DOCUMENT RESUME

ED 417 713 IR 018 780

AUTHOR	McCain, Ted D. E.; Ekelund, Mark
TITLE	Computer Networking for Educators. Second Edition.
INSTITUTION	International Society for Technology in Education, Eugene, OR.
ISBN	ISBN-1-56484-107-3
PUB DATE	1996-00-00
NOTE	215p.
AVAILABLE FROM	International Society for Technology in Education, Customer Service Office, 480 Charnelton Street, Eugene, OR 97401-2626; phone: 800-336-5191; World Wide Web: http://isteonline.uoregon.edu
PUB TYPE	Books (010) -- Guides - Non-Classroom (055)
EDRS PRICE	MF01/PC09 Plus Postage.
DESCRIPTORS	*Computer Networks; *Computer Uses in Education; Educational Administration; Educational Planning; *Educational Technology; Elementary Secondary Education; Information Networks; Strategic Planning
IDENTIFIERS	Apple Macintosh; *Connectivity; IBM PC Compatibles; Technology Integration; *Technology Plans

ABSTRACT

This book introduces the basic concepts of connecting computers together and provides technical background necessary for constructing small networks. For those already experienced with creating and maintaining computer networks, the book is intended to encourage the creation of a school-wide network. The book is divided into two main sections: an introduction to networking in schools (6 chapters) and an introduction to the technical side of networking (10 chapters). The chapter headings are as follows: (1) "How Can Networking Enhance the Use of Computers?"; (2) "Advanced Networking Concepts"; (3) "Why Network Classroom Computers?"; (4) "Why Network Computers for Administration?"; (5) "Why Consider School-Wide Networking?"; (6) "How To Plan a School-Wide Network"; (7) "Technical Details for both IBM-compatible and Macintosh Computers"; (8) "Technical Issues in IBM-compatible Networking"; (9) "Simple Network Examples for IBM-compatible Computers"; (10) "Complex Network Examples for IBM-compatible Computers"; (11) "Technical Issues in Macintosh Networking"; (12) "Simple Network Examples for Macintosh Computers"; (13) "Complex Network Examples for Macintosh Computers"; (14) "A Complex Network Example Using IBM-compatible and Macintosh Computers"; (15) "Networking Apple II Computers"; and (16) "The Never-Ending Chapter". A glossary of terms is included. (AEF)

