CHAPTER 2

An Overview of High-Speed LAN Technologies

As we write this book—and in particular this chapter—we are making certain assumptions about you, the reader. We assume you are using 10Mbps Ethernet as a LAN technology. Because Ethernet is approaching 90% market share today, we think that's a safe assumption. If you aren't using Ethernet today, you are probably considering it; otherwise, you wouldn't have bought this book. We also assume you want to migrate to a higher-speed LAN technology at some point because whatever you are using today is running out or has run out of bandwidth. Obviously, one of the higher-speed Ethernet derivatives is on your shopping list, which may be why you bought this book. Therefore, the big question is this: Which LAN technology is the right choice? As a LAN manager today, you essentially face two choices for high-speed technologies and standards: Ethernet or Asynchronous Transfer Mode (ATM). That hasn't always been the case. As recently as three years ago, about half a dozen different high-speed contenders were vying for a share of the high-speed networking hardware market.

This chapter provides a broad overview of high-speed LAN technologies, past and present. The discussion focuses primarily on a comparison between ATM and Ethernet. In the end, we hope to convince you that Ethernet is the only choice.

This chapter is divided into three sections. Because choosing the right high-speed technology to migrate to is a strategic decision that requires careful consideration, the first section of this chapter discusses the different selection criteria an Ethernet LAN manager needs to think about when selecting his or her next-generation high-speed LAN standard.

The second section discusses the different high-speed technologies that have at one point or another attempted to compete with Ethernet: FDDI, TCNS, Iso-Ethernet, 100VG-AnyLAN, Fibre Channel, and HIPPI. Although you might consider some of this material redundant, proponents of some of these technologies positioned them as being viable alternatives to Ethernet as recently as three years ago. We can learn some valuable lessons by considering why none of these technologies have been commercially successful.

The third section of this chapter contrasts Ethernet with ATM (the only two viable contenders today). In this section, we also discuss why we think Ethernet is the smart choice for local area networking hardware.

We wrote this book from the perspective of an Ethernet user. Because we do realize that about 10% of all networks still use Token Ring, we have included a short section for Token Ring LAN managers at the end of this chapter, including some discussion of Switched and High-Speed Token Ring.

A Historical and Somewhat-Biased Perspective

This chapter briefly contrasts several high-speed technologies. Some of these technologies, such as Iso-Ethernet, never even made it off the drawing board, some did and have since disappeared, and some are still shipping but are in a state of decline. The notable exception is, of course, ATM. This chapter covers ATM in detail.

A review of the three editions of this book helps illustrate how the high-speed networking landscape has changed over the past five years. You will find a strong element of "we told you so" in this chapter—only because we have been advocating the superiority of Ethernet for many years, and will continue to do so. This third edition is no exception. Now we have the benefit of hindsight, having seen various competitors come and go. We also aren't going to pretend that we are offering you an unbiased view of technology. Some authors or journalists may pretend to offer you "just the facts," but even facts often contain a huge element of bias you may not recognize as such. We would never reprint the results of a reputable market research firm that concluded that ATM would outship Ethernet in 1999, for example. These might be considered facts, but we choose to find "our own set of facts" instead: a market research report that concluded the opposite. You would be impressed with our data, of course, and would never know that some study out there had concluded the exact opposite!

M

When we wrote the first edition of this book in 1995, Fast Ethernet was just being introduced. In that edition, we extensively debated the pros and cons of 100BASE-T versus 100VG-AnyLAN, the competing proposal from Hewlett-Packard. Well, 100VG-AnyLAN has come and gone. At the time, most people used the term *coexistence*, with 100BASE-T taking a slightly larger market share. We didn't buy into the coexistence theory because people migrate to a single standard. Our favorite article on 100VG-AnyLAN was one titled "100VG-NowhereLAN," which was written by Kevin Tolly and printed in the November 1994 issue of *Data Communications*.

For our second edition, published in 1996, our publisher asked us to include a chapter on ATM. Readers of our first edition told us likewise. That's because in 1996, ATM was being hyped as the be-all and end-all of networking technologies. We didn't quite see it that way, but in the interest of constant improvement, we agreed, albeit reluctantly, to include a chapter on ATM integration.

For this third edition, we have become Ethernet purists again. The ATM integration chapter has been cut. Instead, we have included two new chapters on Ethernet Layer 2 and 3 switching, as well as Gigabit Ethernet. Looking back, this microcosm trend within the context of our book also represents what has happened out there in the real world. For a few years, ATM was on the rise, but with the advent of Gigabit Ethernet and numerous other Ethernet innovations, ATM seems to be very much on the defensive. Just to put things in absolute numbers: The installed base of Ethernet nodes is now around 200 million, whereas the installed base of ATM is less than a million. That's a 200:1 ratio.

Migration Issues

Imagine that you are the network manager for a LAN that is running out of bandwidth and want to migrate or move your users, servers, and peripherals to a higher-speed network. Unfortunately, the number of competing high-speed technologies has created a lot of confusion in the marketplace. Some of these technologies are complementary, whereas others clearly compete with each other. This section outlines the most important issues you should consider before deciding on a particular high-speed technology. We assume in this chapter that you are running an Ethernet LAN, probably shared. That's a safe assumption, considering that well over 80% of all installations are Ethernet-based. (The balance includes Token Ring, FDDI, and ATM, with Token Ring accounting for the lion's share.)

