1

First, the Basics

1.1 LANs, WANs, What Is the Difference?

LAN means Local Area Network, or basically "everything inside." WAN means Wide Area Network, or basically "everything outside" (the telephone company). There have been other manifestations of the name as well. Campus Area Network (CAN) is sometimes used, and the meaning is self-evident. Metropolitan Area Network or MAN is also sometimes used, sometimes in the context of a citywide LAN. (Wait, I thought LANs were inside!) Some of the more silly network acronyms include DAN (Desk Area Network) and PAN (Personal Area Network).

Of course there is also POTS, which is not a network at all, but instead stands for "Plain Ordinary Telephone Service" and really is a term used in the industry. Therefore one can argue that we use pots and pans. This can get just about as silly as you care to make it. See Table 1.1. All kidding aside, however, we mostly use two terms: LAN and WAN.

The other point that you have probably already deduced is that the acronyms in our business of networks will just about drive a new user crazy. Do not worry, we will try to acclimate you to the new language in the course of the book through repetition and immersion. By the time you are through the first few chapters, we will ask you if you can run B8ZS and ESF on an xDSL T1 circuit with CCIS7, and you know what? You may not know for sure but you will be able to offer an educated guess! If you get really stuck in the meantime, there is a glossary and list of common acronyms in the back of this book for your review. Use them as needed as you read.

1.2 So This Is Basic?

For a chapter that is supposedly dealing with the basics, where do we get off introducing a figure like the one on page 3? If you are new to telecommunications it probably does not look that basic to you! And you are probably right.

Imagine, however, that you really wanted to learn Spanish. The best way to do it quickly would be to live with a family in Mexico City for three months, and completely immerse yourself in the language. That is kind of what we are trying to do here. By repeating the acronyms, and demonstrating a particular component's role in the big picture, we hope to help you remember these things by framing them in a context.

LAN	Local Area Network
WAN	Wide Area Network
MAN	Metropolitan Area Network
CAN	Campus Area Network
FAN	Foreign Area Network
DAN	Desk Area Network
PAN	Personal Area Network
POTS	Plain Old Telephone Service
POTS and PANS	Kitchen Implements

Table 1.1: Terminology. All kidding aside, all of these terms have been used at one time or another. Realistically speaking, we are most concerned with LAN and WAN.

In the same way that someone immersing themselves in a language learns quickly, having a framework in which to reference the words, rather than just a book from which to memorize them, will help you to learn quickly. That is also why some of the more important information repeats, albeit under a different context each time. It is the best way I have found to bring an overwhelmed technical service manager up to speed quickly. So with this in mind, let's evaluate the major components in this framework, in the context of a complex network diagram.

1.2.1 What Do Mainframes Do? ① [1]

Mainframes do a couple of things really well. They are the cheapest way to deploy lots of computer (CPU) cycles to a large number of users. Mainframes are also good for applications that encompass repetitious tasks which are easily automated. Examples include airline ticket sales, rental car sales, and just about any kind of business where customers call in or walk in, and where an agent fills in fields in a form to sell a service. Insurance policies are another good example. Everyone fills in the same information: name, date of birth, address, and so forth.

If your business is characterized by any of these processes, mainframes are probably the ticket. We will discuss more about mainframes, and their role in future networks later in this book.

1.2.2 What Is a Front-End Processor? ②

You can imagine that a mainframe supporting hundreds of telemarketers, for example, is probably a pretty busy box. Therefore, mainframe manufacturers invented a "traffic cop" to

[1] In Chapter 1, all circled numbers appearing beside section headings refer to the circled numbers in Figure 1.1.

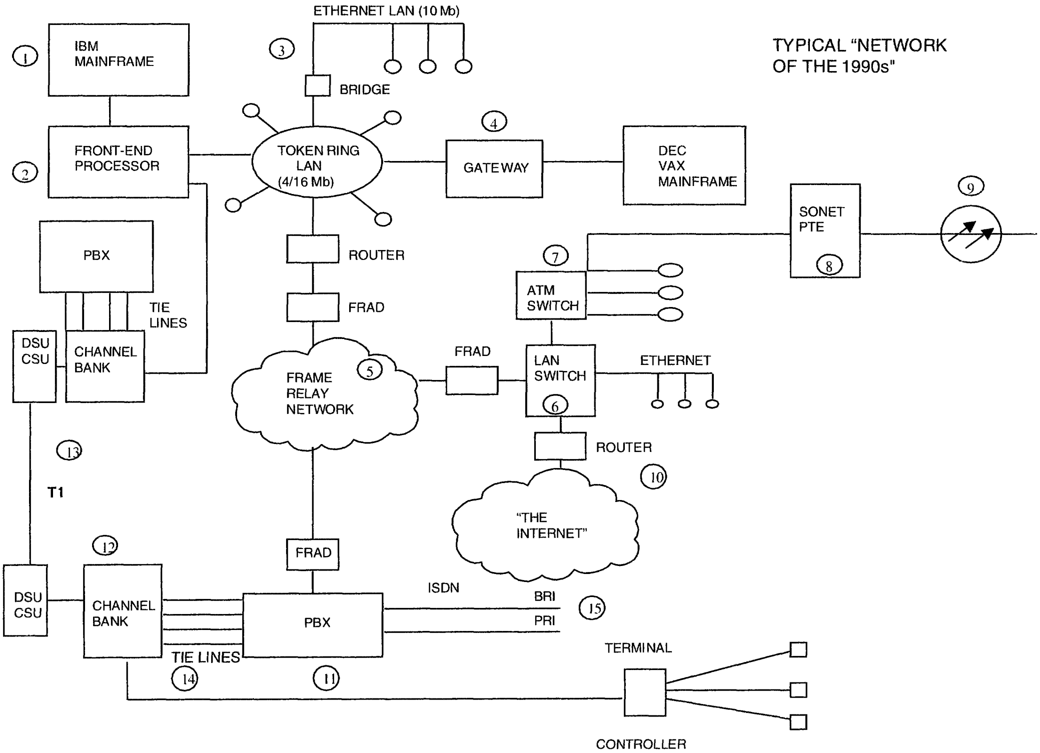

Figure 1.1: This is "basic?!?"

direct all the messages from the users into and out of the mainframe to avoid bogging the mainframe down. The front-end processor is therefore just what the name implies, it is a gatekeeper to the mainframe.

Now with this in mind, you may already know that there are other boxes that perform a similar function in today's networks. Routers and LAN switches come to mind. And indeed, many of these devices are overtaking and usurping the front end due to their flexibility and ease of integration into distributed LAN networks. For this reason, the front end might just be an endangered species. When you look into your crystal ball, ten years from now you will still see a mainframe. They are just too good at what they do and just too inexpensive in a large organization to dispense with. Front ends, however, do not seem to be on the radar.

By the way, a common name for both of these boxes is the *legacy* network. Legacy, as the name implies, means the "holdover" technology from a bygone era, and is loosely interpreted to mean the mainframe environment.

So what is replacing the "legacy" network? It is the LAN network.