ISTE Publications

Computer networking for educators

Second edition

Ted D.E. McCain & Mark Ekelund

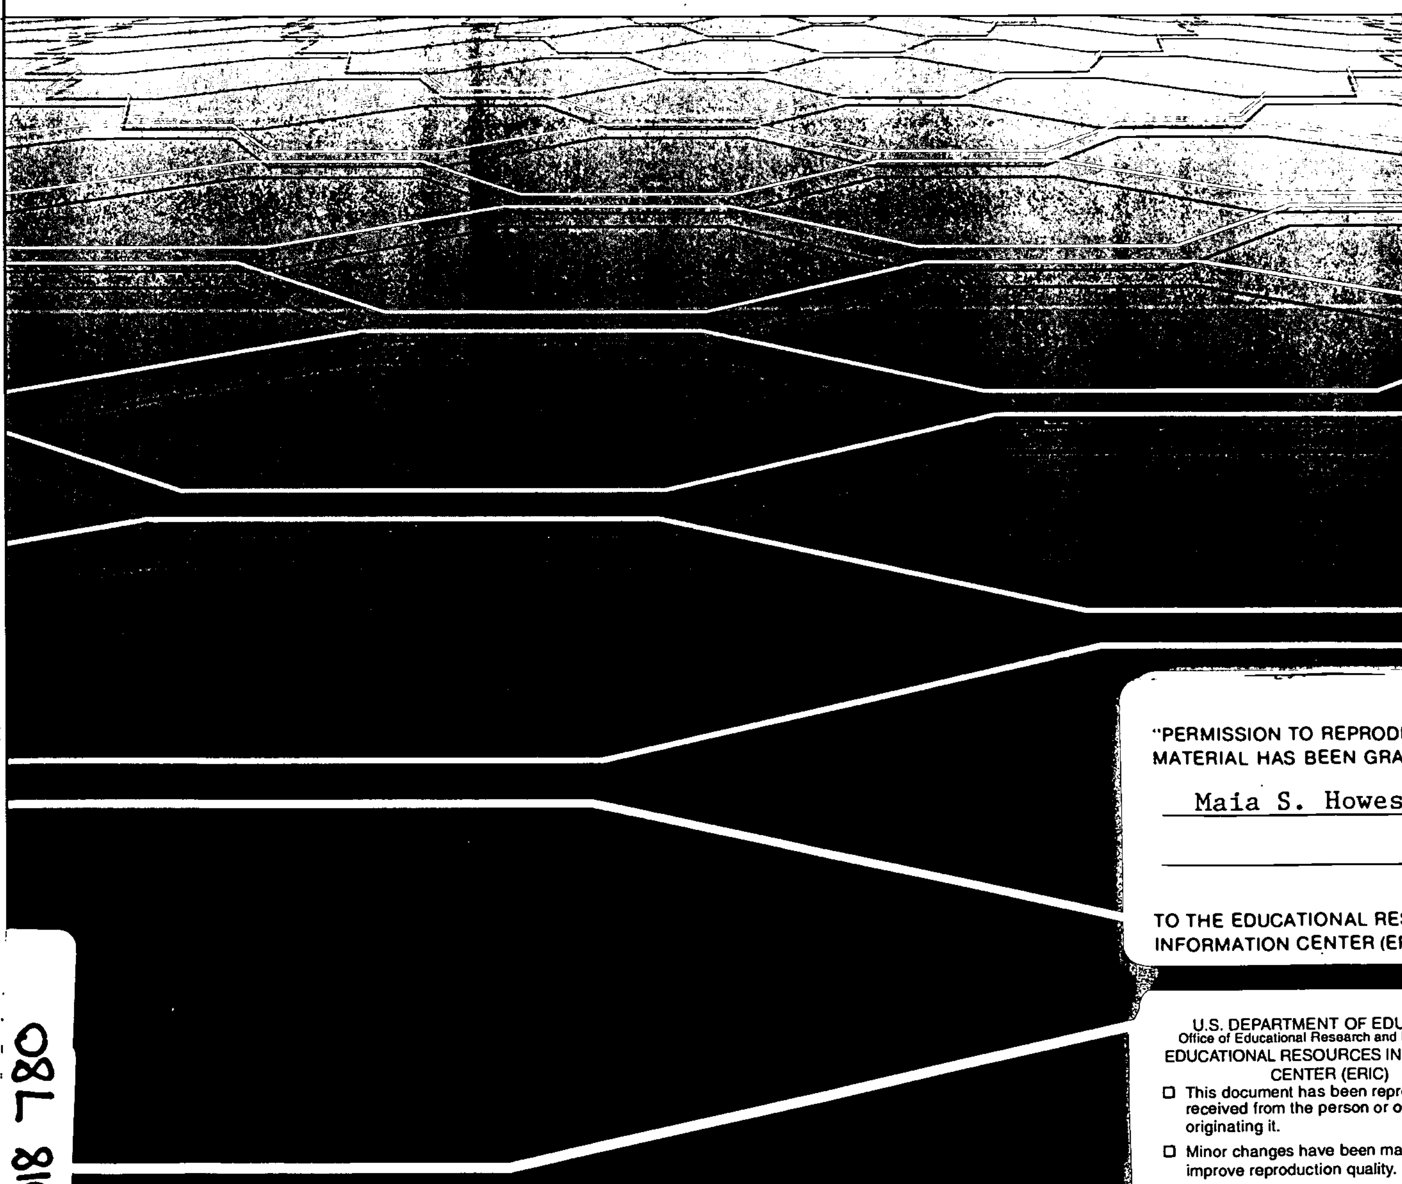

Project Coordinator

Jean Hall, ISTE

Editors

Bob Barber, ISTE; Christy McMannis, The Electronic Page; Carl Durance, WatNet Technologies; Ian Jukes, Futureview Consulting; Rick Withers and Mike Hoebel, Education Technology Centre.

Book Design and Production

Ted D.E. McCain

This book was produced using *Microsoft Word*, *Adobe PageMaker*, and *Macromedia FreeHand* running on a PowerBook 5300cs computer equipped with 36 megabytes of RAM and a 500 megabyte hard disk drive.

Commercial EPS clip art used in this book created by: Artbeats, Art Right Image Portfolios, Artware, C.A.R. Inc, Clipables, Dream Maker Software, Fine Art Impressions, Image Club Graphics Inc., Metro Creative Graphics Inc., and 3G Graphics.

These materials were originally produced for The Education Technology Centre of British Columbia.

Administrative Office
1787 Agate Street
Eugene, Oregon 97403-1923
Phone: 541-346-4414
Fax: 541-346-5890

Customer Service Office
480 Charnelton Street
Eugene, Oregon 97401-2626
Order Desk: 800-336-5191
Order Fax: 541-302-3778

America Online: ISTE
CompuServe: 70014,2117
Internet: iste@oregon.uoregon.edu
World Wide Web: http://isteonline.uoregon.edu

Second Edition
ISBN 1-56484-107-3

Computer networking for educators

Ted D. E. McCain & Mark Ekelund

Table of contents

Preface

Section One
An introduction to networking in schools

Chapter 1 - How can networking enhance the use of computers?

Chapter 2 - Advanced networking concepts

Chapter 3 - Why network classroom computers?

Chapter 4 - Why network computers for administration?

Chapter 5 - Why consider school-wide networking?

