

# Gigabit Ethernet Technical Brief

---

## Achieving End-to-End Performance



Alteon Networks, Inc.  
6351 San Ignacio Avenue  
San Jose, CA 95119

1-408-574-5500

First Edition  
September 1996

© 1996 by Alteon Networks, Inc. All rights reserved.

# Contents

---

<b>Gigabit Ethernet Technical Brief</b>	
<b>Achieving End-to-End Performance</b>	
Goals of Gigabit Ethernet	1
Uses of Gigabit Ethernet	2
Gigabit Ethernet History	3
Gigabit Ethernet Momentum	4
Workstation Speed	4
Desktop Multimedia	5
Multivendor Interoperability	6
Migrating to Gigabit Ethernet	6
Gigabit Ethernet Components	6
Upgrade Scenarios	6
Upgrading Connections to Centralized File and Compute Servers	7
Upgrading Connections between Switches	8
Upgrading a Switched Fast Ethernet Backbone	8
Upgrading a Shared FDDI Backbone	9
Protocol Architecture	10
Physical Interface Characteristics	11
Serializer/Deserializer	11
8B/10B Encoding	11
MAC Layer	11
Upper Layers	11
Cabling types and distances	12
Flow Control	12
Full-Duplex Transmission	12
Half-duplex Transmission	13
Carrier Extension	14
Technology Advances	15
Ethernet Adapters	15
First Generation Ethernet Adapters	16
Second Generation Ethernet Adapters	17
Third Generation Ethernet Adapters—The Alteon Difference	18
Ethernet Switches	20
Specialized ASICs to meet the Performance Demands of Gigabit Ethernet	21
Seamless Integration with Ethernet and Fast Ethernet	21
Standard Ethernet Management	21
Conclusion	22



# Gigabit Ethernet Technical Brief

## Achieving End-to-End Performance

Intranet Web servers, centralized file and compute servers, data warehousing, groupware, medical imaging, CAD/CAM applications, 3-D modeling, animation, video, pre-press applications, server farms, seismic processing...

The list goes on and on. The demand for high-speed network connections is proliferating at a pace almost as rapid as the speed requirements of the applications themselves. Evidence is everywhere: the rapid acceptance of 10/100 Mbps connections on today's desktop computers, Ethernet switching at the department level, and the deployment of Fast Ethernet switches in corporate backbones are a few examples of the need for faster and faster networks.

And still, bottlenecks remain. Server network connections have been limited to 100 Mbps since FDDI was shipping in volume in the late 1980's. Fast Ethernet made it easier to build internetworking products, but did not provide a faster server interface. Today, centralized servers are often configured with multiple 100 Mbps network connections to meet bandwidth requirements.

Enter Gigabit Ethernet.

Gigabit Ethernet is a new technology that will provide seamless interoperability with Ethernet and Fast Ethernet. Gigabit Ethernet transfers data at a blazingly fast speed: one gigabit per second, or 100 times the rate of standard Ethernet. Gigabit Ethernet is designed to deliver the same benefits as Fast Ethernet: seamless integration with installed Ethernets, dramatically higher performance than the previous standard, and a familiar management environment.

It couldn't happen at a better time. Multimedia over IP is just starting to take off with the proposed IETF standards, real-time transfer protocol (RTP) and resource reservation protocol (RSVP). These protocols, combined with gigabit networks, high-performance desktops, and Internet technology, will quickly change the way corporations access information.

Ethernet. Fast Ethernet. Gigabit Ethernet. Networking made simple!

### Goals of Gigabit Ethernet

Under development in the IEEE by the 802.3z Task Force, Gigabit Ethernet has the following primary goals:

- Complete interoperability with Ethernet and Fast Ethernet
  - Retain the installed base of NICs
  - Leverage the investment in hubs, switches, and routers
  - Leverage the network management environment

## *Gigabit Ethernet Technical Brief*

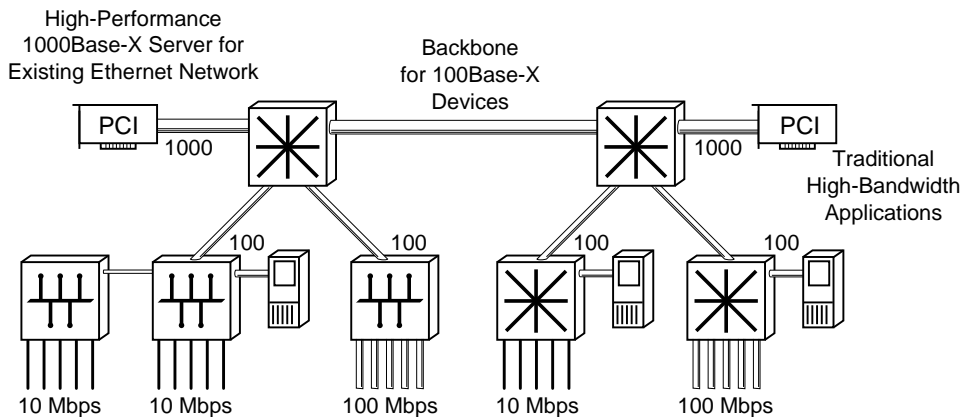
- Conform to the Ethernet standard
  - Frame format
  - Minimum and maximum frame sizes
  - CSMA/CD access method
  - 802.2 LLC specifications
- Provide a simple forwarding mechanism between 10, 100, and 1000 Mbps
  - No fragmentation of frames
  - No encapsulation of frames
  - No translation of frames
- Provide 10 times the performance of Fast Ethernet

### **Uses of Gigabit Ethernet**

Initially, Gigabit Ethernet will be deployed in three key areas:

- Aggregating traffic between Ethernet clients and centralized file or compute servers
- Connecting multiple 100Base-T Fast Ethernet switches through 100/1000 Mbps switches
- Connecting both workstations and servers with Gigabit Ethernet to run high-bandwidth applications, such as CAD/CAM, medical imaging, and pre-press

Figure 1 illustrates the projected initial uses of Gigabit Ethernet.



**Figure 1. Initial deployment of Gigabit Ethernet**

In the future, the uses of Gigabit Ethernet will be extended to incorporate switches based entirely on Gigabit Ethernet on the campus backbone, centralized server farms connected to these all-gigabit switches, and specialized high-performance client workstations. Figure 2 illustrates some future Gigabit Ethernet possibilities.