1.2.3 What Are LANS? ③

LANs are characterized by the following:

- They are *local*, meaning within a building or at most a campus environment.
- You, not the telephone company, own them. That means the bits are "free" once the initial investment has been made. It also means that in case of trouble you are on your own.
- LANs usually run on copper cable media, like twisted pair or coaxial cable. Older technologies used coaxial cable, while most new ones use unshielded twisted pair (UTP) cable rated for the particular speed of the LAN.

The two most prevalent LAN technologies out there today are Ethernet and Token Ring. Ethernet was designed to run at 10 Mbps, but has been all but replaced in large organizations by Fast Ethernet, which operates at 100 Mbps. This has chiefly been the result of a major drop in the price of Fast Ethernet cards and greater bandwidth (capacity) demands of the user.

Token Ring was originally embraced by IBM. It too is being replaced in most organizations by Fast Ethernet. Token Ring operates at 4 or 16 Mbps. It has been commonly held that if the decision on LAN technology was made while the mainframe bigots were still in charge, the company went with Token Ring. In later years, the swing was to Ethernet.

Both of these technologies will be discussed in greater detail in Chapter 2.

1.2.4 How Do We Connect LANS? ④

In much the same way that a front-end processor is the gatekeeper to the legacy environment, a bridge, router, or gateway is the gatekeeper to the LAN environment. Bridges basically isolate and filter different LANS from one another. This way if one LAN segment gets really busy, it does not slow down the whole organization. On the other hand, if the two segments need to communicate with one another, the bridge will forward a message to the other segment. That is why people commonly describe the function of a bridge as two things: filter and forward.

Routers and gateways act similarly but have a few more brains that bridges. For example, bridges are not particularly good at dealing with ambiguity. If more than one path exists to a user (common since networks are often engineered to provide multiple paths for fault protection), a bridge is not particularly well suited for selecting the best one. A router on the other hand is much more capable in this area.

1.2.5 What Is Frame Relay? ⑤

Frame Relay is one of the hottest technologies available today. In Figure 1.1 we have it illustrated as a "cloud." I hate clouds. My first question to a vendor is what kind of network is inside the cloud, and in later chapters we will look into exactly that.

So why is Frame Relay so hot? Imagine for a moment that I am a competitive *postal* provider and offer you a deal. I will deliver a letter anywhere in the U.S. for five cents. Sound interesting? Good, I will go on.

In order to get this rate, you have to follow some rules. First, you will have to commit to some minimum volume of letters you will let me deliver. Second, you will need to use *my* envelopes. Don't worry, I'll provide you with the envelopes. Third, you can put no more than one sheet of paper in each envelope. If you want to send a three-page letter, you must use three envelopes. Even so, that is still only 15 cents, less than half of what the U.S. Post Office charges. For 100-page documents from the legal department, you can always drop them in Priority Mail using your old postal carrier.

In keeping with this analogy, consider the following. If you mail five letters to Dallas on the way to work in the morning, then five more at lunchtime, then five more at 5:00 p.m. on the way home:

A. Will all 15 arrive in Dallas at the same time?

B. Will all 15 get to Dallas in the same way?

C. Will all 15 even get to Dallas on the same *day*?

D. Is there any guarantee that all 15 will even get to Dallas at all?

The answers to A, B, C, and D above are all "no." And the same holds true for Frame Relay.

First off, 15 data packets sent on a Frame Relay network to Dallas will not necessarily get to Dallas at the same time. That is because not all 15 will take the same path through the network. The carrier decides on the path based on network conditions at the instant you send the data. Just like your letters may take a train, aircraft, or truck depending on the time you send them, your data similarly will take different paths.

There is also no guarantee that the data will get there at all. If the Frame Relay network gets too busy, a few packets might be thrown overboard. But when you think about it, even with a dedicated circuit there is no guarantee that your data will get to the far end. That is why there are higher level protocols like TCP/IP or SDLC which check for lost data or packets delivered in the wrong order. That means Frame Relay works great for data, and it is cheap. But what about voice or video?

Voice and video require a continuous bit stream with the timing relationship of the bits maintained. Another word for this is *isochronous.*

For example, if I am watching the Dallas Cowboys on TV here in Dallas, and Emmitt Smith rushes for a touchdown, I get to see it about half a second earlier than someone in New York who would be receiving their view of the game via network television satellite. We see the same play however, and even though the New York viewer sees the play delayed by half a second (because of the propagation time to the satellite) he or she sees the same thing I do.

Imagine however that the video image was sent over a Frame Relay network, where there were no guarantees that the bits would arrive in order. I could just hear John Madden now:

> He's at the 40!
> The 30!
> The 35!
> The 20!
> The 5!
> The 10!

Okay, so I made a lame joke. However, this does illustrate the point and why a user contemplating voice or video on a Frame Relay network, or other *non*isochronous means of transport could have problems. We will discuss this technology in greater detail later in the book.

1.2.6 What Do LAN Switches Do? ⑥

For lack of better words, LAN switches are souped up bridges. Not only do they filter and forward, but they have advanced network management capabilities which allow the technologist to keep an eye on the network through proactive maintenance. Moreover, LAN switches are *fast*, and for that reason they often replace routers. Lastly, today's LAN switches allow myriad possibilities for mixing and matching different equipment, components, and environments for truly seamless solutions. These are discussed later in the book.

1.2.7 What Is ATM? ⑦

What is a user to do if he desires the economies offered by Frame Relay, combined with isochronous capabilities for voice and video? What if a user wants to eliminate the distinction entirely between his LAN and WAN environments? What if a user needs to move multiple gigabyte diagrams, such as computer-aided designs or X-ray images? What if users need a single, integrated network access for voice, data, video, image, multimedia, and so on, with virtually limitless speed and flexibility? These describe the characteristics of the ATM user. If comedian Tim Allen was allowed to "rewire" Frame Relay, he might come up with ATM.

Why do LANs and WANs have different rules? When ATM becomes prevalent, the signal will stay the same from a workstation in New York to a workstation in Dallas and the distinction between LANs and WANs will be virtually gone. We will describe how this will happen in later chapters.

1.2.8 What Is SONET? ⑧

SONET will replace T1 and T3 as the backbone transmission media in the U.S.

Since T1 and T3 have been around for almost 40 years, it sounds like a safe bet that by the time there is a replacement for SONET, we will all be wandering around shopping malls looking for the perfect frozen yogurt. And how often do we get a planning window that long? That is the thing about transport technologies. They have a nice long life.

SONET stands for Synchronous Optical NETwork, and was designed as a transport technology for fiber optics. SONET is far smarter than T1. T1 was originally designed as a means of combining 24 channels between Bell central offices onto four wires instead of 96. T1 was designed for *voice*.

T1 is woefully dumb. In fact, if you put a T1 and a cinderblock into an intelligence contest, it would be a tie.

SONET on the other hand is *smart*. SONET is clean and error free. SONET has the capability to repair itself if the fiber is cut through a feature called fault protection switching. SONET has nonintrusive maintenance features—meaning you do not always have to take it down to test if it has trouble. And SONET is designed for glass—fiber optics—the mainline telecommunications medium in place today.

Your users will be willing to pay for the features in SONET, and the technology is more available today than ever. We will discuss the features of this exciting and long-lived technology, as well as how to get it into your organization in later chapters.