When making your migration decision, you should consider several factors, including connection cost, performance, quality of service, ease of migration, your understanding of the technology, scalability, multivendor support, and standards. The following sections discuss these factors one at a time. After that, the discussion focuses on the capabilities of Ethernet and ATM with respect to these factors.

Connection Costs

The cost of a new LAN technology should be measured in terms of *connection costs*, which include the cost of the client or server NIC as well as the workgroup hub or switch port on the other end of the wire. You also need to factor into this calculation the cost of new network management software or any other additional hardware, such as backbone switches or routers, that need to be purchased.

Performance Considerations

Because you are reading a book about high-speed Ethernet options, performance is probably your number-one concern. Different vendors measure and promote their products' performance in various ways. Here are a few thoughts on what to look for in the small print:

- *Wire speed and data throughput rate*—You can measure performance in data networks in different ways. The *wire speed* is the maximum theoretical bit transmission rate for the network, inclusive of any overhead. For shared Ethernet, the wire speed is 10Mbps only after a node gets access to the network. The term *wire speed* is used because it implies the number of bits on the wire, regardless of the kind of bits, data, or overhead. The data throughput rate is always much lower, depending on such things as access method and protocol efficiency, whether it's a shared or switched connection, and whether it's half-duplex or full-duplex transmission.
- *Shared bandwidth*—Most of today's Ethernet networks share the total available bandwidth among many users, and the average data throughput rate available to each user is significantly less than the wire speed. One way to calculate average data throughput would be to take the wire speed number and divide it by the number of stations that share the network. A fully loaded Ethernet segment that has 200 users contending for the same 10Mbps channel delivers only 10Mbps/200 = 0.05Mbps of average throughput per station. This of course assumes that all stations generate the same amount of traffic all the time. The reality is that most of today's data traffic is very uneven or bursty at all times.
- *Dedicated bandwidth*—Some LAN technologies do not use a shared-media approach, and in such cases the available bandwidth can actually approach the wire speed. Switched Ethernet provides for a dedicated

connection without sharing the bandwidth with other users, so the available bandwidth is close to the wire speed.

- *Utilization*—Utilization is defined as how much of the theoretical available bandwidth is being used up. Shared Ethernet uses the carrier sense media access with collision detection (CSMA/CD) access method. CSMA/CD works extremely well with smaller numbers of users, large frames, and light traffic; with an increasing number of users and heavy traffic, however, collisions become a limitation. As a result, the utilization rate of a large shared Ethernet network has to be significantly less than 100%. Only in a point-to-point or Switched Ethernet environment does the actual throughput rate come close to the wire speed because a switched network in essence contains only two nodes and minimizes collisions.
- *Inherent frame or cell inefficiencies*—All networking transport methods require some kind of overhead, reducing efficiency to less than 100%. That's because such things as the address and error checking are sent over the wire but don't represent actual data. Ethernet efficiency varies widely, for example, depending on the frame size, from 40% for small frames to as high as 98% for maximum frame size. ATM efficiency is 90% at all times. Note that this efficiency refers to the data within the packet or frame and does not include efficiency calculations that result from different access methods.
- *Full-duplex (FDX) throughput*—*Full-duplex* means that data can be simultaneously transmitted and received, effectively doubling the nominal wire speed. Full-duplex is a relatively new phenomenon for data networks, having been made possible by twisted pair or fiber cabling, where one media pair can send while the second pair is receiving data. Coaxial cable, on the other hand, uses one single wire for either transmitting or receiving, and hence cannot accommodate full-duplex traffic. Note that full-duplex works only in a switched or point-to-point environment where both stations can support this capability.

  Many Ethernet vendors now advertise FDX as a feature and claim 20 or 200Mbps throughput. Switch-to-switch connections operating in full-duplex mode can yield this 100% improvement, but a full-duplex server or desktop connection will show hardly any improvement at all. The reason full-duplex shows marginal gains over half-duplex in these environments is that today's data flow is very lopsided. Applications such as Internet access are typically 90% downloads and 10% uploads, for example. For these applications, full-duplex merely adds a performance improvement of a few percent. Many networking protocols also cannot take full advantage of bidirectional simultaneous traffic flows.

The bottom line is that wire speed does not tell the entire performance story. The type of access method used and its efficiency are equally important considerations. A switched 100Mbps full-duplex Ethernet connection using 1.5KB frames, for example, will have a data throughput rate of 100Mbps × 2 × 98%, or 196Mbps. We cover this topic in more detail in Chapter 7, "Bandwidth: How Much Is Enough?"

Quality of Service

Quality of service (*QoS*) is a relatively new buzzword in the LAN industry, but has been in existence in the telecommunications industry for quite some time. QoS means that the recipient of the data gets the data *when* and *where* he or she needs it. Telephone companies measure QoS in terms of delay time or latency, signal-to-noise ratio, echo, wrong numbers dialed, and so on. QoS never used to be important for data networks, and with today's networked applications it is mostly not an issue. Take a very busy Ethernet network as an example, where a user is trying to send email over the network. The sender attempts to transmit his or her data stream right away but cannot because the LAN is busy. At some point, a time slot becomes available on the wire, and the transmission occurs. This process is transparent to the receiver, who doesn't care or know that the email message arrived a fraction of a second later.

