Increasing Network Bandwidth with Switched and Full Duplex Ethernet

adaptec®

The need for network bandwidth

Today's networks are performance hungry. The explosion of network technology has resulted in more users, more applications, and more intensive use of corporate networks. Networking applications range from the mundane to the inspired, but even traditional applications such as word processing continue to demand more bandwidth from the network. Whether loading an application, printing a file, or conducting an interactive video conference, networks are being challenged to provide the highest possible throughput in the most cost-effective manner.

The need for increased network performance has provided the catalyst for the development of a number of technologies. In fact, the sheer number of technologies developed to tackle the problem suggests its size. The most popular new technologies include ATM (Asynchronous Transfer Mode) and Fast Ethernet (or 100 Megabit Ethernet), and in the future, Gigabit Ethernet. These technologies are receiving a good deal of attention. In addition, two developments, network switching and Full Duplex Ethernet, are further extending the capabilities of Ethernet. Today, Ethernet accounts for up to 75% of all LANs and accounts for more than 89 million network nodes annually. (Source: IDC 1996)

The evolution of Ethernet

As networking has grown increasingly vital to businesses, Ethernet has continued to evolve. The need for manageable cabling spawned the introduction and acceptance of an industry-standard structured wiring system, 10BASE-T, the Ethernet standard for unshielded twisted pair (UTP) cabling. 10BASE-T has become the de facto standard for new Ethernet nodes in the U.S.

With the advent of Ethernet switching, Ethernet performance has further evolved. Switched Ethernet is completely compatible with existing coax and UTP cable installations. Perhaps best of all, switched Ethernet opens the door to Full-Duplex Transmission (FDE) . FDE is a cost-effective technology that can deliver more than double the performance of traditional Ethernet.

Fast Ethernet, which uses the same packet format as Ethernet but delivers data 10 times faster, became a networking force in 1995. Fast Ethernet is expected to continue capturing networking market share, with more than 100% growth through 1997. (IDC 1996). Fast Ethernet maintains backward compatibility with standard Ethernet through network adapters and switches that support both 10 and 100 Mbps. Both Full-Duplex transmissions and switching technology are also supported by the Fast Ethernet specification.

Ethernet switching technology

Ethernet's original design was based on a bus topology. In a bus topology, all users have access to a single network cable at their own discretion, rather than being granted access (or being passed a "token"). With Ethernet, (and now Fast Ethernet) this access is arbitrated through a mechanism called CSMA/CD, or Carrier Sense Multiple Access/Collision Detection.

CSMA/CD ensures that only one network node on the Ethernet bus is transmitting data at any given time. All nodes can receive simultaneously, but only one can transmit. For this reason, any node that wants to transmit data must first listen to the cable. If the cable is not in use, it can then send its data. If the cable is busy, the node must wait a few moments, then begin the cycle again. In practice, this causes only small delays. Greater access delays occur when the cable appears to be available and two nodes both try to transmit at the same time. This results in an event called a collision. Collisions are normal in the CSMA/CD mechanism, but will cause transmission delays, especially in networks with many active nodes.

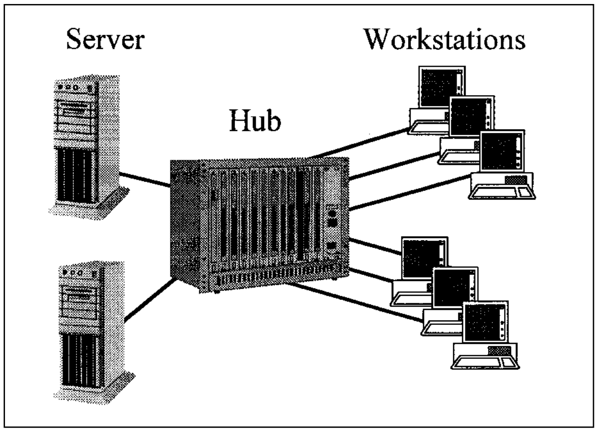

Conventional Ethernet configurations must share 10 Mbps of bandwidth

With the addition of an Ethernet switch, an existing network is dynamically divided into multiple parallel networks. The result is a multiplicative effect on throughput. Theoretically, using an Ethernet switch can increase the maximum network throughput from 10 Mbps to a combined rate of 70 to 150 Mbps. The same multiplicative effect occurs with Fast Ethernet switches, only at even higher throughput rates.