1.2.9 What Is Fiber? (9)

Fiber optics is the medium used for virtually all high capacity communications services. By high capacity, we mean essentially anything above the T3 rate of 44.736 Mbps. Fiber is also used in the LAN environment as a high speed connection between LAN switches and other components. Fiber can also be used to connect mainframes together at the channel level in the legacy environment. There are two kinds of fiber:

1. Single Mode. This is what the phone companies use.

2. Multimode. This is probably what you have in your building.

We will discuss their differences in later chapters.

1.2.10 What Is the Internet? (10)

Imagine you are the U.S. Government in the 1960s and 1970s, planning for the prospect of incoming Soviet ICBMs. Intuitively, the war planners of the time knew that 20 minutes after a launch, virtually every telecommunications hub in the U.S. would probably be in the upper atmosphere. This presented an obvious problem.

Since any potential opponent would probably hit telecom hubs first (thereby taking out the eyes and ears of their opponent) some survivability had to be planned for. The first thing the government did was build some big telephone centers in relatively outlying locations. For example, consider Ennis, Texas, population 14,000, about 30 miles south of Dallas. A nice little town, good Czech sausage, and one of the biggest AT&T offices I've laid eyes on. I'm not sure if it was part of the 1960's effort to harden the public network, but it sure fits the bill: uncharacteristically big for the town, no windows, and 30 miles from ground zero.

Even with precautions like these however, it would be difficult for data communications to take place using the traditional "master-slave, host-terminal" environment. It would be logical to assume the host would be in a big city. That means that even with the path in place, data communications would still be out of the question.

Just what if, for example, each node in the network was just a little bit smart? Distributed intelligence with nobody in charge. (Kind of like my office.) The idea behind the Internet was that data could be put in packets, with a header address to where it was destined. This packet is called a *datagram*.

After a first strike, some parts of the network might still be intact. Using a distributed intelligence topology like this one, messages might eventually get through. This was the beginning of the Internet we know today.

Routers are just a little bit smart. They are better at dealing with ambiguity than bridges. For example, they can select an "alternate" path when the usual one is out based on internal lists called forwarding tables. Routers, like everything else, require special software called

TCP/IP. Routers do not care what the data is. They just route datagrams to the next router available.

When you access a Web server today, this "bucket brigade" system of distributed intelligence routes datagrams to and from their destination. Only today the government is out of the picture, the Web is commercially supported now. And nobody owns it. Since the Internet is in fact a collection of servers, "distributed intelligence with nobody in charge," it is difficult to control. This will be a policy decision taken up later; for example, can a company put mission critical data on a network that nobody controls, and where there are no quality of service guarantees? On the other hand, the World Wide Web has more users today after just 10 years than the telephone network had after 50. It is clearly positioned to be revolutionary. We will delve into this topic in greater detail later as well.

1.2.11 What Is a PBX? (11)

Much of what we have covered so far has been data oriented. PBXs are for voice. PBX stands for Private Branch Exchange. It is a private telephone switch owned and maintained by your organization, on its own premises. It is also probably a misstatement to say that PBXs are only for voice. Some of them also do data, video, and so on. This is not surprising however. In today's networks, voice is really data. So is data. Any questions? Relax, we will cover the distinctions in this (true) statement later in the book.

1.2.12 What Is a Channel Bank? (12)

Channel banks take a T1 (described below) and break it into 24 64-Kbps voice channels. This is what T1s were originally designed for. A T1 was a cost effective means of taking 24 interswitch trunks between Bell or AT&T central offices, and combining them onto a four wire copper circuit. Since the technology worked so well, and since it has been out there for so many years, it has found its way to the end user as well. Today virtually all end users utilize T1 access lines to combine dial tone lines, trunk lines, and other voice and data services.

1.2.13 What Is a T1? (13)

If you take 24 64-Kbps channels, and multiply the numbers together you get 1.536 Mbps. Add 8 Kbps more that the phone company uses to tell them all where to line up and you get 1.544 Mbps. That is the T1 speed. In fact, you can think of a T1 as a 1.544 Mbps *data circuit*. Twenty-four circuits share this circuit, and each one gets exactly 1/24 of the capacity. There is a T1 section in this book as well. About halfway through it you will probably want to shoot me, or shoot yourself, but the level of detail is important to illustrate what is unquestionably the most widely deployed transport technology today. Moreover, it will provide the basis for comparison with what has changed with SONET, the technology which will replace T1, and how we can best exploit this change.

1.2.14 What Are Private Lines? (14)

T1s will carry more than switched voice traffic. They also carry "private lines," otherwise known as leased lines, custom circuits, or by other names.

A 64 Kbps voice channel can be used for other things than voice. One would be a digital data circuit operating at 56 Kbps or 64 Kbps. Essentially, creating one of these is no more difficult than plugging the appropriate channel card into the T1 channel bank.

Rather than expecting off hook and on hook information from a PBX, the digital data card tells the T1 to expect data bits in that particular time slot on the T1.

On the other hand, the phone company or user might choose to plug a four wire analog circuit card into the channel bank. The card is still designed to carry data, but is looking for an analog signal now instead of a stream of digital bits. This would require a modem.

There are over 200 different kinds of private line circuits. These range from "hoot and holler" lines to digital data circuits.

A hoot and holler line is pretty simple; in fact, it is just a pipe. It does not even ring. The user just picks it up and yells, hence the name. "I want to trade 100 shares of AT&T at 52!" or "I need a fender for a 1964 Chevy!" Junkyards and stock brokers often use this kind of circuit.

1.2.15 What Is ISDN? (15)

ISDN is designed to replace POTS (remember, that is plain ordinary telephone service) by bringing digital to the doorstep of the small office/home office (SOHO) user. Remember, that 64 Kbps channel we spoke of? Well, it can do more than just carry a voice. We have already shown that a channel can carry data at 64 Kbps as well. ISDN takes this concept a step further.

The copper cable which serves the typical small office or residence is really capable of carrying more than just a single voice call. ISDN BRI (Basic Rate Interface) takes advantage of this by providing *two* 64 Kbps channels on the same two wires which previously carried POTS. In fact, it looks and acts like a miniature two channel T1. The common name is 2B+D. That means two bearer channels with one delta channel. The "B" channels carry the actual information. The "D" channel can be thought of as the "supervisor" that instructs the two "worker Bs" what to do, how, and when.

ISDN also comes in a large size. Primary Rate Interface or PRI can be termed "T1 Deluxe." It looks, acts, and quacks much like a T1, but has some additional features. It has 23 channels that carry either voice or data traffic, plus a 24th channel (another supervisor) that controls the whole show and tells the rest what they are and what to do. The common phrase is 23B+D, or 23 bearer channels with one delta channel.