Figure 2. Future deployment of Gigabit Ethernet

## Gigabit Ethernet History

In November of 1995, the IEEE 802.3 committee formed a High-Speed Study Group to examine the future direction of high-speed networking. The IEEE approved the 802.3z Task Force, focused exclusively on Gigabit Ethernet, in March of 1996. The IEEE 802.3z Task Force will develop the standard for Gigabit Ethernet.

The Gigabit Ethernet Alliance (GEA) was formed in May of 1996. Modeling itself after the Fast Ethernet Alliance, the GEA is an open forum whose goal is to support the development, standardization, and deployment of Gigabit Ethernet.

The primary goals of the Gigabit Ethernet Alliance are:

- Fully support the Gigabit Ethernet standards activities of the IEEE 802.3z working group
- Ensure product interoperability
- Provide a mechanism for Gigabit Ethernet suppliers and consumers to communicate
- Increase its membership

More than 60 companies, including some of the biggest networking companies in the industry, participate in regular meetings of the GEA. The GEA home page is located on the World Wide Web at <http://www.gigabit-ethernet.org>.

## Gigabit Ethernet Momentum

There are a number of factors driving the development of Gigabit Ethernet. These include the increased speed of workstations using the PCI bus, the need for multimedia on the desktop, and corporate demands for multivendor interoperability.

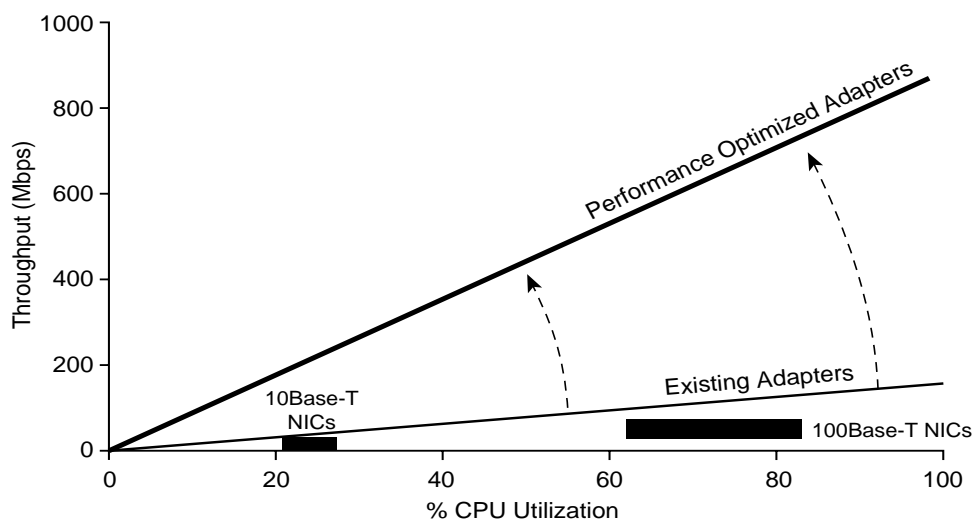
### Workstation Speed

With the advent of the PCI bus and high-performance CPUs, workstations are getting faster and faster. In fact, PCI-bus equipped workstations have more than enough compute power to handle data transfers at gigabit speeds. Table 1 shows the theoretical upper limit of throughput for four bus types that have been used in standard workstations over the years.

**Table 1. Workstation Throughput**

Bus Type	Throughput
ISA	64 Mbps
EISA	264 Mbps
MCA	320 Mbps
PCI (32 bit)	1,056 Mbps
PCI (64 bit)	2,112 Mbps

Faster bus speed is one important factor in achieving gigabit-speed readiness for standard workstations. In addition to having a faster bus, the interaction between the host workstation and the NIC must be optimized. The graph shown in Figure 3 illustrates the impact of greater throughput on workstation host CPU utilization.



**Figure 3. Host throughput versus CPU utilization**



Using current NIC technology, as throughput approaches 100 Mbps, CPU utilization quickly approaches 100%. In order to support data transfers at gigabit speeds, the design of NICs cannot simply be scaled based on existing Ethernet or Fast Ethernet cards. Instead, a new, optimized method of transferring data between the network and the host must be implemented. Later in this paper, we will discuss how Alteon Networks is creating performance optimized adapters to answer the demands of Gigabit Ethernet.

## **Desktop Multimedia**

Up until now, desktop multimedia applications have not experienced the explosion of use that has been predicted for the past few years. The reason for the sluggish implementation of desktop multimedia is simple: there has not been a combination of technologies available within most network managers' budgets to support it.

Multimedia applications need a mixture of ultra fast networking technology and an accompanying suite of protocols designed to simultaneously handle bursty data traffic and streaming voice/video traffic.

The combination of Gigabit Ethernet technology and a new set of standard IP-based protocols is paving the way for desktop multimedia to become an affordable reality.

The Internet Engineering Task Force (IETF) has standardized three new protocols that work together to provide desktop multimedia:

- Real-time transfer protocol (RTP)
- Real-time transfer control protocol (RTCP)
- Resource reservation protocol (RSVP)

RTP is a streaming-oriented protocol that works with the TCP/IP protocol suite to provide end-to-end delivery of multimedia traffic. RTP adds a 10-byte header to each packet. The header contains a time-stamp and sequencing information that allows time-sensitive data to be synchronized and reassembled in the correct order.

RTCP is a control protocol that is used in conjunction with RTP. RTCP is responsible for providing feedback on the quality of network conditions, and allows applications to adapt to those conditions.

RSVP is a protocol used for reserving network resources, such as buffers or bandwidth, for a session. It works by allowing applications to submit a bandwidth reservation request to the host PC. The host PC then requests a specific quality of service from RSVP-equipped routers along the data path, providing end-to-end allocation of network resources.

These three protocols are leading the way for true desktop multimedia to become a reality in the very near future. Gigabit Ethernet is poised to handle the very large bandwidth requirements that desktop multimedia will require.

### **Multivendor Interoperability**

One of the key goals of the Gigabit Ethernet Alliance is to provide end-to-end multivendor interoperability for Ethernet, Fast Ethernet, and Gigabit Ethernet equipment.

Initially, multivendor interoperability will be achieved using fiber optic cabling and Fibre Channel standards at the lowest levels of the Gigabit Ethernet standard. Standardizing on the use of fiber will help spur the early development of interoperable products, without impacting the deployment of Gigabit Ethernet into campus backbones and high-performance server environments. Vendors are committed to shipping interoperable products by the end of 1996 or early 1997.