Switched Ethernet can reduce or completely eliminate collisions, depending on the implementation design. Since an Ethernet switch consists of multiple "ports," workstations, network servers, or an entire Ethernet segment can be plugged into each port. Every device that is plugged into the switch has access to every other device on the switch, with a collision-free full 10 or 100 Mbps connection. When the switch detects a transmission on one of its ports, it creates a physical connection between the sending port and the destination port. For the brief time that the data is flowing from the sender to the receiver, this connection is completely private. No other stations are on the same private "cable," hence there is no possibility of collision.

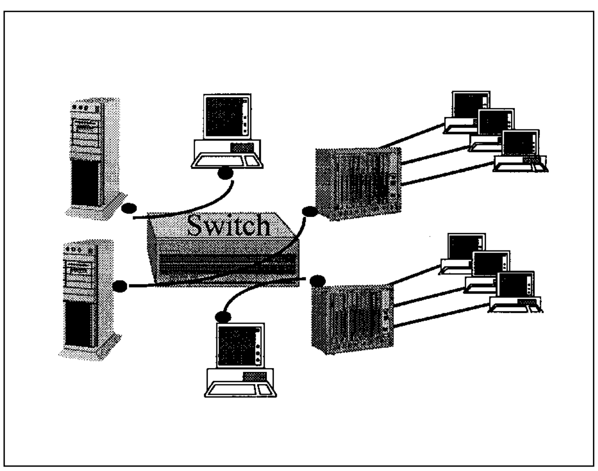

Switched standard Ethernet and Fast Ethernet configurations can have multiple concurrent 10 or 100 Mbps connections achieving rates as high as 150 or 1500 Mbps respectively.

Today's switches support many concurrent conversations. Stated another way, up to 1500 Mbps can pass through the switch at once—clearly a huge improvement over non-switched environments. The physical implication is that a network with a switch can use this increased bandwidth to provide faster access or to expand into a larger network.

Switched Ethernet can support connections to any Standard Ethernet medium, including thin coax, thick coax, unshielded twisted pair (10BASE-T), or fiber optic (10BASE-FL) cabling types.

In addition, for environments that include both standard and Fast Ethernet, switches also serve as bridges between 10 and 100 Mbps line speeds. For CSMA/CD to work properly, all devices sharing a common bus must understand the signals from every other device on the same bus. Since devices on a repeater share a common bus, they must all run at the same line speed. However, a switch creates a separate bus on each of its ports. So, devices plugged into different ports on the switch can run at different line speeds.

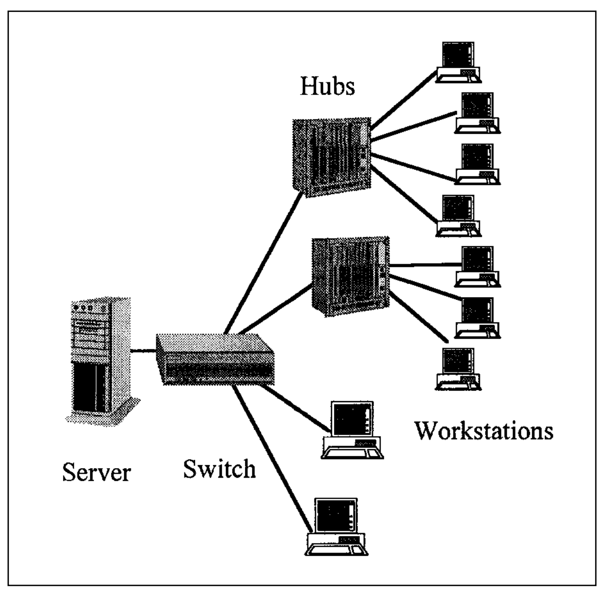

Both hubs and workstations can be connected to an Ethernet switch.