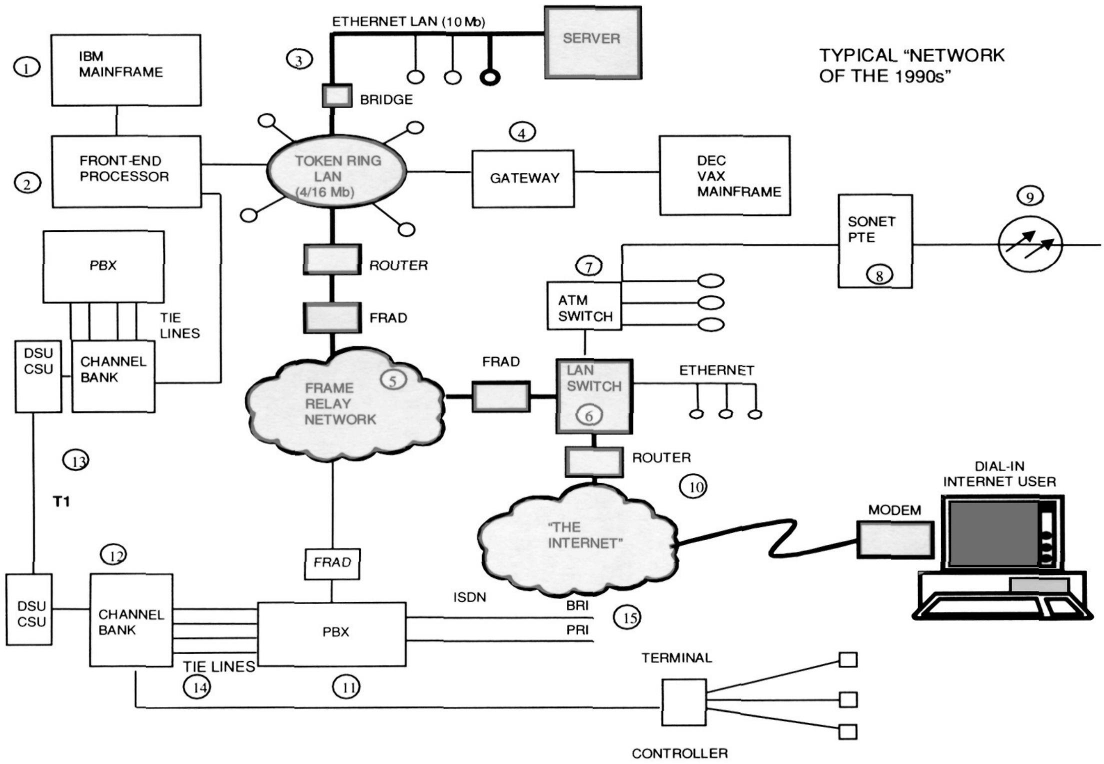

Figure 1.2: Data transaction from internet user to company server.

For your convenience in remembering which is which, just think BRI and PRI. B is for baby, P is for Papa. Both technologies are enjoying broad acceptance. For now however, PRI is used primarily for PBX access just like a smart T1. BRI, however, is surging in use as a means of surfing the Internet at 64 or 128 Kbps.

1.3 How Does Data Flow Through This Network?

Okay, now that we have all the players formally introduced, how does data actually flow through this network? For that to happen, we need to have a set of rules, otherwise known as a protocol. The problem is, there are many protocols. Figure 1.2 shows the devices which utilize the protocols shown in the following list.

- Mainframe—SNA
- Legacy WAN—SDLC

- PBX—PCM
- DSU/CSU—AMI or B8ZS
- Channel Bank—D4 or ESF
- Frame Relay—Carrier Proprietary
- Token Ring
- Ethernet
- Fast Ethernet
- ATM
- TCP/IP
- SONET

The permutations are endless. For illustrative purposes, let's trace a data transaction using the TCP/IP protocol from a homebound user dialing into an ISP (Internet Service Provider) to a server on Ethernet LAN shown as ③ in Figure 1.2.

First it is important to understand that the protocol will change many times in transmission. One way to think about this is to picture the data being put into an envelope, and that envelope either delivered, or stuffed into yet another envelope that allows another system of delivery (see Figure 1.3, TCP/IP protocol implementation).

The story begins with the Internet user, who generates the data. This has no address or envelope, it is just data. This data drops into the first envelope called a TCP segment. TCP puts an address on it, and hands it to IP.

TCP/IP can be thought of as the Abbot and Costello of telecommunications. If you think back to the old movies, Costello was the screw up. Costello could lose things. But he always had his buddy Abbot to get him out of jams. That is kind of what TCP/IP does. IP is the screwball. He can lose things or deliver things out of sequence. TCP, on the other hand, keeps tabs on IP to assure data is delivered, delivered in the correct sequence, and delivered without errors.

IP takes the TCP segment and drops it into an IP datagram. The datagram is the most elementary routable data packet. IP neither knows nor cares what is in the packet. It just routes datagrams.

Since the user is on a dial-up modem, he is using a PPP (point to point protocol) to keep things in order on the dial-up line. Therefore, the datagram drops into the PPP frame, and is transported via the dial-up connection to the Internet Service Provider (ISP).

With me so far?

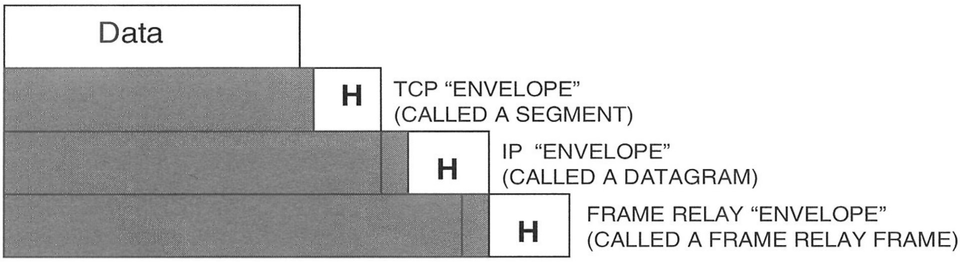

Figure 1.3: TCP/IP protocol implementation.

Once at the ISP, the equipment there opens the PPP frame and finds the datagram. The datagram is routed through the Internet via any number of routers. Based on information in the routing tables, the routers figure out an optimal path given network conditions at that instant in time. They route to an IP address specified in the header of the datagram which corresponds with a port on the LAN switch ((6) in Figure 1.2).

The LAN switch has a router card in it that can move the datagram based on information in its routing tables. This tells the LAN switch that the most optimal path is via the Frame Relay network.

In order to move the datagram over the Frame Relay network, it will be necessary to generate a new envelope. The FRAD (Frame Relay Access Device) opens a new Frame Relay frame, and inserts the datagram in the Frame Relay envelope. This envelope is routed through the public Frame Relay network based on network conditions at that instant to a DLCI (Data link connection identifier) at the far end which corresponds to a FRAD at the other company location.

The FRAD speaks Frame Relay out one side, and IP out the other. It opens the Frame Relay envelope and finds a datagram. This is much to the delight of the IP side of FRAD, which says, "Hey this is a *datagram*, and my reason for existence! I *route* these!" and promptly hands it to the next router upstream.

The next router speaks IP out one side, but Token Ring out the other. You see, to get to our destination we now need to open a Token Ring frame. The router generates a Token Ring envelope, and sends it over the Token Ring network. At the other end is a bridge.

The bridge speaks Token Ring out one side, and Ethernet out the other. It opens the Token Ring frame, and generates an Ethernet frame. It pulls out the datagram, places it into the Ethernet envelope, and sends it to the server.