QoS is becoming more important because today's data networks are increasingly utilized for time-critical applications, such as real-time voice and video transmissions. Latencies are acceptable for time-critical applications as long as they are relatively small and constant. A voice transmission with a constant delay of 0.1 second from sender to receiver will not be noticed, for example. Longer latencies could become an issue. Variable delays are a real problem for multimedia data transmissions. A changing latency is not acceptable because it will make audio transmissions sound like a tape recorder whose speed is varying, which is a sign of bad quality. Video transmissions are also particularly sensitive to variable latencies because jumpy picture quality, known as *jitter*, will result.

Another important part of QoS is congestion control. What happens if the network is very busy or overloaded? Is the data still guaranteed to arrive, will the sender be told to wait until bandwidth becomes available, or will the message be discarded? Guaranteed availability of a minimum amount of bandwidth is a key ingredient of QoS.

In general, dedicated connections provide the best QoS, allowing for guaranteed bandwidth at all times and constant latency. All shared-media technologies inherently exhibit variable delay times because the transmission channel is shared with other users.

Ease of Migration

Ease of migration is probably the most important consideration in determining the best technology for upgrade. If you are building a new network from scratch, this section doesn't apply to you. If you already have a network that you are growing or that needs a performance boost, however, this section is probably the most important point to consider.

You need to leverage your existing investment as much as possible. This investment includes the hardware itself, but also the tools, and your knowledge of the technology itself.

To ensure an easy migration, ask yourself these five questions:

- *Can my cabling plant support the higher-speed networking technology?*

  Before you decide to purchase high-speed networking hardware, make sure you understand the capabilities of your existing cabling plant. Your investment in cable, conduits, wiring closets, and patch panels can exceed the cost of the networking hardware itself! You need to make sure your new high-speed networking gear can run on your existing wiring wherever possible. If that's not the case, factor in potentially huge additional costs for upgrading your cabling plant to accommodate the new LAN standard.

- *How do I boost performance of my existing clients and servers with minimum cost and disruption to my users?*

  Most LAN managers will want to keep as much of their existing equipment as possible because it's working, proven, and already paid for. Replacing equipment is always disruptive and time-consuming and should be avoided wherever possible. Replacing network adapters should be avoided at all cost, for example; whereas the price of a new NIC alone may not seem that high, the cost of installation, configuration, and associated user disruption often exceeds the cost of the NIC itself. If redesigning your network and replacing a single strategic hub instead can increase the performance of the network, choose that way.

  You should avoid replacing equipment prematurely. Networking gear is part of a company's capital budget, meaning that the equipment needs to last for a period of five years before it is, in effect, paid for. Your accountants will tell you that replacing equipment before the five-year depreciation period is over can be prohibitively expensive because the equipment needs to be depreciated in one go for those purposes, which is a costly undertaking.

- *How do I connect new users or servers to the new network infrastructure?*

  When you add new users to your network, choose the best equipment available at the time. These new users will likely have faster machines, requiring a higher-speed network connection. Make sure the new technology scales to accommodate faster user and server connections.

- *How do I join a new section of my network with an older part?*

  New sections of your network will need to be seamlessly connected to your existing network. Think about how you are going to integrate the old and new sections. If you need to purchase additional hardware or software equipment, this will affect your overall cost. User disruption or server downtime should be minimized, if possible. Don't forget network management: Make sure your new networking gear can blend seamlessly into your existing network management map.

- *What happens to the replaced equipment? Can I use it somewhere else in my LAN?*

  Another point to consider when upgrading to a new high-speed LAN is what to do with the old equipment. Often, new high-speed equipment is added one step at a time, replacing at least some existing equipment. Your cost analysis needs to reflect whether the replaced equipment can be used somewhere else or whether it becomes obsolete.

Understanding the Technology

Your technical understanding of a new technology needs to be a key part of deciding which hardware technology to migrate. If you buy something completely new, you and your network/IT staff will need to learn about the new technology before you can deploy it. A steep learning curve will accompany a new technology, and years of familiarity will have to be relearned.

Maintaining a new technology presents its own set of challenges, too. You will need to buy and learn new network management and troubleshooting tricks and tools. Some people say that today's investment in IT is 20% hardware and software and 80% worker knowledge on how to use this hardware. Therefore, although it may sometimes seem to pay to buy something completely new, make sure the benefits outweigh the hidden costs.

Scalability

Your technology should grow as your network continues to grow. Can you upgrade one more step or are you buying a technology that has reached its limits? Does your network design lend itself to further upgrades? Scalability, for example, can mean that the technology can support a higher speed, that you can upgrade from a shared-media environment to a faster point-to-point or switched environment, or that you can add different hubs somewhere on your network to improve the overall throughput capability.

Multivendor Support

Make sure as many vendors as possible support whatever you buy. The following are some reasons why you should buy products manufactured by multiple vendors:

- *Lower prices*—Multivendor support means you have choices, and the more choices you have, the better. Multiple vendors are likely to compete more aggressively, ensuring lower prices for you. An early lead in terms of market share can turn into a formidable lead for a particular technology. Lower prices early on will drive up the volume of goods sold. This in turn will lead to economies of scale for a technology. Economies of scale occur in the high-tech manufacturing industry when the volume of goods consumed causes significant reductions in the associated cost. In this case, the cost of manufacturing hubs, switches, routers, and NICs will decrease significantly as the volume of product manufactured increases. That's why first-to-market products and standards often become so firmly entrenched in the market: They develop an early market share lead, prices decline, sales then skyrocket, and the lead widens even more. Sometimes nothing but a technological breakthrough can ever compete again.
- *Innovation*—Choice means that your particular supplier needs to work harder to earn your dollars. Innovation is just as important as lower prices because it ensures that future products will provide more features, higher performance, and other improvements that will benefit you in the long run. The more vendors that support a technology, the more innovation there will be, as competitors need to innovate to differentiate their products.
- *Availability of the necessary building blocks*—Today's network consists of many building blocks—chassis or stackable hubs and switches, bridges, routers, desktop, server and notebook NICs, MAUs, management software, and so on. No single vendor can supply all the building blocks, no matter what they tell you. Choosing technologies with broad vendor

support means that you can buy all the building blocks you need for your network, not just some or most of them.