Differences between switches, routers and hubs

An important feature of the Ethernet switch is its ability to quickly switch packets without introducing delays. The best Ethernet switches are able to handle multiple continuous flows of 10 or 100 Mbps traffic, without needing to store each packet and then forward it to its destination.

Routers are examples of "switches" that have high propagation delays, but routers serve a different function than switches. Routers are used to filter traffic or translate protocols in order to transfer data packets between stations located on different LANs. The router must read and understand the entire packet before it can be processed, so it acts as a store-and-forward message relay service. This may introduce considerable delays in the flow of traffic. Routers use an internetwork protocol that is recognized by the attached LANs, while Ethernet switches are protocol-independent.

Like a hub, a switch can be the central connection point of a network. However, a hub is the central device in a star-configured LAN segment. Each station in the segment is connected directly to the hub, through which all data must pass at the same line speed. A typical LAN may connect several hub-centered segments to one or more servers. These hub-dependent star networks mean limited cable access for each user. Since the hub is involved in every data transaction, it, too, becomes part of the processing bottleneck. Instead, if each of these hubs and servers is interconnected through a switch, considerable throughput gains are possible.

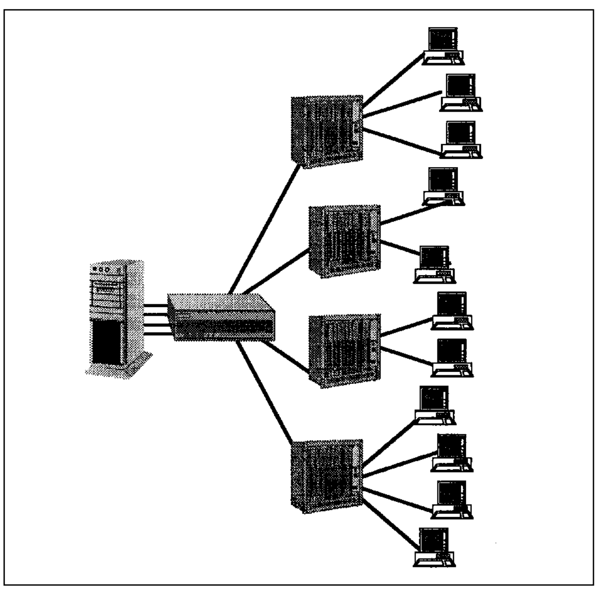

Multiple server adapters have a multiplicative effect on bandwidth.

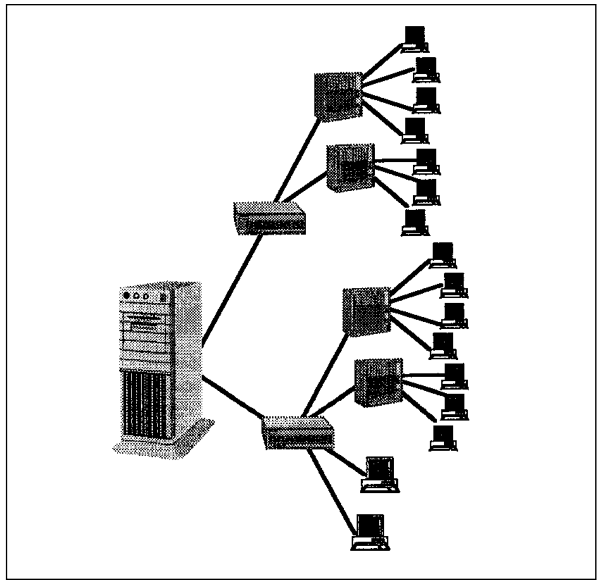

Scalability is an important benefit of Ethernet switching.