The network interface card (NIC) in the server matches the MAC address (Ethernet address) on the Ethernet frame. The NIC opens the Ethernet envelope and out pops a datagram. IP software hands the datagram to TCP. TCP opens the IP datagram, assures the

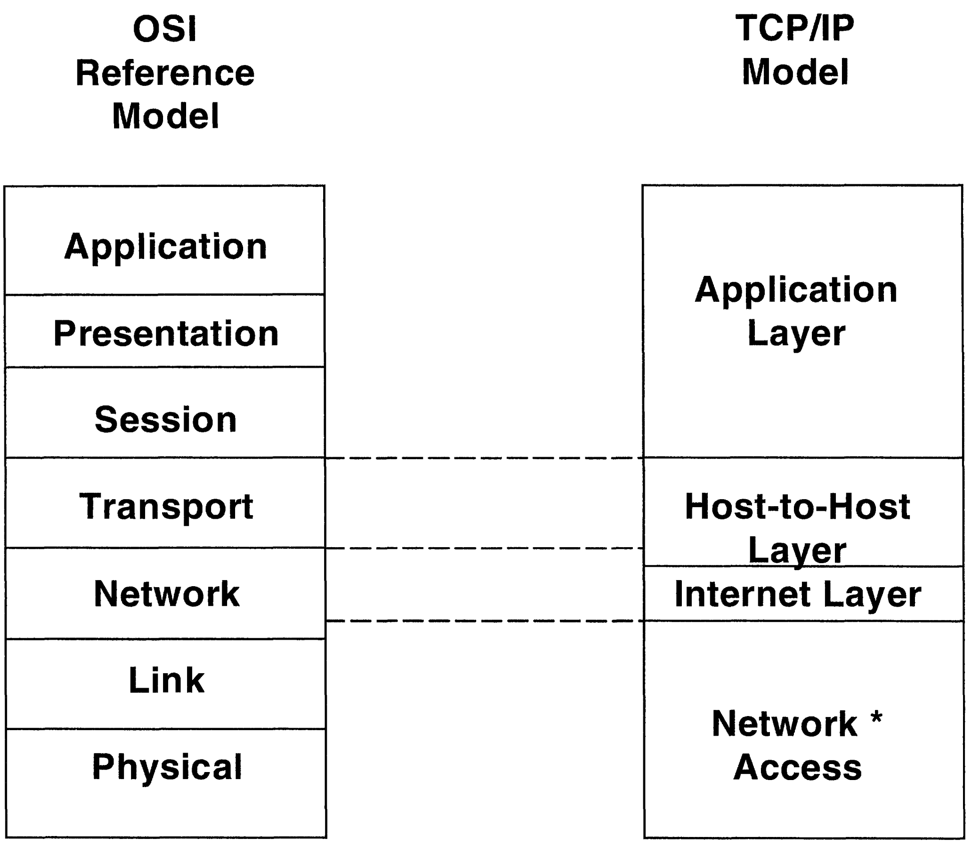

Figure 1.4: Comparison of OSI and TCP/IP protocol models. Notice that the "network" got smarter? Rather than just "dumb" layer 1 transport, the "network" today does lots of other things!

data is error free, and hands the segment to the application. The application in turn pulls out the original data, which in this case might be a request to access a Web page on the server by the original user.

So as you can see, in addition to lots of hardware, software interactions take on a significant role in the transfer of data. Protocols change innumerable times. It is no wonder why networks become so complex and why users often "black box" themselves to death! Figure 1.4 might help put things into perspective on the software side.

The OSI reference model is shown on the left. T1 lives in layer 1—the physical layer. Ethernet and Token Ring live in layer 2. TCP/IP lives in layers 3 and 4. The application lives in layer 7. The TCP/IP model on the right is simpler and shows only four layers, the application layer, TCP layer, IP layer, and network access layer.

Notice that the network layer got bigger? What used to be just a dumb layer 1 network layer in the OSI model now goes to layer 2 1/2. This is because networks have become smarter. Frame Relay (Figure 1.5), for example, does routing, switching, segmenting, error control, and other higher level functions that used to be I.S. rather than telecommunications functions. We will discuss the implications of these smarter networks later as well.

The "legacy" version of data transfer is shown in Figure 1.6. Mainframes and front ends use SDLC, and the applicable components are highlighted for your review. Figure 1.6 also shows the "envelope in an envelope" example in a comparison of IP and the U.S. Post Office.

1.4 PBX to PBX

Voice communications between PBXs are handled differently as shown in Figure 1.7. The backbone link between the two is a T1. The T1 terminates in a DSU/CSU that converts the digital signal from the T1 to one that can be understood by the PBX. The DSU/CSU is an integral component of the channel bank, or can even be internal to the PBXs themselves. If a PBX has a T1 card, a channel bank is not necessary, nor is a DSU/CSU. This is shown in Figure 1.7 as the "traditional" way to break a single T1 into 24 individual channels—in this case, *tie lines* between PBXs. Tie lines are used to connect two PBXs in neighboring buildings, or even different cities. One benefit of using them is the ability to provide four or five digit dialing transparently and simultaneously in both locations.

1.5 Token Ring User to Local Ethernet User

Figure 1.8 is highlighted to show a local Token Ring user connecting with a local Ethernet user via a bridge. Bridges are now smart enough to not only do the filtering and forwarding but to also do the format translation (Token Ring to Ethernet).

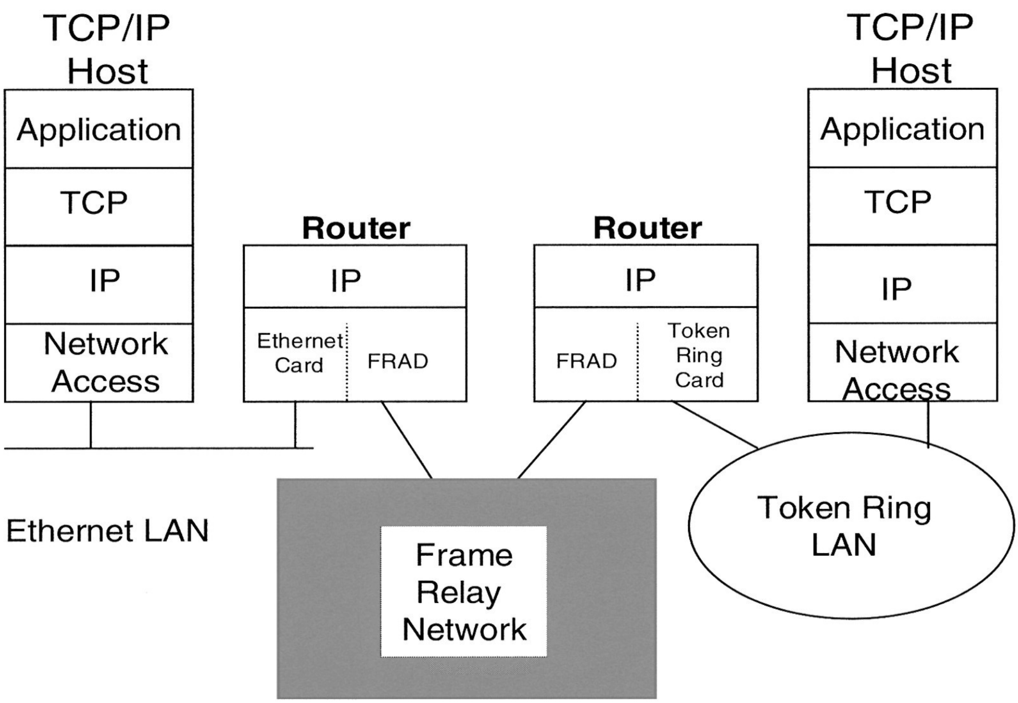

Figure 1.5: TCP/IP Internetworking.