Chapter 6 - How to plan a school-wide network

Section Two
An introduction to the technical side of networking

Chapter 7 - Technical details for both IBM-compatible and Macintosh computers

Choosing workstation software

Chapter 9 - Simple network examples for IBM-compatible computers

Chapter 10 - Complex network examples for IBM-compatible computers

Chapter 11 - Technical issues in Macintosh networking

Chapter 12 - Simple network examples for Macintosh computers

Chapter 13 - Complex network examples for Macintosh computers

Chapter 14 - A complex network example using IBM-compatible and Macintosh computers

Preface

Why we wrote this book

The primary goal of using technology is to empower people to work more effectively. Computers empower business people to store and retrieve information, perform mathematical calculations, produce graphics, etc. Computers help teachers keep track of student performance, report to parents, and present instructional material to students. Computers help students access information, produce written work, and learn new concepts. Computer technology has become so powerful that it is changing virtually every aspect of modern life.

Children growing up today must be equipped with skills to operate and succeed in an electronic world. That brings us directly to the reason we wrote this book: networking computers opens up a whole new world of possibilities for communicating with others and for accessing needed information. Students should be exposed to the world of computer networks because using these tools is a life skill.

Teachers and administrators can certainly use networks, too. Networks provide an infrastructure that gives teachers and administrators new ways of monitoring student performance and that facilitates new ways of learning. Networks even make some administrative tasks in a school easier – and wouldn't that be good news!

We believe that networking computers creates the potential for fantastic new learning opportunities. Don't be fooled by the uses that networks have when first introduced in a school – exciting new uses are just around the corner. That is not to say that networks are not useful already; in fact, the advantages of networking are already too great to ignore. But even more dramatic uses will emerge in the near future, and we don't want you to miss out.

There are two main purposes for this book. First, if you are new to computer networking, we hope to introduce you to the basic concepts of connecting computers together and equip you with some of the technical background necessary to begin constructing small networks. Second, if you are experienced with creating and maintaining computer networks, we hope to help you consider the creation of a school-wide network. Regardless of your previous experience with networking, we hope this book will take you further, so that you will be able not only to take advantage of the power that networks have right now but also to utilize the truly startling developments in world-wide communication that will soon be available.

So who did what?

While both of us have had experience with networking various computer systems, to ensure some normality for our personal lives during the writing process, we divided up the responsibility for the writing of specific areas. Mark has more expertise with networking IBM computers, so he was responsible for the chapters dealing with the technical details and examples of IBM-compatible networking. In addition, Mark tackled the chapters on planning school-wide networks, and the glossary.

Ted has more expertise with Macintosh networking, so he wrote the chapters dealing with Macs. Ted also wrote this preface, the first two chapters dealing with basic and advanced networking concepts, as well as the chapters covering networking classroom computers, networking administrative computers, and networking Apple II

computers. The chapters covering school-wide networks, the technical details shared by both Mac and IBM-compatible computers, the example of a combined Mac and IBM network, and the look in the crystal ball were collaborative efforts. Finally, Ted was responsible for creating all of the diagrams, assembling the overall book into a coherent whole, and creating the book layout.

Who are these guys anyway?

We thought you might like to know a little bit about our backgrounds. We hope these biographies will convince you that we are qualified to write this book because we have led sufficiently misdirected lives by spending inordinate amounts of time fooling around on computers. So, in abbreviated form, here are our lives.

Ted McCain

Ted graduated from Simon Fraser University with a degree in computer cartography. After some initial work in graduate school, he worked for several years in the computer industry as a programmer, salesperson, and consultant. He felt the desire to move on to other things and his much-checkered past includes running a hunting and fishing lodge with his wife in the remote Chilcotin region of British Columbia. He entered the teaching profession in 1980 and began teaching Computer Science 11, Data Processing 11, and Mathematics. In the early 1980s, he developed courses in Computer Science 12 and Data Processing 12 for the Maple Ridge School District. In 1984 and 1985, Ted was a member of the British Columbia Provincial Curriculum Committees for Computer Studies 11 and Computer Science 12. As an Administrative Assistant at Maple Ridge Secondary, he supervised pilot-school development of two major school administrative packages (HARTS for Apple IIe computers and MacSchool for Macintosh computers).

Since 1986, Ted has been the school-based Coordinator of Instructional Technology for Maple Ridge Secondary (MRSS). This position is a half-time release from the classroom to coordinate and implement instructional technology in all curricular areas in the school (MRSS has over 300 computers, 12 networked labs, 5 different types of hardware, and a school-wide Ethernet network). Ted is also the Business Education Department Head at MRSS. Recently, he developed courses in Desktop Publishing 11 and 12 for the Maple Ridge School District.