Multiple adapters

A cost-effective means of increasing performance is placing more than one adapter in the server. Each of these adapters can access the server at full bandwidth, taking better advantage of the server's I/O and CPU capacity. Each is connected to a separate switch port or separate LAN segment if a switch is not present.

Multiport Ethernet adapters

Throughput can be further increased by using one or more multiport Ethernet adapters in place of the traditional single-port adapters in the server. These adapters take advantage of the 132 MByte/sec bandwidth and sophisticated architecture of the PCI Local Bus machines that are now commonly found in servers. The Adaptec® four-port Cogent™ Quartet™ adapter uses the cascadable peripheral bus architecture that is unique to PCI. Quartet presents four adapters to the server's CPU through a single PCI slot. Therefore, traffic through this single slot may be increased four-fold. Multiply this by the number of PCI slots that are available-usually three-and you significantly increase the number of LAN segments that can directly access the server.

Load balancing

Today's high-performance servers and LAN adapters may be capable of extremely high throughput, but I/O bottlenecks can still occur between the server adapters and their connections to the many workstations in the network. Fortunately, load-balancing software can be used to further increase server throughput. The software resides on the network server at the transport protocol layer to distribute network traffic evenly between any number of network segments connected to the server.

With incoming data, the load-balancing software assigns users to different adapters in a round-robin fashion as they log-in to the server. This effectively broadens bandwidth by a factor equal to the number of adapters on the server. With traffic leaving the server, this multiplexing action serves as an internal router that intercepts packets sent by the server and distributes them across multiple interface adapters onto the network. To the server, these multiple adapters appear as a single network interface. To the workstation, the server appears more immediately available, i.e., without the delay-producing congestion typical during high-access periods.

Implementation of load-balancing requires a server connected to multiple segments, an Ethernet or Fast Ethernet switch, and the load-balancing software itself. All network adapters and switches support load-balancing. Quartet adapters are particularly suited to load-balancing because they provide the necessary multiple segments, all on one adapter. Load-balancing software is operating system and protocol specific. As of this writing, load-balancing software is available for NetWare servers version 3.12 and 4.x. Novell, Cisco, and NSI all offer IPX load-balancing modules in NetWare Loadable Module (NLM) format.

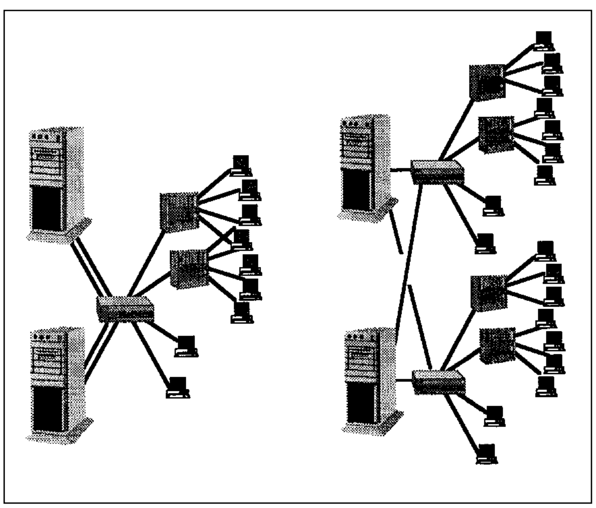

Multiserver topologies can be designed for cost efficiency or fault tolerance.

Full Duplex Ethernet

As Ethernet has evolved we have seen the physical media, thick coax, being changed to thin coax, 10BASE-T, and fiber optic cable types. These changes have created new opportunities to take advantage of the unique features of the cable.