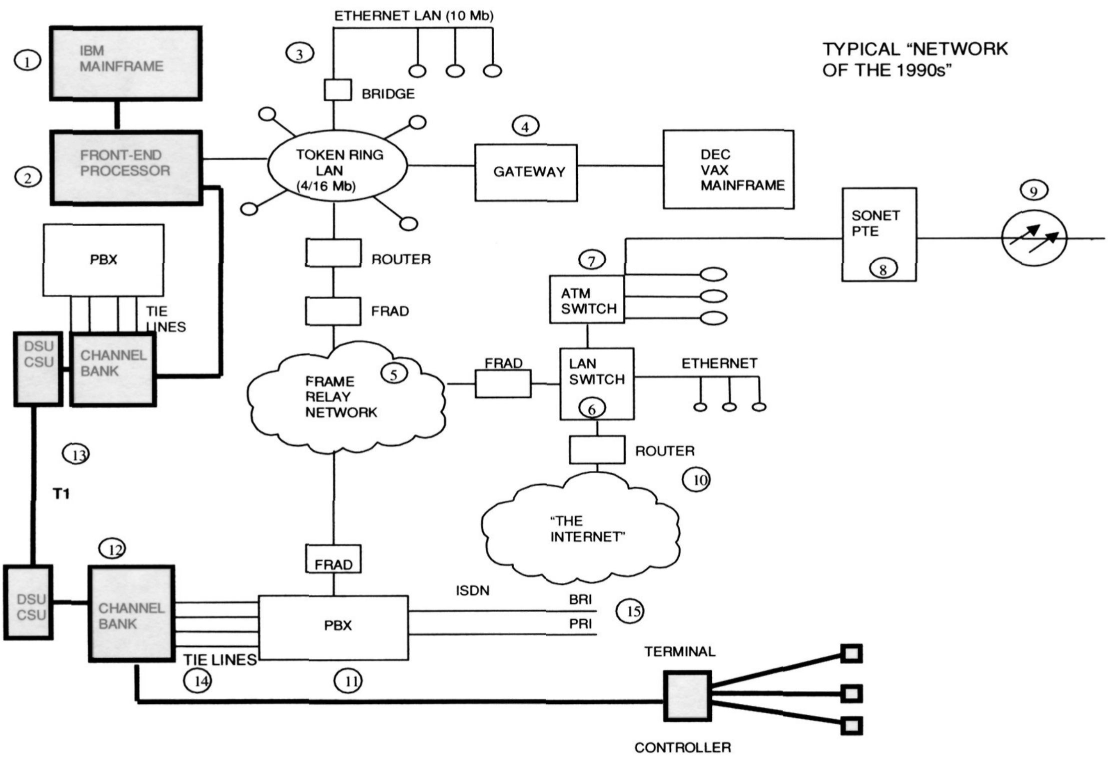

Figure 1.6: SDLC private line.

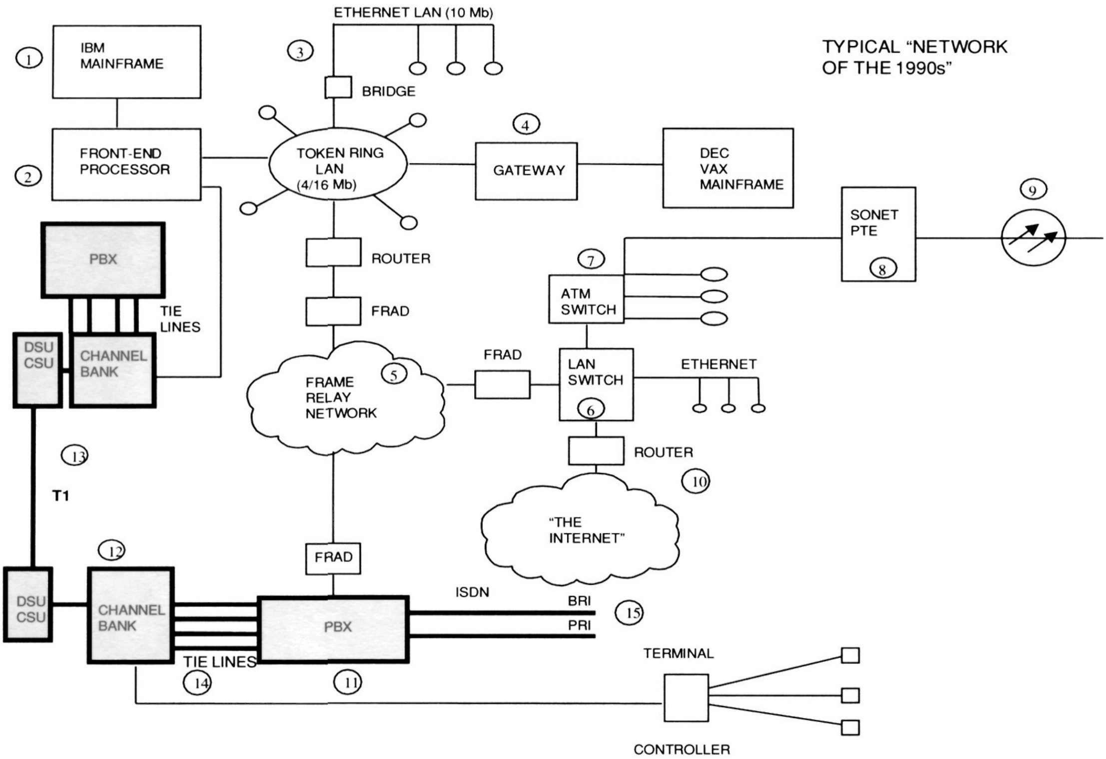

Figure 1.7: PBX to PBX.