The IEEE 802.3z standard will include specifications for running Gigabit Ethernet on Category 5 unshielded twisted pair. Due to greater technological challenges, twisted-pair implementations will follow the initial deployment of Gigabit Ethernet over fiber. The Task Force is also considering a 25 meter connection using twinax or similar cable. The standard will, once again, leverage Fibre Channel (FC-0) for this low-cost cabling solution.

### **Migrating to Gigabit Ethernet**

---

Gigabit Ethernet is a campus technology, designed initially for use in the campus backbone. Gigabit Ethernet will be used between routers, switches, and hubs. It will also be used to connect individual servers, server farms, and power workstations requiring more throughput to the high-speed backbone. Eventually, Gigabit Ethernet will be used to connect both shared and point-to-point workgroups.

### **Gigabit Ethernet Components**

To achieve the objectives of gigabit migration, four types of hardware are required:

- Gigabit Ethernet Network Interface Cards (NICs)
- Aggregating switches that connect 100Mbps Fast Ethernet and 1000Mbps Gigabit Ethernet segments
- All-Gigabit Ethernet switches
- All-Gigabit Ethernet Repeaters

In the next section, we will discuss how to use Gigabit Ethernet components to migrate a variety of network designs to Gigabit Ethernet.

### **Upgrade Scenarios**

In general, there are five scenarios for upgrading to Gigabit Ethernet:

- Upgrading connections to centralized file and compute servers
- Upgrading high performance workstations with Gigabit Ethernet NICs, and connecting them to Gigabit Ethernet switches or repeaters
- Upgrading connections between switches to use 1000 Mbps fat pipes between 100/1000 switches
- Upgrading a switched Fast Ethernet backbone by connecting Fast Ethernet switches with a Gigabit Ethernet switch or repeater

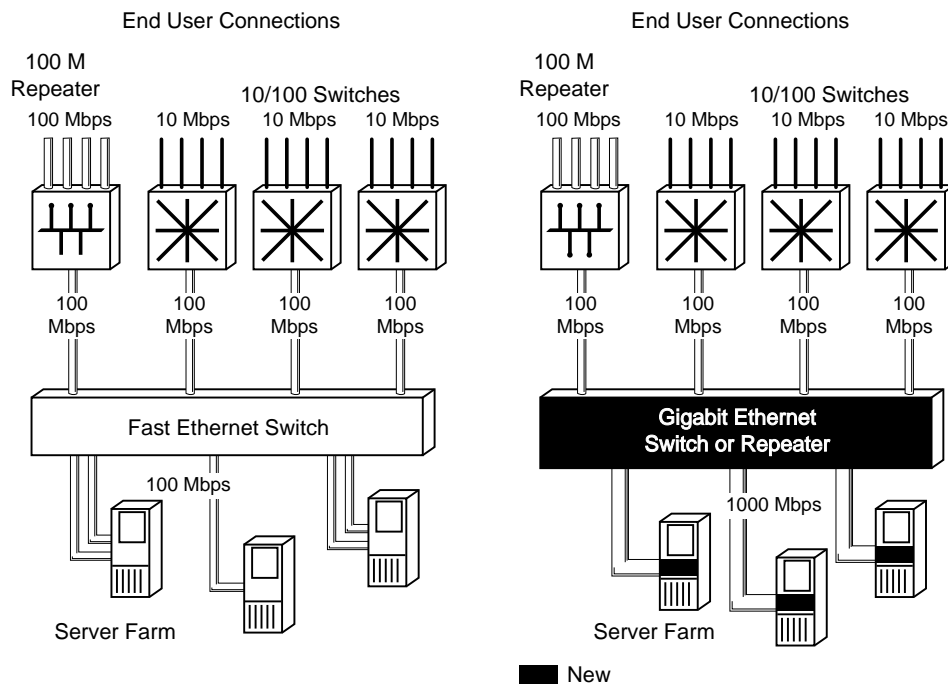
- Upgrading a shared FDDI backbone by connecting FDDI concentrators or Ethernet-to-FDDI routers with Gigabit Ethernet switches or repeaters

In all of the upgrade scenarios listed above, the network operating system, applications, and NIC drivers remain unchanged at the desktop. This allows companies to retain their investment in network management applications, tools, network hardware, and software. Existing multi-mode fiber runs, originally used for the existing backbone, are leveraged as well.

### **Upgrading Connections to Centralized File and Compute Servers**

The simplest way to take advantage of the increased speed afforded by Gigabit Ethernet is to upgrade one or more Fast Ethernet switches to an aggregating Gigabit Ethernet switch. 10/100 switches can be upgraded to 100/1000 switches, and high-performance super servers can be connected to the high-speed switches by installing Gigabit Ethernet NICs.

In Figure 4, the illustration on the left shows the original network. Some servers have multiple 100 Mbps connections to meet the bandwidth requirement of the applications. The illustration on the right shows the network using Gigabit Ethernet technology. The servers each have a single Gigabit Ethernet connection, simplifying network administration and providing much faster response time.



**Figure 4. Upgrading Connections to Servers and High-Performance Workstations**

### Upgrading Connections between Switches

Another simple way to take advantage of gigabit speeds is to upgrade the 100Mbps links between Fast Ethernet switches to 1,000Mbps. Fast Ethernet switches are replaced by 100/1000 aggregating switches. This allows the 100/1000 aggregating switches to support more switched and shared Fast Ethernet segments, without replacing any desktop hardware or software.

In Figure 5, the illustration on the left shows the original network, the one on the right shows the upgraded network.