When Ethernet relied on coax cable, all traffic traveled on the same two wires. This meant that the CSMA/CD mechanism was critical to preserving order on the cable. With the introduction of 10BASE-T and fiber-optic cabling, the channels for sending and receiving were separated. The collision-detection model was retained because all of the network nodes were still wired together with hubs in a logical bus topology. When a switch is added to the network, collisions between nodes or subnets are eliminated. When the collision-elimination attribute of the switch is combined with the send/receive channel separation of 10BASE-T, a new transmission mode is possible-Full Duplex Ethernet (FDE). FDE allows stations to send and receive simultaneously. Collision detection remains important, however, for any stations that are connected to a hub because they do not have the collision protection afforded by the switch.

Fast Ethernet was developed for use with twisted pair wiring from the beginning which means that Fast Ethernet exhibits the send/receive channel separation required for Full-Duplex transmissions. In addition, Fast Ethernet now supports multimode fiber optic cable. Since this cable also provides separate cable channels for send and receive, it too supports Full-Duplex transmissions. The only Fast Ethernet standard that does not support Full Duplex is T4 Fast Ethernet.

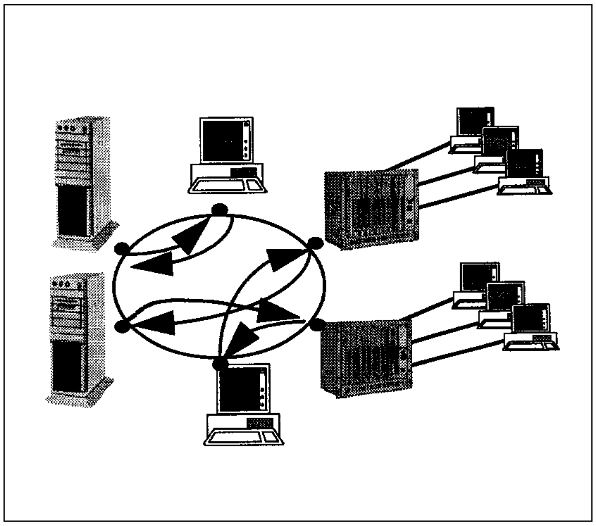

Full Duplex Ethernet allows simultaneous bi-directional connections at 20 Mbps

Full Duplex increases Ethernet's throughput by creating two collision-free 10 Mbps paths, one for sending and one for receiving, between two Ethernet stations on a single UTP segment. It is simply switched Ethernet with the collision-detection feature removed. Many of today's Ethernet vendors have designed and manufactured their Ethernet adapters and drivers to support Full Duplex, but not all of these manufacturers support Full Duplex on both 10 Mpbs and 100 Mbps. All of Adaptec's Cogent products for PCI support Full-Duplexing, both at 10 Mbps, and at 100 Mbps where appropriate. (Again, twisted pair cabling is required; and for 100 Mbps Full-Duplexing, TX products are required.) In addition to the appropriate hardware to support FDE, vendors must also provide multi-threaded device drivers that can handle the simultaneous sending and receiving of data packets.

By allowing a network node to simultaneously send and receive data, Full Duplex Ethernet doubles the theoretical bandwidth of switched Ethernet from 10 Mbps to 20 Mbps. How does this increase in cable bandwidth translate into performance benefits? The answer, although it seems obvious, will be examined next, in the context of typical network operations.

How Full Duplex benefits servers

The first benefit takes place in the network server. Generally the hardest working component in a network, the server must both receive requests and respond by sending data. While data requests comprise relatively few packets, returning the requested file usually means many packets of data. Therefore, file-server traffic tends to be weighted more heavily to sending data (ignoring, for a moment, file saving and print spooling).

When the file server sends without FDE, it is unable to receive data requests from other workstations. If the file server was able to receive requests, it could begin assembling the response (as a background task), even while previous send requests are being processed. FDE allows the server to simultaneously receive requests and queue sends—a definite throughput enhancement. Also, most networks use a Send-Acknowledge method for data transmission. After a packet (or multiple packets) is sent, an acknowledgment is returned by the receiver. Using FDE, this acknowledgment can be returned on the receive channel, without interfering with the "send."