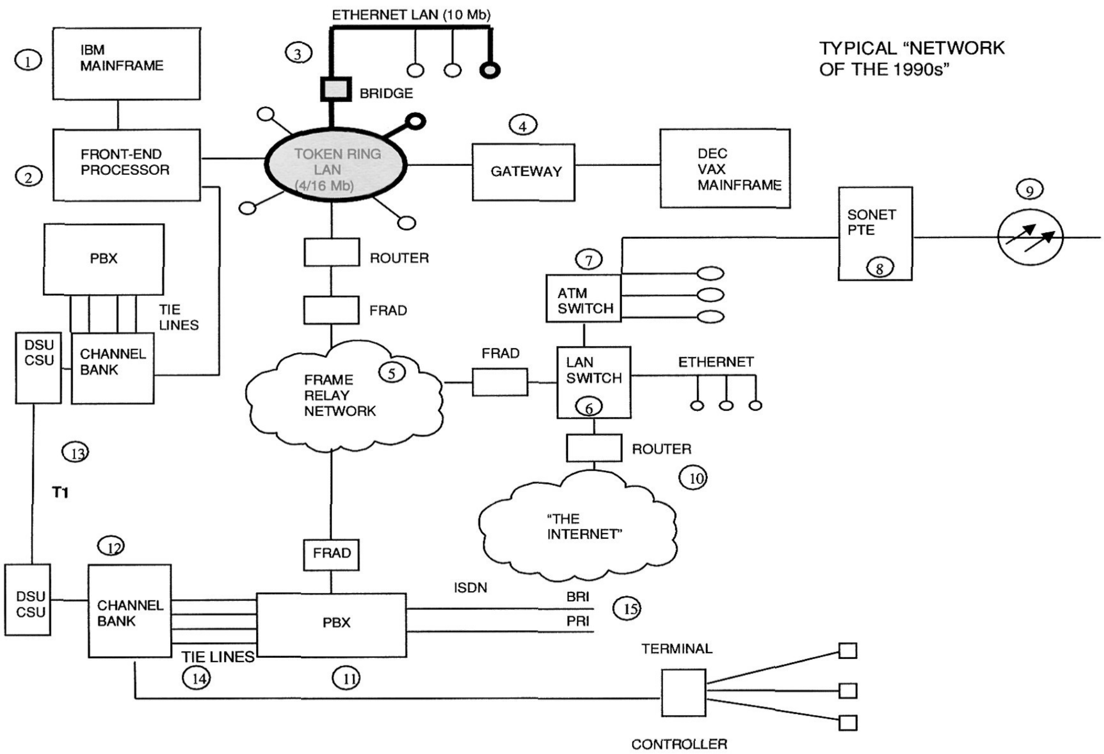

Figure 1.8: Token Ring user to local Ethernet user.

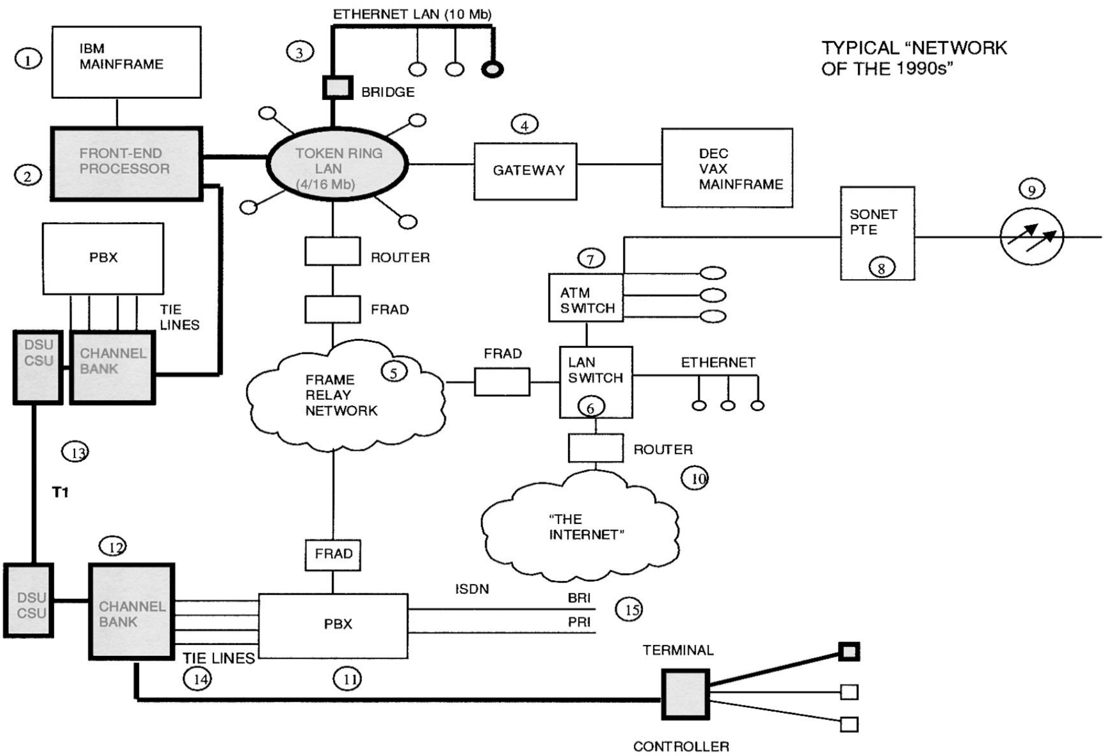

Figure 1.9: Local Ethernet to remote terminal.

1.6 Local Ethernet User to Remote Terminal User via Private Line

Imagine that the Ethernet user needed to communicate with a distant SDLC terminal device. Several format translations must take place, but through each of them (Ethernet→Token Ring→SDLC) the data stays the same as shown in Figure 1.9. Only the envelopes change.

1.7 Local Ethernet User to Remote Ethernet User via Frame Relay

Frame Relay is a wonderful technology for connecting LAN users together. As we stated earlier, Frame Relay is not *isochronous*, but that is okay. Neither are most LANs. Additionally, if traffic is discarded due to network conditions, the intelligent higher level data protocols simply retransmit. This is illustrated in Figure 1.10.

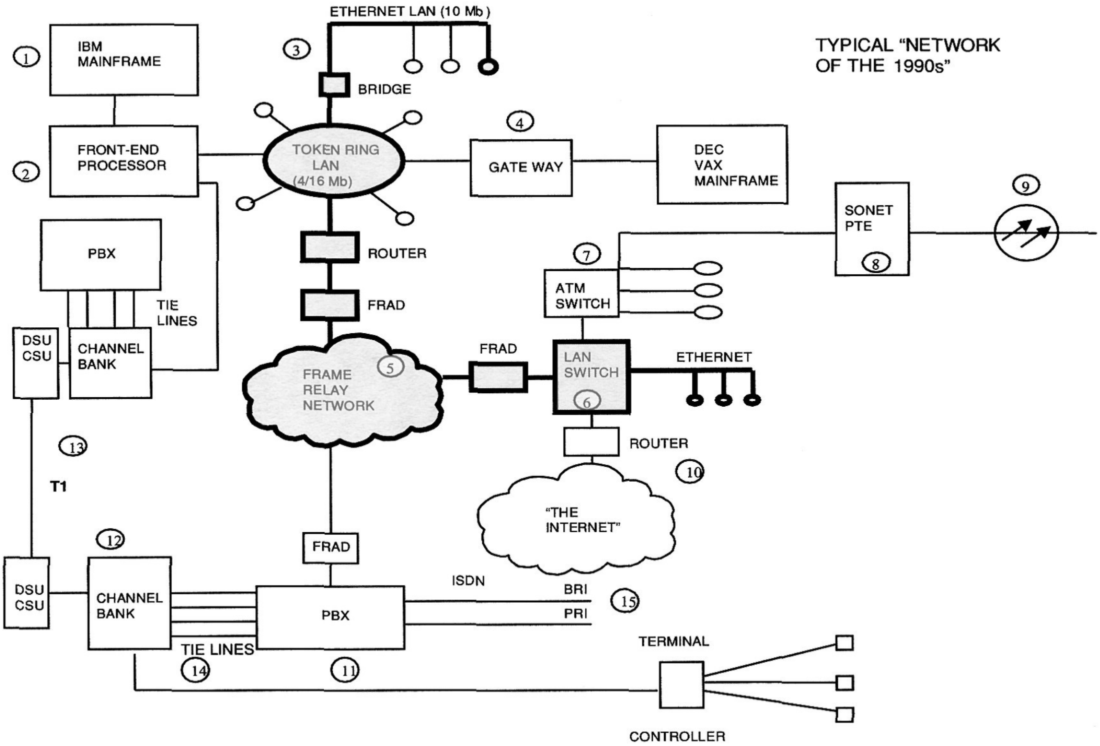

Figure 1.10: Local Ethernet user to remote Ethernet user via Frame Relay.