Standards

Most LAN hardware sold adheres to some industry standard, yet the quality standards vary. A *good* standard means that hardware vendors can build equipment that is truly interoperable. *Interoperable* means you can buy two pieces of hardware from two different vendors, you can connect the two devices together, and they will work without fine-tuning or detailed configuration work. To accomplish that, the standards documents must be clear, unambiguous, and, above all, leave little open to interpretation. A standard, however, needs to be open enough to enable vendors to add their own features to differentiate their products. In addition, a good standard needs to be respected, or *authoritative*. This means that vendors building products will actually go to the trouble of obtaining and understanding the standards documents and specifications before building products. Overall, writing good standards takes experience and requires a good sense of balance.

An Ethernet Primer

We need to recap Ethernet for you briefly to compare Ethernet to the other high-speed LAN choices discussed in this chapter. Let's take a look at the inner workings of Ethernet, as well as its current market position:

- *Frame-based*—Ethernet is a frame-based LAN standard that was designed as a data transmission technology 25 years ago. With the exception of ATM, all LAN technologies use frames. Frames are per definition of variable length. In the case of Ethernet, the frame size varies from 64 to 1522 bytes.
- *Speed*—Ethernet was conceived as a 10Mbps shared media LAN technology, but has been modified over the past five years to run at 100Mbps and 1Gbps as well. A 10Gbps version seems inevitable.
- *Market share*—Ethernet is the most widely used LAN standard, with an installed base that now exceeds more than 80% of the installed base. New Ethernet shipments are approaching 90% market share, meaning that the dominance of Ethernet is still increasing. The market share growth of Ethernet is coming at the expense of such technologies as Token Ring,

FDDI, 100VG-AnyLAN, and older technologies, such as TCNS and ArcNET. To put things in perspective: Ethernet has an installed base of almost 200 million nodes. FDDI, ATM, and 100VG-AnyLAN each have an installed base of just several hundred thousand nodes, or about 0.1% of that of Ethernet. The only other LAN standard with a significant market share is Token Ring, which accounts for about 1 in 10 nodes.

Ethernet has reached a stage where it is becoming ubiquitous. Many new PCs now ship with a 10/100Mbps Ethernet built in. In 1997, 13 million PCs, or about 20%, included the familiar 10/100Mbps RJ-45 connector next to the serial and parallel port.

Figures 2.1 and 2.2 show 1997 data for both Ethernet switch port and NIC sales. Shared-media hubs still account for about 50% of all port shipments, although this number is declining rapidly as switch prices decline. Because ATM is only a switching technology, shared media hub sales are not shown.

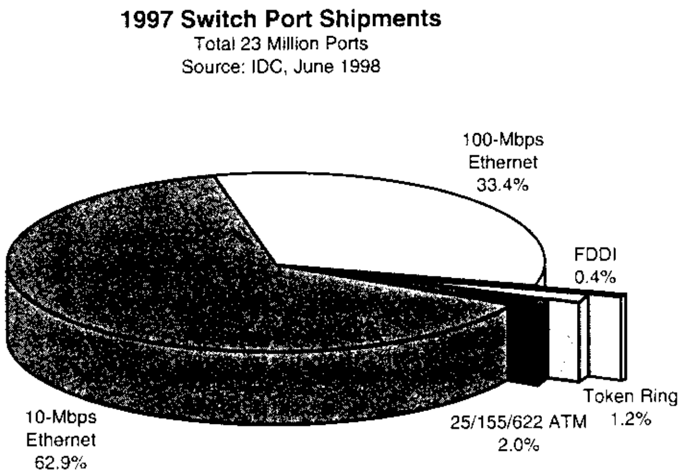

FIGURE 2.1 *Actual 1997 switch port shipments, according to International Data Corporation (IDC). Ethernet switch ports accounted for 96.3%, whereas ATM accounted for only 2%. Total port shipments were 23 million units, which excludes another 20 million or so shared-media port shipments.*

