrates the network after upgrading to Gigabit Ethernet

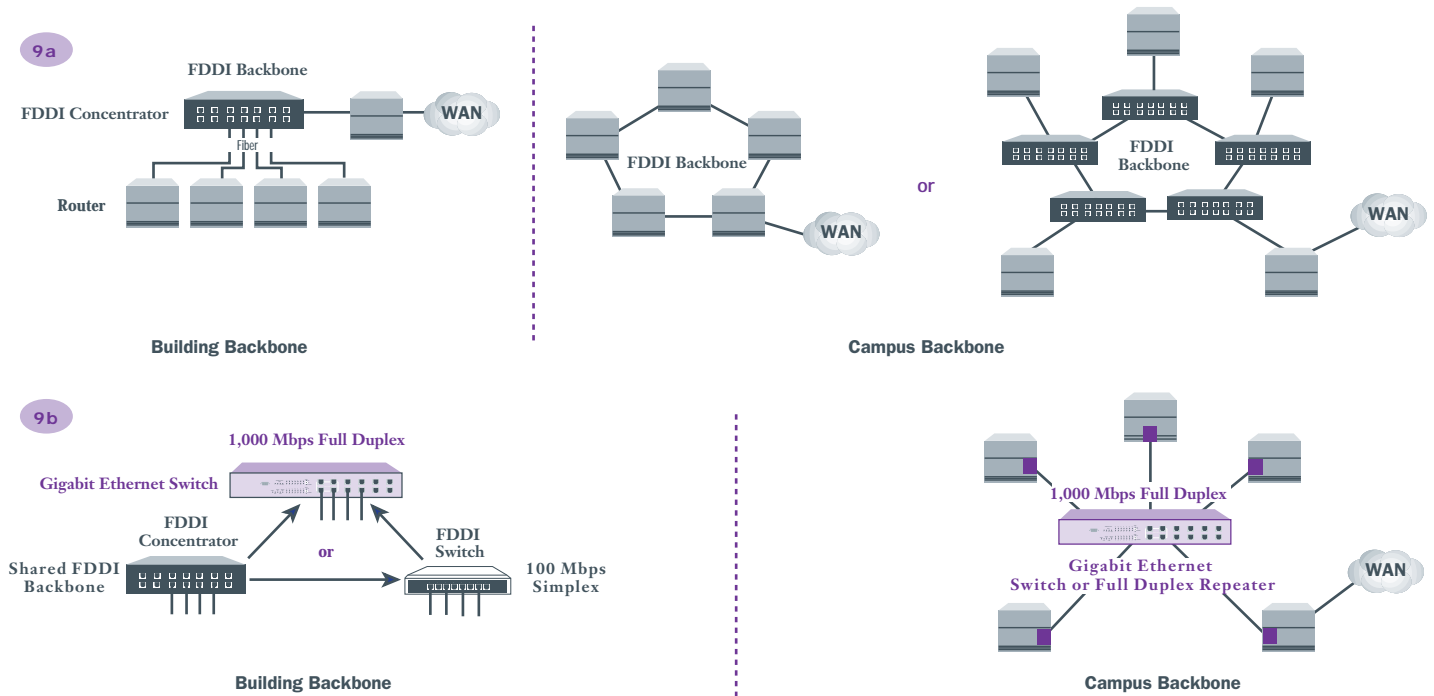


Figure 9. Upgrading a Shared FDDI Backbone

Upgrading High-Performance Workgroups

In the later phases of Gigabit Ethernet adoption, as Fast Ethernet or FDDI connected desktops run out of bandwidth, Gigabit Ethernet NICs will be used to upgrade high-performance desktop computers with Gigabit Ethernet connectivity. The high-performance desktop computers are then connected to Gigabit Ethernet switches or buffered distributors. Figure 10a illustrates the network prior to upgrade, while Figure 10b illustrates the network after upgrading to Gigabit Ethernet.



Figure 10. Upgrading High-Performance Workgroups

About the Gigabit Ethernet Alliance

The Gigabit Ethernet Alliance is an open forum whose purpose is to promote industry cooperation in the development of Gigabit Ethernet. The primary objectives of the Alliance are to:

- 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
- Foster two-way communications between potential suppliers and consumers of Gigabit Ethernet products
- Recruit new members to participate in the Alliance

The Gigabit Ethernet Alliance builds on its members' past experience and success with the Fast Ethernet Alliance. The Alliance has the following organizational structure:

- A steering committee, responsible for oversight of all Alliance activities
- A technical sub-group
- A marketing and communications sub-group

Membership in the Alliance and participation in Alliance activities is open to all interested parties. For more information on Gigabit Ethernet and the Gigabit Ethernet Alliance, visit our World Wide Web site at: <http://www.gigabit-ethernet.org>.

Abbreviations and Acronyms

ASIC	application-specific integrated circuit	MPOA	Multiprotocol over ATM
ATM	Asynchronous Transfer Mode	NIC	network interface card
CSMA/CD	Carrier Sense Multiple Access with Collision Detection	NOS	network operating system
FDDI	Fiber Distributed Data Interface	PAR	Project Authorization Request
GMI	Gigabit Media Independent Interface	PNNI	Private Network to Network Interface
IC	integrated circuits	QoS	Quality of Service
IEEE	Institute of Electrical and Electronic Engineers	RMON	remote monitoring
IETF	Internet Engineering Task Force	RSVP	Resource Reservation Protocol
IFG	interframe gap	SNMP	Simple Network Management Protocol
ISO	International Standards Organization	TCP/IP	Transmission Control Protocol/ Internet Protocol
LANE	LAN Emulation	UTP	unshielded twisted pair
MAC	media access control	VLAN	virtual LAN
MIB	management information base		