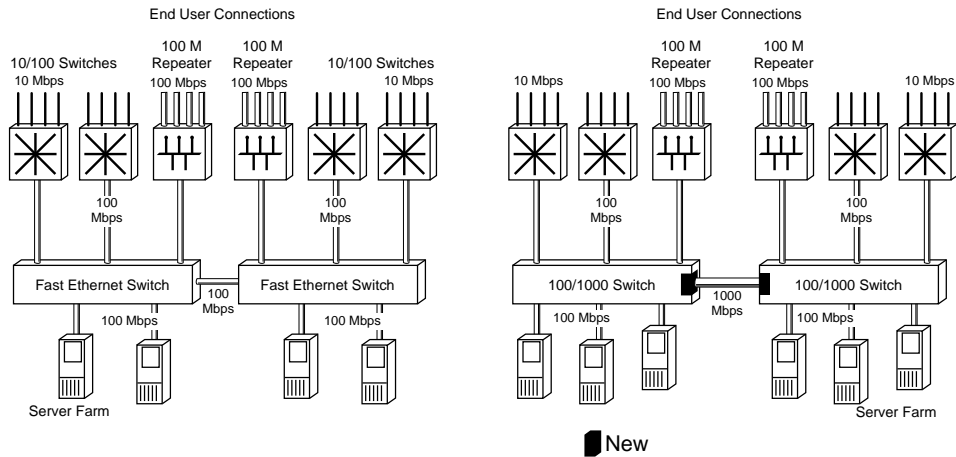



Figure 5. Upgrading Connections between Switches

### Upgrading a Switched Fast Ethernet Backbone

In this scenario, a Fast Ethernet backbone is upgraded to a Gigabit Ethernet switch or repeater. The 10/100 feeder switches are upgraded to 100/1000 aggregating switches. Once the backbone switch is upgraded to 1,000 Mbps, high-performance server farms can be connected directly to it by installing Gigabit Ethernet NICs. Additionally, the network can now support a greater number of segments, more bandwidth on each segment, and a greater number of nodes per segment.

In Figure 6, the illustration on the left shows the original network, the one on the right shows the network migrated to take advantage of Gigabit Ethernet.




**Figure 6. Upgrading a Switched Fast Ethernet Backbone**

### Upgrading a Shared FDDI Backbone

FDDI backbone networks are usually designed in one of two ways: a centralized backbone with FDDI concentrator and routers feeding into it, or a backbone that is distributed among a ring of FDDI concentrators or routers.

To upgrade either of these FDDI backbone designs to Gigabit Ethernet, the FDDI concentrators and routers are replaced with a Gigabit Ethernet switch. In the centralized backbone approach, the FDDI concentrator is simply replaced with a Gigabit Ethernet switch or repeater, and the routers are upgraded with a Gigabit Ethernet connection. Figure 7 shows the before and after illustrations of a centralized backbone.



**Figure 7. Upgrading a centralized FDDI backbone**

To upgrade a distributed backbone, each FDDI concentrator and/or router is replaced with a Gigabit Switch. Upgrades can be done on a case by case basis, or all at the same time. If the upgrade is accomplished one piece at a time, existing FDDI routers and concentrators will need to have a Gigabit Ethernet interface installed in order to connect to the Gigabit Switch. Figure 8 shows a distributed FDDI backbone upgrade. The one on the left is the original network. The network on the right shows the addition of a single Gigabit Ethernet switch.

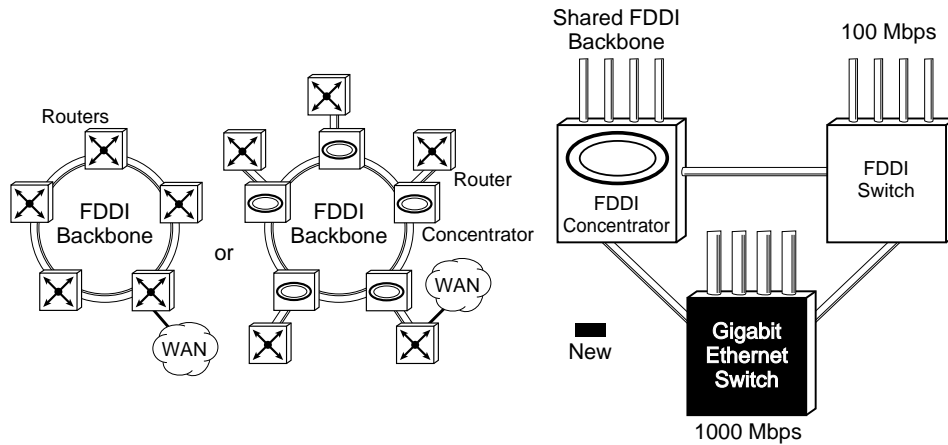


Figure 8. Upgrading a distributed FDDI backbone

## Protocol Architecture

Gigabit Ethernet uses a mixture of proven protocol technologies adopted from both the original IEEE 802.3 Ethernet specification and the ANSI X3T11 Fibre Channel specification. Figure 9 illustrates the way in which Gigabit Ethernet joins the two standards.

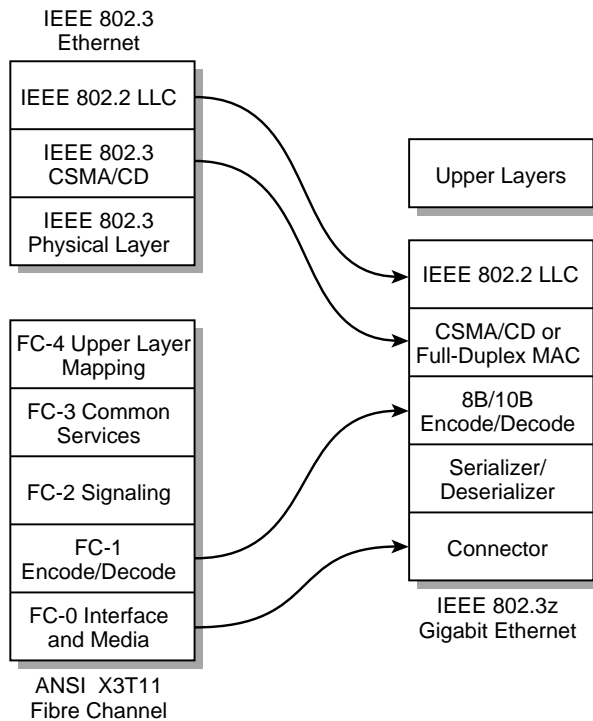


Figure 9. Gigabit Ethernet Protocol Stack

## **Physical Interface Characteristics**

The lowest layer of the initial fiber version of the Gigabit Ethernet protocol stack uses layer FC-0 from the Fibre Channel specification. FC-0 defines the physical characteristics of the interface and media, including the cables, connectors, drivers, transmitters and receivers. It also describes the speed and distance combinations of each media type.

## **Serializer/Deserializer**

On top of the Physical (PHY) layer is the serializer/deserializer. The serializer/deserializer is a layer that supports multiple encoding schemes—one of which is 8B/10B, as specified in the Fibre Channel specification. Encoding is media dependent; 8B/10B encoding was designed for fiber optic media. When support of twisted pair cabling is added to the Gigabit Ethernet specification, the serializer/deserializer will provide a mechanism for the encoding scheme used for twisted pair to work with the Gigabit Ethernet PHY.