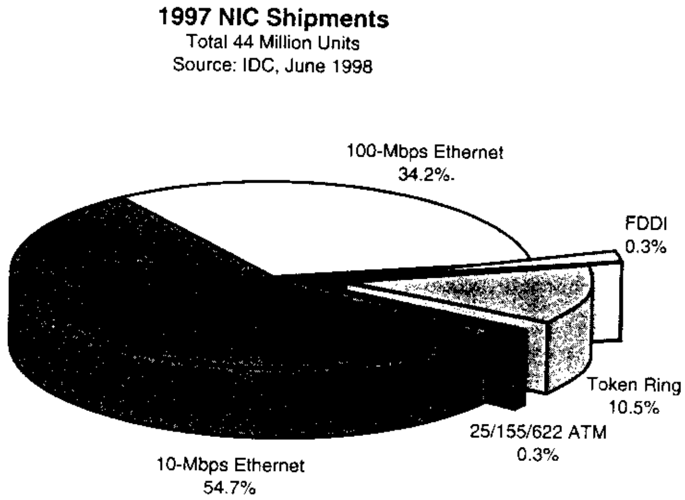

FIGURE 2.2 *According to IDC, 44 million NICs were sold in 1997. Note that ATM accounted for a mere 0.3% of all NIC shipments, whereas 10Mbps and 10Mbps/100Mbps Ethernet NICs made up almost 90%.*

- *Standard*—Since 1980, the IEEE 802.3 committee has been setting Ethernet standards. In the past few years in particular, the IEEE has been very active with improving Ethernet. As a result, six major addenda have been published in the past five years.
- *Pricing*—Today, a single 10Mbps shared Ethernet connection costs as little as $50 (for a NIC and a hub). Compare that to the cost of the original 3Com EtherLink NIC 15 years ago: It cost $1,000 back then! A move to 100Mbps or a switched 10Mbps connection costs about twice as much. Gigabit Ethernet connection cost is still relatively high, around $2,000 per node. In terms of cost per Mbps, Ethernet is by far the cheapest LAN standard out there.
- *Vendor support*—Every local area networking vendor in the business supports Ethernet. That means lower prices, choices, and the availability of all the necessary building blocks.
- *Shared, Switched, and Full-Duplex Ethernet*—Ethernet was originally designed as a shared-media technology in which multiple nodes share the same bandwidth. In this configuration, the CSMA/CD MAC is operative and controls all data transmissions. Over the past few years, Ethernet switching has emerged whereby only two nodes reside on one single Ethernet. All three Ethernet speed versions can be switched—that is,

operated in a point-to-point method. In addition, multiple switched Ethernet connections can now be connected in parallel to increase the overall throughput (called Link Aggregation). Ethernet is relatively simple and, because of its huge installed base, is well understood and documented.

- *Works with IP*—All data transmissions use the OSI Reference Model, which is covered in more detail in Chapter 3, "Ethernet, Fast Ethernet, and Gigabit Ethernet Standards." Ethernet operates at Layer 2 and works extremely well with all Layer 3 software, including the Internet Protocol (IP).

Ethernet does have some disadvantages as well. Ethernet has been designed as a data transmission technology, as opposed to real-time voice technology. This means that Ethernet does not have any built-in quality-of- or class-of-service guarantees. Recent efforts, such as the IEEE 802.1p and IETF RSVP standards, will improve Ethernet's capabilities in this respect, but true QoS will remain elusive.

Note

In this chapter only, we refer to the different 10Mbps, 100Mbps, 1000Mbps Shared and Switched Ethernet options collectively as Ethernet.

Pricing in this chapter refers to late 1998 data.

Revolutionary Versus Evolutionary Technologies

Today, the computer and communications industry experiences essentially two types of technological innovations. *Evolutionary* innovations build on the installed base and provide some migration path. The innovation typically occurs in smaller steps. The PC industry as we know it is built on evolutionary progress. Windows was designed to run on top of DOS, not in its place, so that people could still run their existing DOS applications. Likewise, Windows NT/95/98 all run the huge collection of existing DOS and Windows applications. All Intel microprocessors are backward compatible (the "X86 family"), yet each new generation adds some additional features.

Revolutionary technology, on the other hand, implies a really radical breakthrough of some kind. This typically requires the forklift upgrade, which means scrapping the installed base of product. With revolutionary technologies, the benefit of moving to the technology must clearly outweigh the costs of making obsolete the existing technology. A fairly recent example of a revolutionary technology is the compact disc, which made the vinyl record obsolete.

As our society has become more technologically oriented, revolutionary technologies have become few and far between. That's because the installed base has become so large that the cost of obsolescence for a large installed base has become insurmountable.

Note

In the networking industry, the term forklift upgrade *often describes the concept of wholesale replacement of one technology by another. The term is derived from the concept of ripping or lifting out all the existing equipment, carting it away, and then bringing in truckloads of new equipment.*

An example of a failed revolutionary product is the video laser disc. The laser disc used a completely new size disc, which meant purchasing a new player as well as new media. There was no backward-compatibility with anything. Whereas laser discs offered superb video and audio quality, that wasn't enough to cause everyone to throw away their VHS players and cassette collections and buy an expensive laser disc player instead. The inventors of DVD players chose a different path: DVD offers laser disc-type quality, yet is backward compatible with today's audio and data CD formats. The success of DVD as the replacement of today's audio CD players, laser discs, VHS players, and CD-ROMs is almost guaranteed.

The Most Common High-Speed Technologies Compared

This section compares yesterday's and today's high-speed technologies and discusses their strengths and weaknesses and current market positions. We use the previously discussed migration issues to guide us through this analysis. We look at FDDI, 100VG-AnyLAN, TCNS, Fibre Channel, HIPPI, and Iso-Ethernet. Many of these technologies have tried to compete head-on with Ethernet and have fallen by the wayside. Other technologies focused on specialized niche areas, however, and are still in production but with negligible overall market share.

We have dedicated an entire separate section to ATM, which is the only viable alternative to Ethernet at this point.