For the past 11 years, Ted has also run his own computer consulting company, through which he develops custom software packages, consults on the purchase and implementation of computers and networks, does desktop publishing, writes articles on the use of technology and graphic design, and gives presentations and workshops on the business and educational uses of computers. Ted has done consulting and training for a wide range of clients including businesses such as Apple Computer Inc., Westerly Yacht Sales, as well as several Toyota dealerships and school districts in Canada and the United States, The Education Technology Centre of British Columbia, Okanagan College and the University of Alaska. Ted is the author of *Designing for Communication – The Key to Successful Desktop Publishing* and *Teaching Graphic Design in All Subjects,* both published by The International Society For Technology Education, and is a co-author of *The Microcomputer – A New Tool For Tilling The Intellectual Garden,* published by Simon Fraser University.

But if you think Ted has had a wild past, just check out this next guy!

Mark Ekelund

Mark has had a varied career (this might qualify for understatement of the year, as you will soon see). After completing a B.Sc. degree in Physics at the University of British Columbia, he earned his commercial pilot's license at an aviation college. Subsequently, he worked as a research pilot investigating thunderstorm activity in Alberta (we are fairly sure he was struck by lightning at this time, which would explain the erratic behavior he has displayed ever since). Finding his life unstimulating, Mark made the logical career move and became an engineer building oil drilling platforms from ice in the Arctic. Still unfulfilled, he became an oceanographer studying water currents around Newfoundland. When all of these passing fancies had lost their appeal, Mark finally turned to the real challenge – education! He has been in education for twelve years and he is still in education today, which tells you it must be an interesting field to work in, or he would have been gone long ago!

Mark taught at the secondary school level; his areas of specialty were Physics, Mathematics, and Computer Science. While teaching, he also earned a Masters degree in Educational Administration from the University of British Columbia. Mark is now the computer coordinator for the Quesnel School District in central British Columbia. Quesnel schools have a significant number of computers and Mark has been responsible for developing a district-wide plan for implementing those computers in K-12 classrooms. Quesnel uses primarily IBM and compatible computers in its schools and all schools already have Novell networks. Mark's plans include extending those networks to every classroom in the district, making links between schools, connecting every school to a world-wide network, and accessing many more educational resources. Of course, all this is only a hobby. His only real goals have to do with his wife and his four children, who have an innumerable menagerie of pets.

Ted and Mark met while working together on the Provincial Curriculum Committees for Computer Science courses, almost a decade ago. Therein lies a lesson in the difficulty of change. Those committees lobbied hard for frequent revision of Computer Science courses due to the fast-paced development of the computer field. But while the rest of the world has changed drastically, that curriculum remains unchanged.

It was while working on the curriculum committees that Ted and Mark discovered their unique ability to argue interminably over something they agree on. After several peaceful years of working on separate projects, for reasons that remain unclear to this day, they ended up in a meeting to discuss the need for informing educators about the benefits of networking. They decided at that meeting to write this book and they haven't stopped arguing since.

Section One

An introduction to networking in schools

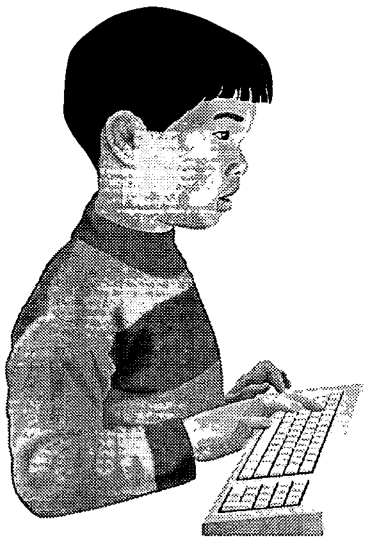

Chapter One

How can networking enhance the use of computers?

Chapter 1

How can networking enhance the use of computers?

Topics covered in this chapter...

- *Basic concepts of what networks can do*
- *Print share networks*
- *Print spooling networks*
- *Disk sharing networks*
- *File sharing networks*

So you want to know about networking...

Networking provides education with a number of powerful features. Your interest in networking is easy to understand. Learning about networking, though, is not so easy. The topic quickly mushrooms into a complex field of endeavor that requires a lifetime of high-level academic study. Since such an abrupt career change is likely not your prime motivation for reading this document, we will try to keep our discussion of networking limited to concepts and practical suggestions required for planning, implementing, and maintaining smaller microcomputer networks.

What is a network?

Let's begin by defining a network. A network is the interconnection of computers for the sharing of information and resources. Some of the resources most commonly shared on a network include:

- printers and plotters
- disk drives
- programs
- data files

Information shared on networks includes:

- electronic mail
- instructional files such as assignments, tests, and schedules
- work group files such as budgets, drafting plans, agendas, etc.
- databases of information such as customer lists, encyclopedias, dictionaries, and directories of people, countries, and companies
- any other information worth sharing that can be stored electronically

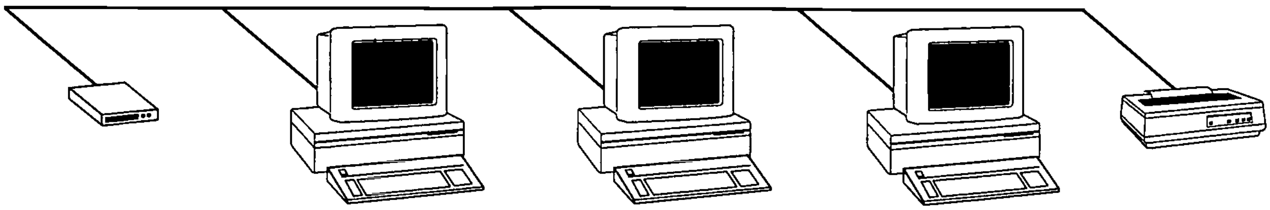

The whole idea behind connecting computers together to form networks is to share information, software, and any resources, including hardware, more efficiently.