FDE and workstations

Workstations also receive substantial benefit from FDE, even when they are not equipped with FDE adapters and are operating in half duplex mode. When their server is FDE-equipped, workstations can transmit their requests (or save and print files) even while their server is busy transmitting data to another node. Conversely, they can receive their data even while the server is receiving another node's request. At the present time, however, there is little reason for most workstations to have FDE adapters themselves. With the notable exceptions of OS/2 and NT machines, most workstations do not have multitasking operating systems or applications which would effectively utilize both the send and receive channels.

Consider a workstation running a word processing application. The application is either reading (opening a file), or writing (saving or printing a file). Word processing applications are not designed to multitask these operations to, say, open one file while printing another. The workflow of these events is sequential, not concurrent. In the very near future, however, there will probably be applications, such as two-way video conferencing, that will require workstations to send and receive data at the same time.

FDE doubles bandwidth by taking advantage of the collision-eliminating feature of the Ethernet switch. It allows data to be simultaneously transmitted and received on different twisted-wire pairs. With FDE and multi-threaded networking drivers, servers become multitasking systems capable of handling larger amounts of traffic than possible with traditional half duplex Ethernet. This increased bandwidth benefits network users by providing workstations with faster service.

Summary

The need to expand the capabilities of Ethernet to support network applications with high volumes of data has fostered an innovative enhancement of the technology. From the early development of the economical 10BASE-T wiring standard, routers have been upstaged by Ethernet switches that provide dedicated links between network nodes. Multiport adapters and load-balancing software have improved server performance by an order of magnitude or more. Now, simply by using Ethernet switches in conjunction with Adaptec's high performance Cogent Ethernet and Fast Ethernet adapters, networks benefit from a far more cost-efficient option—Full Duplex Ethernet.

FDE: More bandwidth, less money

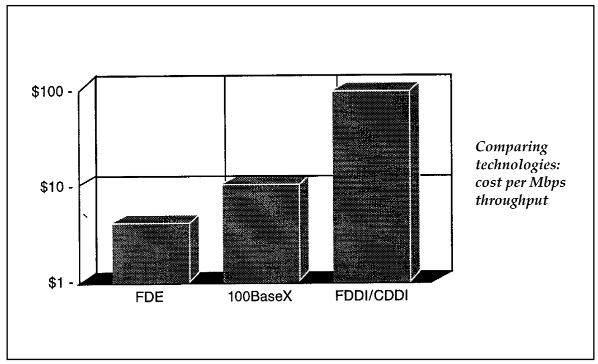

Compatibility with existing 10BASE-T networks means that the cost of upgrading to a full duplex configuration is dramatically less per dollar of throughput. An FDE network can give twice the throughput of half duplex networks using standard twisted pair cable, LAN adapters, and 10BASE-T hubs. If each Mbps of throughput is measured in dollars, FDE is money in the bank.

Adaptec, Inc.
691 South Milpitas Boulevard
Milpitas, California 95035
Tel: (408) 945-8600
Fax: (408) 262-2533

Adaptec Europe-Belgium
Tel: (32) 2-352-34-11
Fax: (32) 2-352-34-00

Adaptec Japan-Tokyo
Tel: (81) 3-5276-9882
Fax: (81-3) 5276-9884

Adaptec Singapore
Tel: (65) 278-7300
Fax: (65) 273-0163

Literature
1-800-934-2766 (USA and Canada)
or (510) 732-3829

Ordering Software:
1-800-442-7274 (USA and Canada)
or (408) 957-7274

World Wide Web:
http://www.adaptec.com/networking

Internet ftp server:
ftp.adaptec.com CompuServe:
GO ADAPTEC

Microsoft® Network:
GO ADAPTEC

Adaptec USA Bulletin Board Service (BBS):
(408) 945-7727 (up to 28,800 baud, using 8 bits, 1 stop bit, no parity)

Interactive Fax:
(408) 957-7150

P/N 700256-011 Printed in USA 9/96