## **8B/10B Encoding**

8B/10B encoding is another layer adapted from the FC-1 Fibre Channel specification. It describes the byte synchronization and the encode/decode scheme. The 8B/10B encoding scheme transmits eight bits as a 10-bit code group. 8B/10B encoding features low-cost component design and provides good transition density for easy clock recovery.

## **MAC Layer**

Gigabit Ethernet supports both full-duplex and half-duplex transmission systems. In full-duplex mode, the Gigabit Ethernet Media Access Control (MAC) uses the IEEE 802.3x Full-Duplex specification, including IEEE 802.3x frame-based flow control.

In half-duplex mode, the Gigabit Ethernet MAC supports the Carrier Sense Multiple Access with Collision Detection (CSMA/CD) access method, standardized in the original IEEE 802.3 Ethernet specification. It is expected that initial uses of Gigabit Ethernet will focus primarily on backbone connections using full-duplex transmission. Therefore, the first Gigabit Ethernet products to ship will support only full-duplex transmission. Support for half-duplex will be added later.

## **Upper Layers**

Above the MAC layer, Gigabit Ethernet is unchanged from the original IEEE 802.3 standard, supporting both Ethernet and IEEE 802.2 Logical Link Control (LLC). Support for LLC allows multiple upper layer protocols (including TCP/IP, SPX/IPX, and so on) to run over Gigabit Ethernet.

## **Cabling types and distances**

---

The IEEE 802.3z Task Force investigated a variety of cabling types and distances to be supported by Gigabit Ethernet. Table 2 summarizes the proposed cabling types and distances.

**Table 2. Cabling Types and Distances**

<b>Cable Type</b>	<b>Distance</b>
Single-mode fiber	2 kilometers
Multimode fiber	500 meters
Category 5 Unshielded Twisted Pair (UTP)	100 meters

Initial Gigabit Ethernet installations will use single-mode and multimode fiber optic cable. Support for 4-pairs of Category 5 UTP will be added later, with a goal of 100 meter cable runs.

## **Flow Control**

---

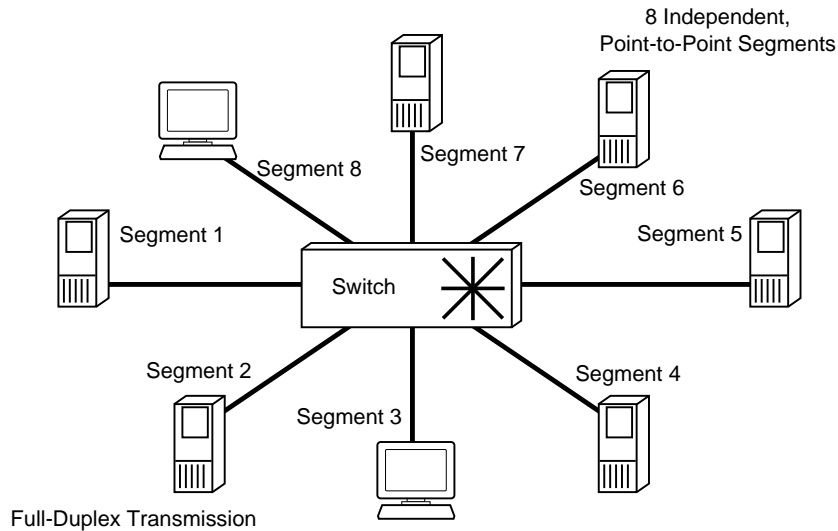
The flow control mechanism for Ethernet networks is different for full-duplex and half-duplex transmission types. The next section discusses these two transmission schemes, and the flow control mechanisms for each.

### **Full-Duplex Transmission**

Using full-duplex transmission, signals travel in both directions on the same connection, at the same time. Simultaneous bi-directional transmission allows the aggregate data rate of an Ethernet network to be doubled. For example, a 10Mbps Ethernet network can achieve an aggregate 20Mbps data rate, a 100Mbps Fast Ethernet network can achieve an aggregate 200Mbps, and a 1000Mbps Gigabit Ethernet network can achieve an aggregate 2Gbps data rate.

Full-duplex transmission is available for point-to-point connections only. Using full-duplex transmission, the issue of collisions on the network is completely eliminated and the CSMA/CD access control mechanism does not need to be invoked. Full-duplex can be used between a single workstation and a switch port, between two switch ports, or between two workstations. Full-duplex cannot be used for shared port connections, such as a repeater or hub port connecting to multiple workstations. Figure 10 shows a full-duplex, point-to-point network.





**Figure 10. Full-duplex, point-to-point Ethernet segments**

An optional flow control mechanism being defined by IEEE 802.3x, is available for full-duplex transmission and works in a way similar to XON/XOFF flow control. A receiving station at one end of the point-to-point connection can send a packet to the sending station at the opposite end of the connection instructing the sending station to stop sending packets for a specified period of time. The sending station ceases to transmit packets until the period of time has passed, or until it receives a new packet from the receiving station with a time of zero, indicating that it is okay to resume transmission.

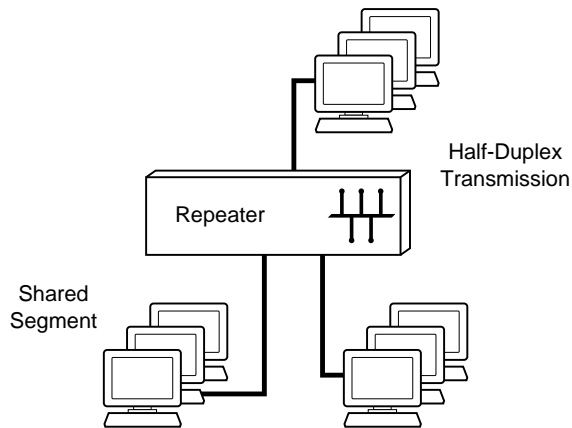
Full-duplex Ethernet is being standardized by the IEEE 802.3x committee. The full-duplex standard is not specific to any particular speed for Ethernet. It can be used for Ethernet, Fast Ethernet, and Gigabit Ethernet.

## Half-duplex Transmission

Using half-duplex transmission, signals travel in both directions on the wire, but not simultaneously. The original 802.3 standard specifies half-duplex transmission.


To gain access to the network in a half-duplex environment, Ethernet has traditionally employed Carrier Sense Multiple Access with Collision Detection (CSMA/CD) as the standard access method. Using CSMA/CD, a station waits for a clear channel. When it detects clear channel, it begins to send frames onto the wire. If two stations start sending data at the same time, a collision occurs. Each station must detect the collision, abort the transmission, and wait for a random interval of time before attempting to transmit data on the network again.