Networks for microcomputers are usually designed to meet the information and resource-sharing needs of users whose computers are located reasonably close to each other – in an instructional lab or the offices of a small business. For this reason, the majority of microcomputer networks are known as "Local Area Networks" or LANs for short.

Why network computers?

You can see that networking allows people to communicate with each other. This is a prime reason for hooking computers together. Another is to cut down the cost of purchasing expensive peripherals like printers, plotters, etc. Yet another is to add a measure of security to the programs and data that are stored on computers. Finally, networking centralizes many of the administrative tasks associated with running computers and thereby introduces a glimmer of hope that those overseeing the use of computers in a business or institution can have some sort of personal life!

All in all, once in place and running smoothly, a network makes life easier for all concerned. The only problem is getting it in place and running smoothly. For the novice, this is a daunting task. Networking takes you below the more friendly surface of your previous computer use, into the technical world of computer systems. If you have never ventured into technical computing, you may be overwhelmed. The "techies" who can make your head spin when they start talking the "computerese" of networking don't help much at all! Relax. One of the major goals of this book is to ease you into this field without giving you a lot of information you don't need and without making you feel uncomfortable in the process.

Those new to networking please take note...

For many of you, subsequent chapters of this book will be unclear because you have not mastered basic networking concepts. Reading this chapter will help, so don't jump ahead! In addition, at times you may find it necessary to come back to this chapter for a quick refresher on a particular feature before continuing. To put your mind at ease, we want you to know that this is allowed, even encouraged!

We've split the discussion of networking concepts into two chapters. This chapter will cover basic concepts and some common features in network software. The second chapter covers a few of the more sophisticated concepts and features of networking. If you are new to networking, you may want to skip over Chapter Two for now. You

will still be able to understand the material presented in Chapters Three through Six and you can return to the information in Chapter Two when you have gained a little more experience with networks.

Now, let's talk about what networking can do.

Building a conceptual foundation

There's no use beating around the bush, so we might as well just come right out and say it: computing is technical. Sure, computer manufacturers and software companies do their best to hide many of the underlying complexities from our sight, but beneath that friendly surface lurks a highly sophisticated machine, a technical marvel. Why mention this? Simply because you cannot enter the arena of school-wide computer networking without embracing some of the technical aspects of computing.

Mastering the technical aspects of computing is a life-long endeavor, so we won't attempt to present too much detail in our discussion. Instead, we will talk in more general terms, providing you with a conceptual foundation. We will cover the basic concepts of network capabilities, computer operating systems, and application software, all from the point of view of network administration. Let's begin.

Networking: what it can do

People have been networking large computers for well over 25 years. The networks on these large computer systems are now run by sophisticated programs capable of performing many complex functions. An examination of these large networks would involve learning a great number of complicated and esoteric computing concepts that go well beyond the scope of this publication. Instead, we will follow the development of networking as it evolved for the microcomputers that burst onto the scene in the late 1970s. This will allow us to cover some of the important conceptual foundations for the networking you will likely be doing without having to consider the problems pertaining to large computer networks.

In the beginning was the stand-alone computer

When microcomputers first appeared, they were stand-alone computers. A stand-alone computer is not connected to any other and is completely dedicated to use by a single person. Hence the terms, stand-alone or dedicated computer. When a person is thinking or otherwise engaged, the computer simply waits until he or she is ready to continue. Furthermore, peripheral equipment such as a printer, sits waiting for the single user to issue a command to print something. While this is an acceptable situation for small offices and in the home, it is inefficient when there are many computer users in an office or institution.

Printer sharing

As microcomputers grew in popularity in the early 1980s, the inefficiency of having many stand-alone computers in a single place, each with its own printer, became glaringly apparent. An inordinate amount of money was being spent for printers that sat idle most of the time. Linking all of the computers together to a single printer solved the problem. This sharing of a printer created the first network for microcomputers. Each computer could send printing down the cable to the printer, but could not communicate with any of the other computers on the network.