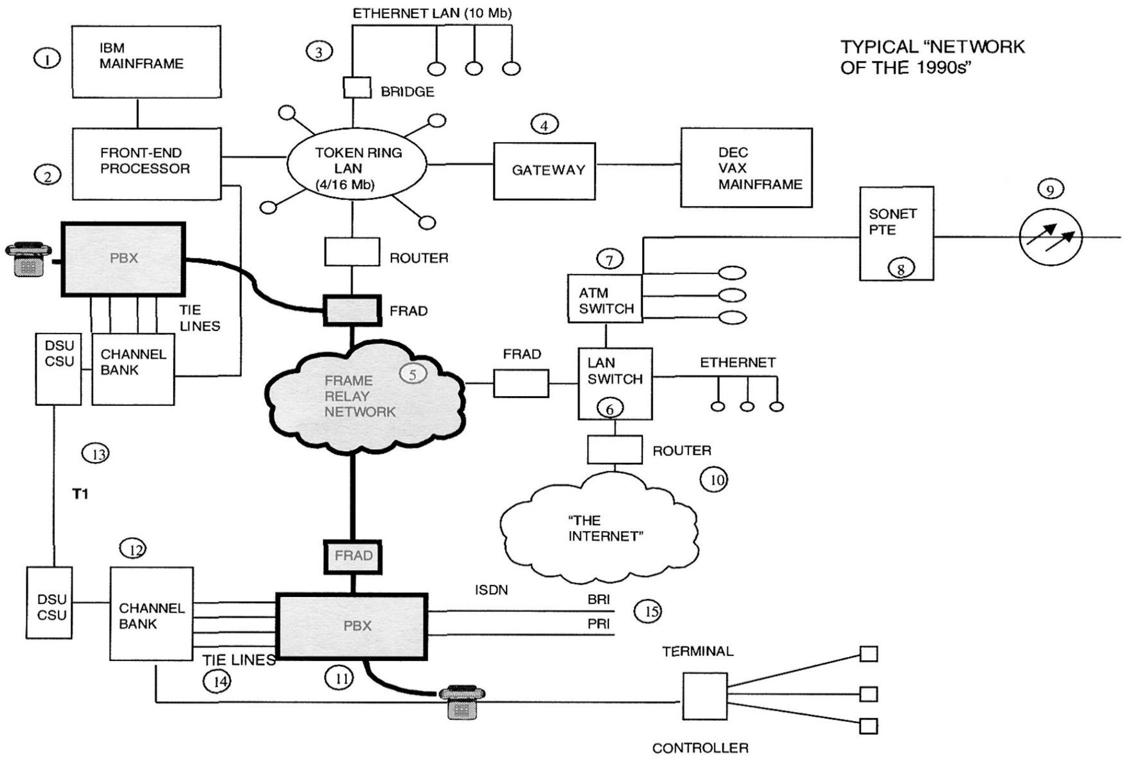

Figure 1.11: Voice over Frame Relay.

1.8 Voice Over Frame Relay

Unlike data, voice poses problems for Frame Relay. In the case of "discard eligible" traffic, which part of my voice is not important? Voice is less forgiving, and in fact, is the *premium* kind of traffic on most kinds of Frame or Cell Relay. For those of us used to cheap voice and expensive data, this represents a major change. This is illustrated in Figure 1.11.

1.9 SONET and ATM

How can one derive the benefits of frame- or cell-based networks but still have *isochronous* capability? For these kinds of applications, ATM will be the ticket. And ATM will ride on SONET. See Figure 1.12.

For the moment, think of SONET as railroad tracks—large fixed investment with carrying capacity dictated by the number of trains run in a given period. ATM cells are the trains. Even though the capacity of the tracks (SONET) stays the same, the carrying capacity goes

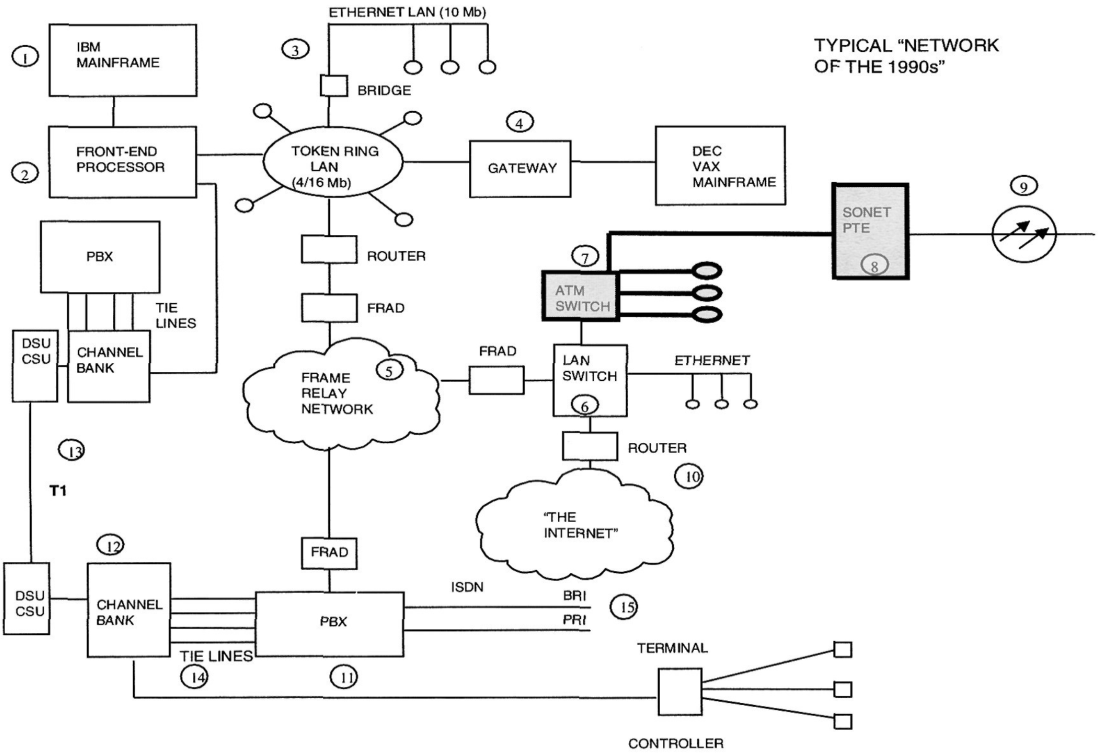

Figure 1.12: SONET and ATM.

up as more ATM cells are sent. ATM therefore provides the perfect way to share SONET bandwidth. When some users are not actively using the tracks, others can run trains—on a second by second basis!

Now that the basics are out of the way, let's take a closer look at the *Local Area Network* in Chapter 2.