Half-duplex transmission is most commonly used on shared Ethernet segments. Unlike point-to-point connections, shared segments have two or more stations sharing a single port. Figure 11 shows a half-duplex, shared segment.



**Figure 11. Half-duplex, shared Ethernet segment**

Most switches manufactured today allow the user to select half-duplex or full-duplex on a port by port basis. This allows you to migrate your network from shared segments to point-to-point full-duplex segments over time. A switch port can be shared by front-ending the port with a repeater or hub. Figure 12 shows a network with both shared and dedicated switch ports.



**Figure 12. Ethernet network with full-duplex and half-duplex ports**

### Carrier Extension

In order to abide by the rules of the CSMA/CD access method, all stations sharing an Ethernet segment must be able to hear and detect a collision that has occurred on the network. More specifically, a station must hear the collision for the frame it is sending before it has completed transmitting the entire frame. The amount of time it takes for a station to send the frame the full length of the wire and have the jam signal resulting from a collision travel back to the station is known as the slot-time.

With very small (64 byte) frames travelling at speeds of 1,000 Mbps, the standard slot time used in the original IEEE 802.3 Ethernet specification is not long enough to accommodate a 100 meter cable run. The small frames are transmitted too quickly, and the sending station is finished transmitting the frame before it is aware of any collision that might have occurred at the other end of the segment.

In order to account for the problems of simply scaling CSMA/CD, and allow for minimum 64-byte frame sizes for compatibility across Ethernet devices, half-duplex Gigabit Ethernet implements a longer slot time using a technique called carrier extension. The frame size is not changed, but the time consumed on the wire is extended. Figure 13 shows a standard Ethernet frame and how extension is used, when necessary, to guarantee at least a 512-byte slot time.

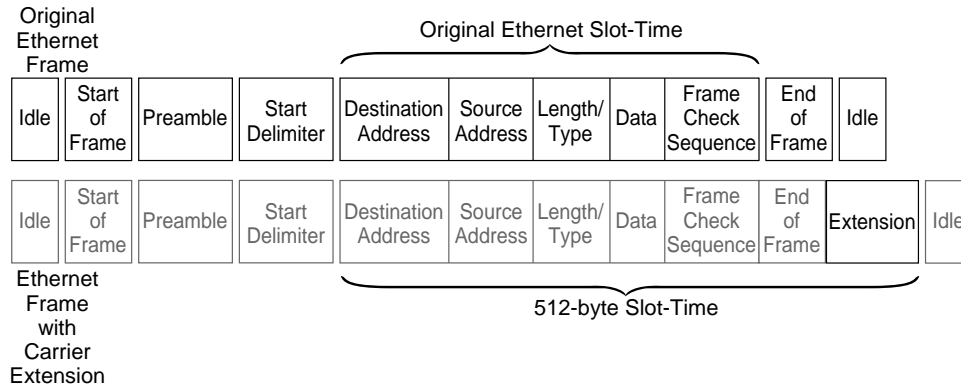


Figure 13. Ethernet Frame using Carrier Extension

## Technology Advances

### Ethernet Adapters

We have seen that workstations using the PCI bus have more than enough power to accommodate gigabit-speed network interface cards (NICs). Although the PCI bus is ready for Gigabit Ethernet NICs, traditional NIC technology is not suitable for Gigabit Ethernet speeds. The problem is that as network speeds accelerate past the speed of the CPU, utilization of the CPU can reach 100%.

Technology innovators are faced with two compelling questions:

- Which processes require the bulk of the CPU time?
- How can we reduce or even eliminate the CPU-intensive processes, thereby enabling the host to concentrate on application processing?

In terms of network interaction, there are two processes that use a great deal of host CPU time:

- Interacting with protocol layers—adding protocol headers, removing protocol headers, generating checksums, and so on
- Moving data within the memory system

Minimizing the involvement of the host CPU in both of these categories is a requirement for scaling Ethernet NICs to gigabit speeds. While it is not possible to completely eliminate protocol interaction and the need to move data within the memory system, it is essential to minimize the impact of each in order to run at gigabit speeds.

At Alteon Networks, we are leveraging years of experience in the gigabit networking arena to produce intelligent network adapters that will match the performance of Gigabit Ethernet. To run at gigabit speeds, we are taking large strides in NIC development, shifting the balance of processing power and time away from the host CPU and to the intelligent NIC itself.

To fully appreciate this technological leap, we will briefly examine first and second generation Ethernet NIC technology. Then we will explore the Alteon Gigabit Ethernet solution.

### First Generation Ethernet Adapters

Originally designed for the ISA bus, first generation Ethernet adapters are unintelligent cards designed with rudimentary functionality. All of the processing is handled exclusively by the host CPU. Figure 14 shows the processing sequence for data being transmitted by a first generation NIC.

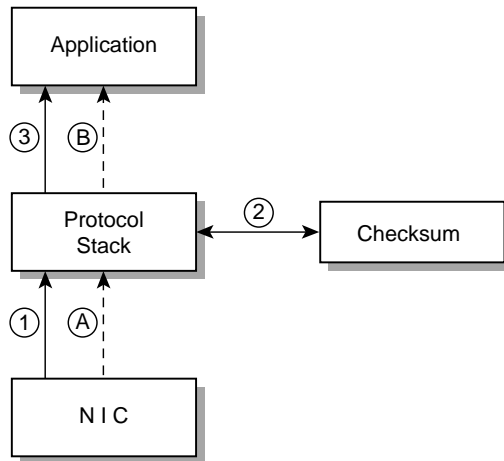


Figure 14. Data reception using first generation NIC

The process of transmitting and receiving data using a first generation NIC requires each packet to be copied multiple times and requires multiple interrupts to be issued for notification along the way. Data reception using a first generation NIC is described in Table 3.

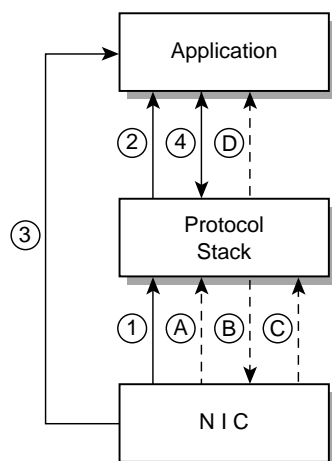
Table 3. Steps in Data Reception using First Generation NIC

Data Steps	Control Steps
1. The NIC passes the data, as it arrives, to the protocol stack entity on the host using a pre-allocated buffer.	A. The NIC notifies the stack that it has transferred the data by issuing an interrupt.
2. The protocol stack performs a checksum on the data.	
3. The protocol stack moves the data to the application memory.	B. The protocol stack informs the application that data has arrived.