FDDI

Mainframe companies, such as Sperry, Burroughs, and Control Data Corporation, first developed Fiber Distributed Data Interface (FDDI) more than 10 years ago. Standardized by the ANSI X3T9.5 committee in 1990, FDDI incorporates many features of IBM's Token Ring technology, such as the Token Ring frame format and shared-media ring architecture. FDDI also has sophisticated management, control, and reliability features not found in classic Ethernet or

Token Ring. An optional second counter-rotating network ring improves overall reliability, for example. FDDI supports cable lengths of up to 2km for multimode fiber. FDDI products first appeared in 1988. For many years, it was the only viable high-speed backbone technology available.

FDDI was created at a time when scientists and engineers were still predicting that fiber cabling would reach every office and home by the year 2000. One of the inventors of FDDI claims that he chose a 100Mbps data rate to make FDDI so fast that no copper-based LAN technology would ever be able to match its speed.

Note

The choice of making FDDI run at 100Mbps so that no copper-based LAN would ever be able to match its speed seems pretty ironic these days. A few years after FDDI was introduced, the UTP-based CDDI technology was introduced. More ironic still is the fact that we are now talking about Gigabit Ethernet over UTP.

Historically, ANSI was always the standards authority for high-speed data transmission technologies, high-speed being 100Mbps. The IEEE, on the other hand, focused on the slower LAN standards, such as Ethernet and Token Ring. Of course, this has changed with Fast and Gigabit Ethernet.

Over the past 10 years, twisted-pair wiring and structured cabling systems rapidly replaced coaxial cable as the most popular LAN media, preventing fiber cabling from becoming widely accepted. In addition, engineers figured out a way to run FDDI over twisted-pair wiring. This technology is called *TP-PMD*, short for *twisted-pair physical media dependent*. The ANSI TP-PMD standard uses a transmission scheme called *MLT-3*, short for *Multilevel Transmission 3*. The standard requires two pairs of Category 5 data-grade wire and supports a maximum distance of 100 meters. Crescendo Communications, now part of Cisco, pioneered this technology and called it *CDDI*, short for *Copper Distributed Data Interface*.

Strengths of FDDI/CDDI are as follows:

- FDDI delivers 100Mbps throughput capability with little overhead. It is a shared-media technology; unlike Ethernet, FDDI does not use a collision-based access method, so there is no performance degradation at high usage rates. This allows utilization rates of 80 to 90%, translating to close to 80 to 90Mbps data throughput.
- The FDDI standard was in the making for almost 10 years. Products have been shipping for more than about 10 years as well. The technology is proven, mature, and well understood.

- At its peak, FDDI had broad multivendor support. All the major networking suppliers, including Cisco, Bay Networks, 3Com, and Cabletron, offer FDDI products. FDDI chipsets from Motorola, AMD, and National Semiconductor provide a source of semiconductor building blocks to the networking industry.
- The second counter rotating network ring provides an element of redundancy to FDDI networks. Should one ring go down, dual attached stations (DAS) can continue communicating over the second ring.

The weaknesses of FDDI/CDDI are the following:

- FDDI is too expensive to compete effectively with Ethernet or even Fast Ethernet. To this day, FDDI and CDDI prices are many times higher than Fast Ethernet's prices.
- Upgrading from Ethernet to FDDI is difficult and expensive because FDDI uses a different frame format—meaning that complex and expensive routers are required.
- Many people view FDDI as a high-performance technology that is expensive and only suitable for backbones. This perception always makes customers leery of buying FDDI for anything other than a backbone technology.

FDDI-II, an improved version of the proven FDDI standard, was supposed to offer improved support for multimedia data transmissions. It offered isochronous data transmission, which means predictable and guaranteed access time. In addition, the FDDI-II proposal included a prioritization scheme to ensure low latency, which is important for video transmission in particular. The ANSI X3 committee set some FDDI-II standards from 1994 to 1997. The vendor community decided to focus on Fast Ethernet, however, and no FDDI II-compatible products ever made it to market (as far as we know).

Note

What exactly is isochronous data? Isochronous data technologies can accept data at specified and guaranteed intervals. Isochronous data transmission provides what ATM calls a constant bit rate (CBR) service. This capability is important for voice and video transmission, where no interruptions or delays in the traffic flow can be tolerated without degrading the quality of service.

Integrated Services Digital Network (ISDN), Iso-Ethernet, FDDI-II, Fibre Channel, and ATM all contain isochronous data-carrying capabilities. Ethernet, on the other hand, does not. Ethernet provides for asynchronous data transfer, which means the data is transferred in a best-effort manner.

When working on a Fast Ethernet standard, the IEEE Fast Ethernet engineers chose to use the existing and proven FDDI and CDDI physical layer technologies. The Fast Ethernet equivalents became the 100BASE-FX and 100BASE-TX standards.

Until Fast Ethernet and ATM appeared on the market in the mid-1990s, FDDI was the only high-speed LAN standard to deliver true 100Mbps performance. As a result, many users, resellers, and manufacturers used FDDI as a backbone technology. FDDI sales grew steadily from 1988 until about 1996, at which point Fast Ethernet and ATM became the de facto standards for new backbone installations. Many manufacturers continue to support FDDI, but the technology is clearly in a state of decline, with no new products being designed. Current sales support existing installations only.