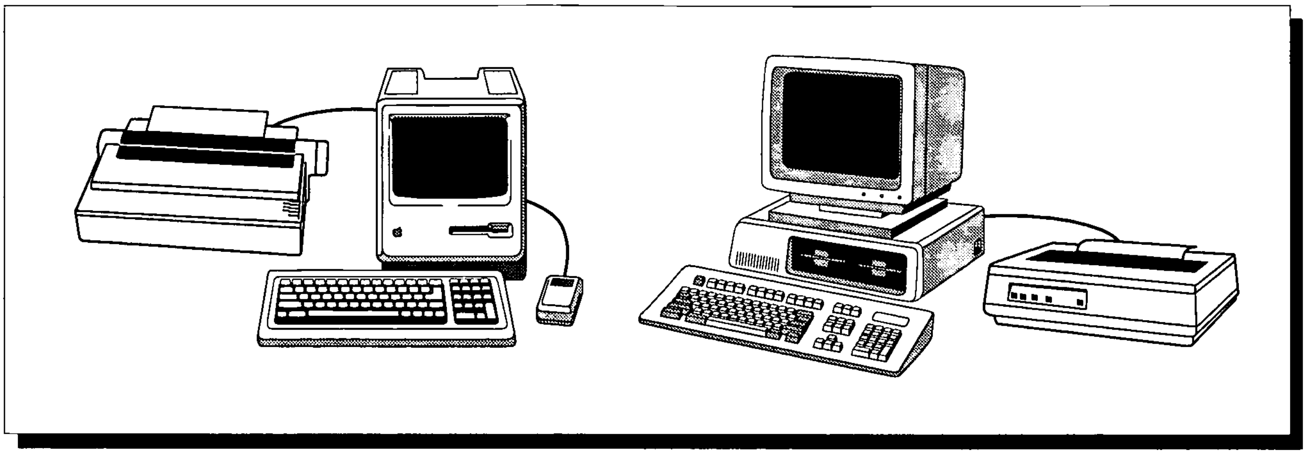

Stand-alone microcomputer systems. The IBM PC was introduced in 1981 and the Mac came along in 1984. You could get a complete system for $5000. These machines revolutionized computer use but inefficiencies were soon discovered when more than one system was placed in a business or institution.

Today, you can share many peripheral devices that were never intended for a networking environment. These devices were originally designed for a single user working on a stand-alone computer and did not contain the instructions necessary for multiple user access. Devices on the market make it possible to connect stand-alone peripherals such as printers and plotters to a network. These devices contain all of the logic missing from stand-alone peripherals. Furthermore, some current network systems allow you to connect a stand-alone peripheral directly to the network server; the server allows the device to be shared by all network users.

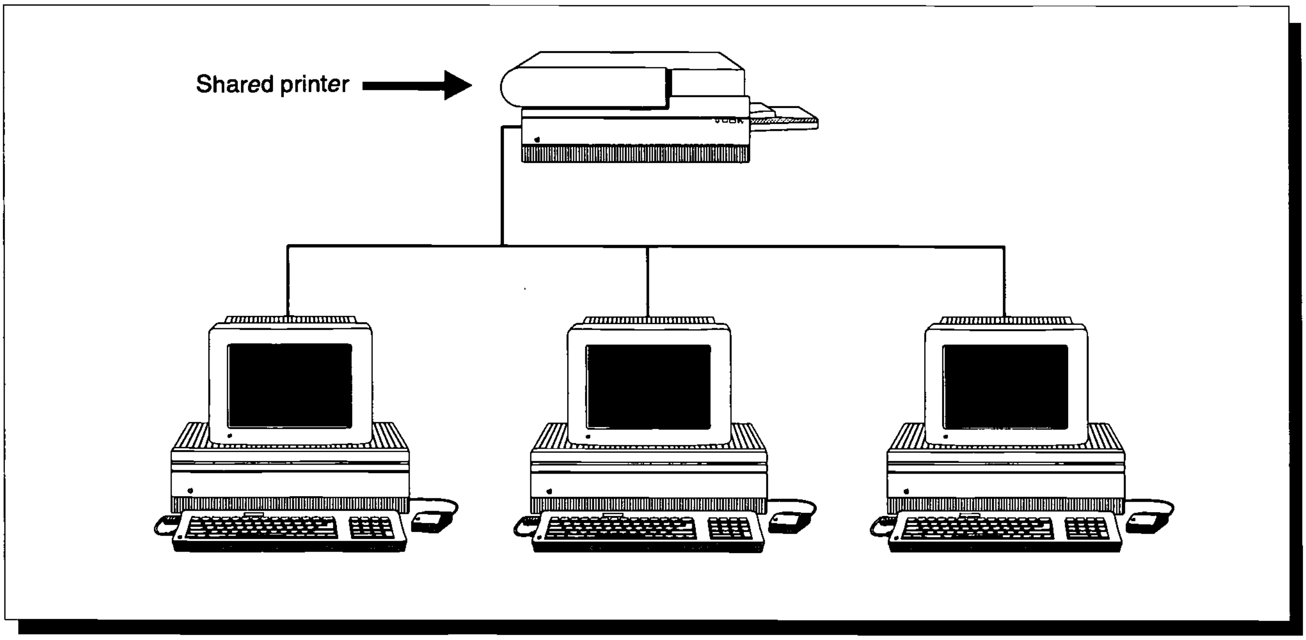

Print sharing is the simplest form of networking. It allows more than one computer to access a single printer. All signals on this network are one-way communications from one of the computers to the printer. While the cost of purchasing printers is reduced with such a scheme, a number of problems still remain.