First generation Ethernet adapters not only rely exclusively on the host CPU, they also require multiple data copies and interrupts to get the job done. Clearly, this technology is slow and cumbersome, and cannot scale to higher speed networks.

## Second Generation Ethernet Adapters

With the advent of the faster applications that needed more CPU power, second generation Ethernet adapter manufacturers recognized the need to limit the amount of host CPU time required to move data around the network. In response, they changed the processing sequence of the NIC by implementing a concept known as “look ahead” buffers, as shown in Figure 15.



**Figure 15. Data transfer using second generation NIC**

Data reception using a second generation Ethernet NIC is described in Table 4.

**Table 4. Steps in Data Reception using Second Generation NIC**

Data Steps	Control Steps
1. The NIC moves the first 64 bytes of the packet to the protocol stack through a pre-allocated buffer. The first 64 bytes includes the header information and some data.	A. The NIC notifies the stack that it has moved 64 bytes of data by issuing an interrupt.
2. The protocol stack moves any data from the initial 64 bytes (minus the header) to the application memory.	B. The protocol stack analyzes the headers and tells the NIC where in application memory to put the remaining data from the packet.
3. The NIC moves the rest of the data into application memory. If application memory is not accessible, then an intermediate buffer is used and the data is copied to the application memory by the host in Step 4.	C. The NIC tells the stack that it has finished moving the rest of the data packet into application memory by issuing an interrupt.
4. The protocol stack performs a checksum on the packet in the application memory space.	D. The protocol stack informs the application that data has arrived.

Transferring only the first 64 bytes of information to the protocol stack, instead of the entire packet, represents a good stepping stone to minimizing data copies and alleviating a small portion of host CPU time. Unfortunately, this method creates a new problem—an increased number of interrupts required among the three processing entities.

The increased number of interrupts needed by NICs using look ahead buffers compromises the ability of these cards to scale to gigabit speeds. As network speeds outpace the speed of the CPU, NICs will request many more interrupts per second. Eventually, the CPU will be unable to keep up with the number of interrupts being issued and will end up dropping packets.

### **Third Generation Ethernet Adapters—The Alteon Difference**

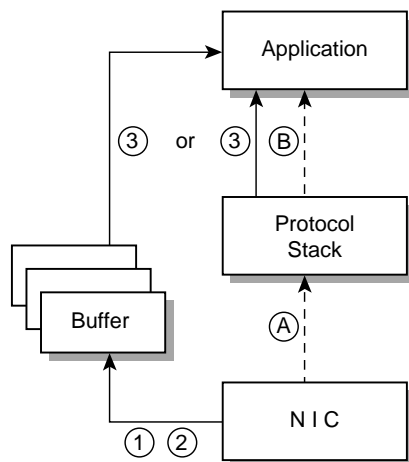
Clearly, a new methodology for moving data between the NIC, protocol stack, and application memory is needed to support Gigabit Ethernet speeds. Alteon Networks is pioneering the next generation of Ethernet NICs in the gigabit arena.

Using an intelligent adapter with an onboard RISC-based processor specially designed for embedded application processing, Alteon's Gigabit Ethernet technology not only reduces the number of times data is copied among processing entities, it allows a single interrupt to be issued for multiple data packets—radically altering the ratio of interrupts to packets, and eliminating the scalability problems inherent in older adapter designs.

Alteon's technology advantages center around the following concepts:

- Off-loading TCP/IP checksum
- Interrupt coalescing
- Scatter/gather
- Buffer alignment for virtual page swapping
- Adaptive DMA on the PCI bus
- Overall reduction in host buffer copying

Figure 16 illustrates the concepts listed above, and illustrates how data is received by the Alteon Gigabit Ethernet NIC and transferred to application memory space.



**Figure 16. Data reception with Alteon Networks third generation NIC**

Data reception using an Alteon third generation Ethernet NIC is described in Table 5.

**Table 5. Steps in Data Reception using Alteon's Third Generation NIC**

<b>Data Steps</b>
1. The NIC extracts the header information from the first packet, and puts it into a buffer space provided by the host operating system.
2. The NIC puts the packet data into a separate buffer space provided by the operating system. While doing Steps 1 and 2, the NIC performs a checksum on the data.

Repeating Data Steps 1, 2, and 3, the NIC puts as many packets as it can into the buffer until one of the following occurs:

- It runs low on buffer space
- The interrupt timer expires

Alteon Gigabit Ethernet adapters have an interrupt timer that determines when to interrupt a host CPU. This allows a single interrupt to be issued for multiple data packets that are sent into the operating system buffer space. It also allows for “adaptive” interrupts; that is, the NIC can alter the number of interrupts issued per second based on network usage.

The NIC card is said to be at idle during any period of time with no work to be performed that extends past the interrupt period. The next time an event occurs that requires the attention of the host, the NIC does not wait for the interrupt timer to expire. Instead, it begins to act on the data immediately. This is another example of adaptive interrupts.

In a busy environment, issuing a smaller number of interrupts over a regularly spaced time interval greatly improves network performance. In a light environment, it is preferable to issue interrupts more frequently to reduce latency. Alteon adapters allow you to configure the optimum number of interrupts for heavy loads on your network, while it dynamically evaluates network activity to determine which method is most appropriate for normal traffic loads.

Control Step A: The NIC issues a single interrupt, on a timed basis, telling the protocol stack how many packets it has copied to the buffer.

Data Step 3: The protocol stack is responsible for ensuring that the data arrives in the application memory buffer. This happens in one of two ways:

- If the application memory is also aligned on a page boundary, the protocol stack and the application memory swap pointers to the data—eliminating the need for the data to be directly copied to the application memory. Swapping memory pointers is possible because of the groundwork that was laid in Data Step 2.
- If the application memory is not aligned on a page boundary, the data is copied to the application memory by the protocol stack, and

Control Step B: The protocol stack informs the memory that data has arrived.