TCNS and 100VG-AnyLAN

TCNS and 100VG-AnyLAN were both market failures, despite the fact that they both delivered true 100Mbps performance at reasonable prices. We mention TCNS and 100VG here because they provide two examples that failed the upgrade test.

TCNS was a proprietary 100Mbps solution developed by Thomas-Conrad in Austin, Texas in the early 1990s. (Compaq Computer purchased Thomas-Conrad in 1995). TCNS stands for Thomas-Conrad Networking System and is based on the ANSI ArcNET 878.1 specification. The technology uses a token-passing bus access method and operates in a shared-media mode. It supports fiber cable of 900 meters, coaxial cable of 150 meters, and STP or UTP of 100 meters. Thomas-Conrad started shipping TCNS NICs and hubs in 1990, and in 1993 more than 20,000 TCNS NICs were sold, representing a significant share of the high-speed NIC market at the time. No other vendor ever built TCNS-compatible equipment.

100VG-AnyLAN is a 100Mbps shared-media technology jointly developed by Hewlett-Packard and AT&T's Microelectronics (now Lucent) from 1991 to 1994. 100VG incorporates an access method protocol called *demand priority*, which allows time-critical applications to transmit ahead of other noncritical packets. The IEEE standardized 100VG-AnyLAN is the new 802.12 standard. But Hewlett-Packard unsuccessfully positioned 100VG as its own version of Fast Ethernet. 100VG products first appeared in 1994.

TCNS and 100VG failed for two reasons:

- *Multivendor support is critical to the survival of any technology*—TCNS was a proprietary technology from the beginning, and Thomas-Conrad was the only vendor selling TCNS equipment. Whereas 100VG is nominally an industry standard, all major industry players backed the competing 802.3u Fast Ethernet proposal, leaving Hewlett-Packard as the only major vendor promoting 100VG. This effectively made both technologies proprietary, which had numerous drawbacks. Customers buying single-vendor products put themselves at significant risk because no single vendor can or wants to provide all possible building blocks of a complete network.

- *Customers want evolutionary technologies; they don't want to replace everything they own today*—TCNS and 100VG also failed in this area because there was no easy upgrade path from shared Ethernet to either 100Mbps technology. TCNS deployment requires that both NICs and hubs be replaced altogether. This was expensive and time consuming. Whereas Hewlett-Packard is still selling 10/100VG NICs, the lack of 10/100Mbps hubs or switches made seamless migration very difficult.

Technically, nothing is wrong with either TCNS or 100VG. But faced with stiff and growing competition from both Fast and Switched Ethernet offered by literally hundreds of networking vendors both small and large, both TCNS and 100VG were doomed to fail. TCNS is no longer in production. 100VG is still in production to serve a niche market consisting of the Hewlett-Packard installed base.

Note

In some ways, 100VG was designed to provide the best of both worlds. It was supposed to provide backward compatibility with both Ethernet and Token Ring. 100VG can use either Ethernet or Token Ring frame formats, supposedly allowing for easy migration from either technology. In addition, Hewlett-Packard engineers changed the Ethernet MAC to improve the utilization and multimedia capabilities. As recently as two years ago, people still thought that 100VG might become a viable alternative to 100BASE-T. Why then did 100VG fail? There were several reasons:

- *Hewlett-Packard lost the first round when the IEEE decided that 100VG was not Ethernet and that a newly formed IEEE 802.12 committee would be put in charge of standardizing 100VG. This made it clear to the industry and customers that 100VG was not Ethernet. The demand priority MAC from Hewlett-Packard was technically more elegant than the Ethernet CSMA/CD MAC. Over the years, however, dozens of attempts have been made to improve the Ethernet MAC algorithm. People lost interest in 100VG because it was called something other than Ethernet.*
- *Hewlett-Packard ignored the fact that one medium-size networking vendor cannot change the world on its own. While Hewlett-Packard was promoting 100VG, the networking giants (Cisco, 3Com, and Bay, as well as dozens of other medium and small vendors) were backing Fast Ethernet.*

continues

- *The standard Ethernet MAC is already on its way out. Ethernet switching effectively replaces the CSMA/CD algorithm with simultaneous transmission and reception of frames.*
- *The Ethernet camp reused many different building blocks, including the proven 10Mbps MAC and the FDDI physical layer. This meant that semiconductor and passive components were readily available and cheap. Hewlett-Packard chose to reinvent practically everything from the ground up. 100VG, on the other hand, had little support from the semiconductor community, and therefore VG was more expensive to produce.*

In early 1996, Hewlett-Packard's networking division announced plans to support 100BASE-T as well. The company went to great lengths to explain that these developments were not diminishing Hewlett-Packard's commitment to 100VG, but most savvy analysts saw that this marked the beginning of the end for 100VG. The 100VG debacle cost Hewlett-Packard millions of dollars in research, development, and marketing. Worse still, it isolated Hewlett-Packard from the mainstream networking business for a number of years. At one time, Hewlett-Packard was a major networking player, including the number-one vendor of stackable hubs. These days, Hewlett-Packard is a mere shadow of its former self in the networking industry.

Iso-Ethernet

Conventional wisdom has it that Ethernet, today's standard for data communications, is unsuitable for time-critical multimedia applications. At the same time, ISDN was supposed to become a pervasive standard for high-speed digital dialup WAN connectivity. In the early 1990s, National Semiconductor developed a technology called *isochronous Ethernet* or *Iso-Ethernet*, which combines the best of Ethernet and ISDN. This technology, which was adopted as the IEEE 802.9a standard in 1995, features a regular 10Mbps Ethernet channel for non-critical data transmission, and additional support for up to 96 ISDN B channels (for data) and one ISDN D channel (for signaling and control). The total capacity of Iso-Ethernet is 97 × 64Kbps + 10Mbps, or 16.16Mbps total.