Print spooling

While printer sharing solved the cost problem of stand-alone printers, it did not address an even more frustrating problem associated with printing – the slow speed of transferring documents to paper. Printers were so slow that even single-page documents could cause a computer to "hang" in an unusable state for several minutes while waiting for the printer to complete its job. This meant that a person would be unproductive if they needed the computer in order to continue doing their work. Furthermore, if several people tried to print at the same time, the printer responded on a first-come, first-serve basis. Then all the computers would "hang" in an unusable state while waiting for the printer. If you were the last person in line, you could wait for a long time. Something had to be done.

Since the speed of the printer could not be increased, another solution had to be found. Some sort of intermediate device was needed to receive and quickly memorize the document to be printed so the computers could be released for the people to continue with productive work. The only high-speed storage device that was capable of receiving and memorizing print files at the desired rate was the hard disk drive.

So a hard disk drive was connected to one of the computers on the print share network and a program was written to intercept all printing requests. When a request was detected, this program stored the information onto the hard disk and then released the computer that initiated the printing. Then the program would send the information to the printer from the computer with the hard disk drive, not the computer that originally requested the printing.

This had two major advantages. First, the hard disk drive could receive and store print information much, much faster than the printer could print. Thus, a computer requesting printing was released much faster than if it had printed directly to the printer. Second, when several people tried to print at the same time, instead of having all of the computers wait for the printer, this new arrangement only tied up a single computer during the printing process – the computer with the hard disk drive. The people working at the Digital Equipment Corporation (DEC) called this process of sharing a printer by using a hard disk drive as an intermediate device Shared Peripherals Operating On Line and consequently the programs that provided this ability to store print information on a hard disk came to be known as print SPOOLers.

Print information was received and printed by the spooler on a first-come, first-serve basis. Those files waiting for printing formed a print queue that was saved on the hard disk. The spooler was programmed to ensure that each file was completely printed. Thus, when printing was interrupted by an operator or a paper jam, the spooler would restart printing from beginning of the file or the beginning of that particular page. Because all of the print files were stored on the hard disk, the computer with the spooler could be turned off overnight or for the weekend before all printing was completed. When the spooler was restarted, it would simply continue printing until the print queue was empty.

Some early spoolers were designed to work in the background on a computer that was being used for some other productive task such as word processing. The idea was that the computer would run the spooler during the idle time when the user was thinking or otherwise engaged. While this was a good idea, most of the early microcomputers were not fast enough to handle the dual role. The person using the computer quickly got frustrated waiting for the computer to quit spooling and respond to their input. Those who were using the spooler for printing were frustrated at its sluggish performance. Most print spoolers had to operate as the only program running