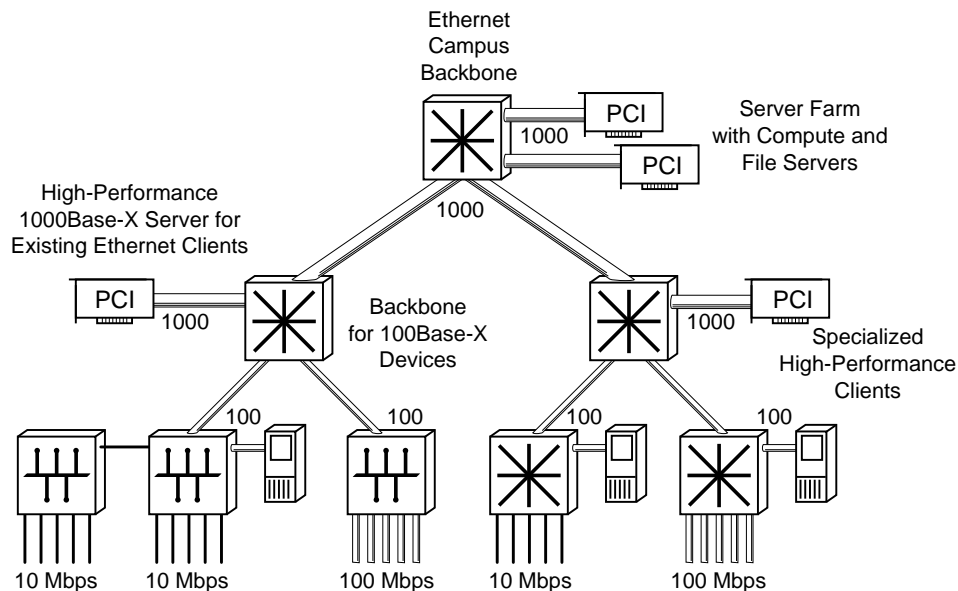
Using groundbreaking concepts such as those listed above, Alteon Networks is clearly distinguishing itself as the leader in Gigabit Ethernet technology. Many years of experience in gigabit networking have allowed us to lead the market with a complete ASIC, NIC, and driver solution for moving data at gigabit speeds.

## Ethernet Switches

Networks are changing quickly to improve the speed and simplicity of data access and enable new applications.

- 10/100 adapters in desktop workstations are providing the bandwidth required by graphic-intensive applications, faster CPUs and faster buses
- Fast Ethernet switching has become the standard switching technology for organizations that have Ethernet on the desktop
- Intranet servers are being expanded to deliver corporate information, training, and on-line documentation, leveraging rapidly improving browser technology and data access protocols
- Multimedia to the desktop is now coming of age with new protocols such as RTP and RSVP for users accessing streaming voice/video traffic

All of these changes will drive the demand for faster connections to centralized servers as well as faster connections between switches. 10/100/1000 switches will provide an easy way to integrate existing Ethernet and Fast Ethernet with Gigabit Ethernet connections. These switches will be used to connect servers to today's networks and can be used in a backbone to interconnect Fast Ethernet switches. The all-gigabit switches provide the ultimate "backbone-in-a-box" for Ethernet networks in a building or campus environment. Figure 17 illustrates Alteon Gigabit Ethernet switches added to an existing Ethernet and Fast Ethernet network.



**Figure 17. 10/100/1000 and All-Gigabit switches**



### **Specialized ASICs to meet the Performance Demands of Gigabit Ethernet**

The majority of initial Gigabit Ethernet products will have full duplex Gigabit Ethernet ports. This is an aggregate of 2 Gigabits/second of traffic over each Gigabit Ethernet connection. In lab tests, using minimum 64 byte frames, each Gigabit Ethernet port will be able to send and receive a total of nearly 3 million packets per second! These performance levels will clearly overwhelm most internetworking switches and routers shipping today.

The Alteon switch architecture is based on intelligent ASICs. Each ASIC includes hardware DMA engines, a 32-bit RISC processor and a complete Ethernet MAC. These highly integrated ASICs are used to provide the high performance, low cost, and reliability required for today's switched networks. The RISC processor is involved in higher-level functions, including packet tracing, accumulating statistics, and other management tasks. Note that every ASIC has an embedded processor to handle these functions, eliminating the bottleneck of a central CPU.

Dedicated switching hardware in each Alteon ASIC is used to handle forwarding, learning, and filtering. Virtual LANs, often used to improve security or to reduce network traffic, are configurable by port number, MAC address, Ethertype, etc. VLAN tagging and frame priorities are supported in the hardware, with final support pending resolution of standards. Spanning Tree support is used for easy installation into existing networks.

### **Seamless Integration with Ethernet and Fast Ethernet**

Gigabit Ethernet utilizes the same frame format as Ethernet and Fast Ethernet, making it easy to forward traffic between dissimilar networks. Using the same frame format avoids the cost, complexity, and performance penalty found when translating between Ethernet and ATM or another network technology. More importantly, this allows Gigabit Ethernet connections to be added only where the performance is needed, maintaining the enterprise's investment in existing NICs, switches, and routers.

Alteon switches are completely interoperable with Ethernet and Fast Ethernet. A dynamic buffering scheme is used when forwarding traffic between network connections. Input and output buffers are dynamically allocated to adapt to changing traffic conditions, minimizing latency and improving overall throughput. And, in order to deliver the maximum performance to all networks and servers, any port on an Alteon switch can run full duplex.

### **Standard Ethernet Management**

Alteon switches are also designed to fit into standard management environments to leverage the investment in training and equipment. Full RMON and SNMP support is included to be compatible with standard network management stations. A private Alteon MIB will be included to deliver product specific configuration and statistics. Remote management and configuration can be even be done through a web browser from a standard desktop workstation.

## **Conclusion**

---

Network speeds are on the rise. Workstation technology has advanced to support gigabit data transfer. Multimedia protocols are in final stages of standardization. Corporations are insisting on leveraging the billions of dollars already spent on networking equipment.

Ethernet is here to stay, at a variety of speeds and for a variety of purposes. Gigabit Ethernet technology, while retaining complete interoperability with all of its Ethernet predecessors, is the natural extension in the ever persistent need for speed.

Based on a combination of Fibre Channel standards at the lowest layers and standard Ethernet above, Gigabit Ethernet provides ultra high-speed backbone connections, high-powered server connections, and, eventually, cost-effective workstation connections on the corporate network.

By combining time-proven, existing, and widely-implemented Ethernet technology with an innovative approach to NIC data transfer, Alteon Networks is leading the industry to market with technology to meet the needs of the future in its soon-to-be-released Gigabit Ethernet switches, NICs, and repeaters.