The concept was powerful, yet elegant and simple. An internal (data) LAN utilized Iso-Ethernet's Ethernet channel to transfer bursty data. Time-critical data was transferred by making an ISDN call within a LAN or to the outside world via an external BRI connection located in the hub. The hub contains more than one BRI ISDN connection to the outside world to allow multiple LAN users to make external ISDN calls simultaneously.

Iso-Ethernet essentially suffered from the same fate that beset TCNS and 100VG:

- Iso-Ethernet required both new hubs and NICs—the notorious forklift upgrade.
- Iso-Ethernet did deliver on its promise of isochronous data transmission capability for time-critical voice and video transmission. Unfortunately, the technology provides no substantive increase in overall data throughput, which is what most LAN managers were seeking. Iso-Ethernet

provides for only 96 64Kbps data channels, a total of 6.16Mbps of isochronous data capacity for an entire network. Therefore, the total data capacity was 16.16Mbps. While this technology was being proposed, the Fast Ethernet people were promoting 100Mbps capacity, at price points similar to Iso-Ethernet.

- ISDN never took off as forecasted.

The rest of the industry chose to ignore Iso-Ethernet, and backed such alternatives as Fast Ethernet or ATM instead.

Fibre Channel and HIPPI

Fibre Channel and HIPPI were developed in the late 1980s and are computer-to-computer or high-speed peripheral connection technologies. Both are used primarily to link different mainframes, supercomputers, and their peripherals together. Fibre Channel, as the name implies, is based primarily on fiber-optic media, but it supports shielded twisted-pair (STP) as well. It runs at speeds ranging from 100Mbps to 800Mbps. *HIPPI* stands for high-performance parallel interface and runs on a 50-pair parallel cable at speeds of either 800Mbps or 1.6Gbps. HIPPI uses a parallel transmission method (a bus of either 32 or 64 bits), but Fibre Channel uses serial data transmission. (Just to confuse things, a serial version of HIPPI was developed later on, which actually makes the two technologies competitors.) Both Fibre Channel and HIPPI are ANSI X3T11 standards.

Both technologies are flexible enough to be used as high-speed LANs as well. Fibre Channel in particular incorporates features that could have made it a very successful high-speed LAN. Neither Fibre Channel nor HIPPI ever achieved any significant volumes, however, and they never really expanded beyond their mainframe, supercomputer, or high-speed peripheral connectivity niche areas.

Fibre Channel uses an 8B/10B encoding scheme that increases the baud rate (bits on the wire) to 1.062Gbaud. Therefore Fibre Channel and HIPPI were actually the first LAN technologies to pioneer Gigabit data rates.

Instead of developing something from scratch, IEEE engineers chose to use the proven Fibre Channel physical layer and encoding standard for the Gigabit Ethernet standard. Therefore, although Fibre Channel was never successful as a LAN technology, parts of it live on under the Gigabit Ethernet name. Table 2.1 compares Fibre Channel and HIPPI.

Both Fibre Channel and HIPPI are still being supported by vendors, such as Sun, Hewlett-Packard, and IBM, for their large computing platforms.

TABLE 2.1 A COMPARISON OF FIBRE CHANNEL AND HIPPI

Variable	Fibre Channel	HIPPI
Data rate	100, 200, 400 or 800Mbps	800 or 1600Mbps
Transmission method	Serial	32- or 64-bit parallel bus
Frame type	0-byte to 2048-byte payload	1K or 2K packets[1]
Standard	ANSI X3T11	ANSI X3T11
Cable options	25m of 2-pair STP (Type 1)	25m of 50-pair STP (1 or 2 cables)
	Two strands of 2km multimode fiber (MMF)	Four strands of 1km MMF[2]
	Two strands of 10km singlemode fiber (SMF)	Four strands of 10km SMF[2]

1 *HIPPI transmits 256 words of 32 or 64 bits each (1KB or 2KB packets). It transmitted in a half-duplex way with flow control.*

2 *Serial HIPPI.*

ATM

ATM has emerged as the only serious and viable long-term alternative to Ethernet. This section discusses ATM as a technology, as well as the pros and cons of migrating to ATM or staying with Ethernet as a LAN solution. This section does not pretend to be a complete and thorough examination of ATM: We cover the basics only. We also focus on the local area networking aspects of ATM rather than the much-touted wide area networking properties of the technology.

A Brief History of ATM

In 1994, when we wrote this book's first edition, the high-speed networking market was filled with numerous contenders. Slowly, one-by-one, they all disappeared. By 1996, when we published this book's second edition, it was becoming increasingly clear that the high-speed LAN race was becoming a two-horse event: Ethernet and ATM. At the time, our publisher convinced us that we needed to include a chapter on ATM integration. We did, although reluctantly. Two years ago, ATM offered some clear benefits over Ethernet. Gigabit Ethernet and 802.1p QoS in particular have eroded this advantage. Therefore, we decided to cut the ATM chapter for this third edition. Overall, we think that ATM is clearly on the defensive, and that there is no place for ATM in the local area network.