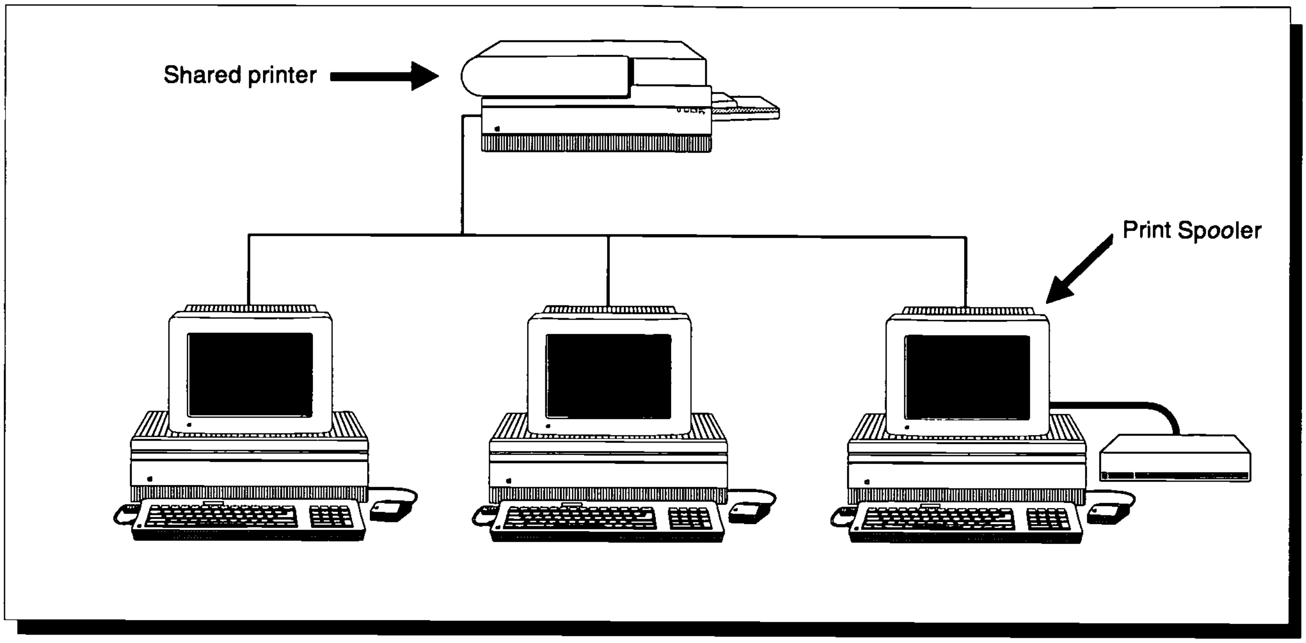

It doesn't look like much of a change, but the addition of a hard disk drive to intercept and store print files made a huge difference for network users. The spooler program ran in the background on the computer with the hard disk. Running in the background meant that the spooler only used idle CPU time to send print files to the printer.

on the computer with the hard disk. This computer had to be completely dedicated to the task of providing the network with print spooling and could not be used for any other work.

Today, two main strategies exist for print spooling. One method is to alter a printer so that it becomes a node on the network. A node is a network device that is given a unique identification number so messages can be directed to that specific location. To make a printer operate as a node requires the addition of a small circuit board with chips containing the instructions necessary for sending and receiving network messages. The server software is programmed to receive print files and then send them to the printer node one at a time. This strategy is used on Macintosh networks. The other method is to connect the printer directly to the network server. In this way, the server software can be programmed to share the use of a stand-alone printer. No additional hardware modifications are required to make the printer operate on the network because the server handles all of the networking signals and then simply sends the print files to the printer one at a time. The Novell network management system for IBM-compatible computers operates in this way.

Using a print spooler took some time to master. If you were the only person printing, the spooler was slower in responding, compared to printing directly to the printer, because the process of storing the print file onto the hard disk and then sending it along to the printer created a delay of a few seconds. Users new to spooling often printed two or more copies of the same document, thinking that the printer had somehow not received the first one or two printing requests. New users were even more likely to print multiple copies when there were already other files lined up in the print queue waiting for their turn to print. The printer seemed to have gone wild when the document was finally printed by the spooler – not once, not twice, but three, four, or more times!

Print spooling also introduced the role of the network administrator – the person who installed the spooler program and ensured that it ran smoothly. This entailed activating the spooler at the appropriate times, deleting files from the print queue when someone inadvertently printed twenty copies of the same file, and moving important files to the front of the print queue if required.

Disk sharing

Printers are not the only peripheral devices in a computer system that are expensive and sit idle most of the time. Disk drives also fit this profile. But as microcomputers were being used for more and more complex tasks, most of the software began to be written for computers with a minimum of two disk drives – one drive for the Disk Operating System and the program and another drive for the user's data disk. In large offices and instructional computer labs, the cost of installing second disk drives on every computer was prohibitive. To address this problem, new network programs were written to allow all the users on a network to share a single hard disk drive.

The hard disk drive was used because it has a large storage capacity. Each user on the network was given a small space of fixed size on the hard disk. The hard disk was large enough to allow for many user spaces. The network program was written to listen for requests from any computer on the network. As long as the request was accompanied by the correct password, the network program would allow the computer requesting access to read and write files to and from that user space. As far as the user was concerned, his or her computer would function as if it had two disk drives. The only difference was the delay involved in sending and receiving information over the network cabling and the possible wait if more than one user was trying to access space on the shared hard disk at the same time.

The computer connected to the shared hard disk drive was completely dedicated to the task of running the network program that gave users access to their user spaces. When a request for access was received, the network computer would send information over the cabling to the computer that made the request. This process was like asking a server to bring you your dinner from the kitchen, so the computer running the network software came to be known as the network server.

The idea was fairly straightforward. First, create a number of user spaces on a hard disk. Actually making the idea work was a bit more complicated and the person with most of the work was the network administrator. The administrator had several difficult tasks. The first major task was to divide the space on the hard disk into user spaces. This process was called partitioning the hard disk to create the user spaces, partitions, volumes, or network disks (all of these terms are synonymous). While the creation of a single user space was relatively easy, the creation of an entire network of user spaces was not. Any user space on the hard disk had to be made up of contiguous sectors on the hard disk and once a space was created, its size could not be adjusted. The process is sort of like putting up fences on a wide open plain. At first there seems to be so much room that it doesn't matter where you put up the fences. However, the fences can't be moved once they are put up. If you don't have an overall plan to guide you, the fences may be placed incorrectly or too much space given away at first, so that you don't have enough room left at the end.

Once the user spaces were created, the network administrator had to maintain the directory of all of the user names and their passwords. This information was needed to check the validity of any request to access a user space (a process called logging on, logging in, signing on, signing in, etc.). Some networks required the administrator

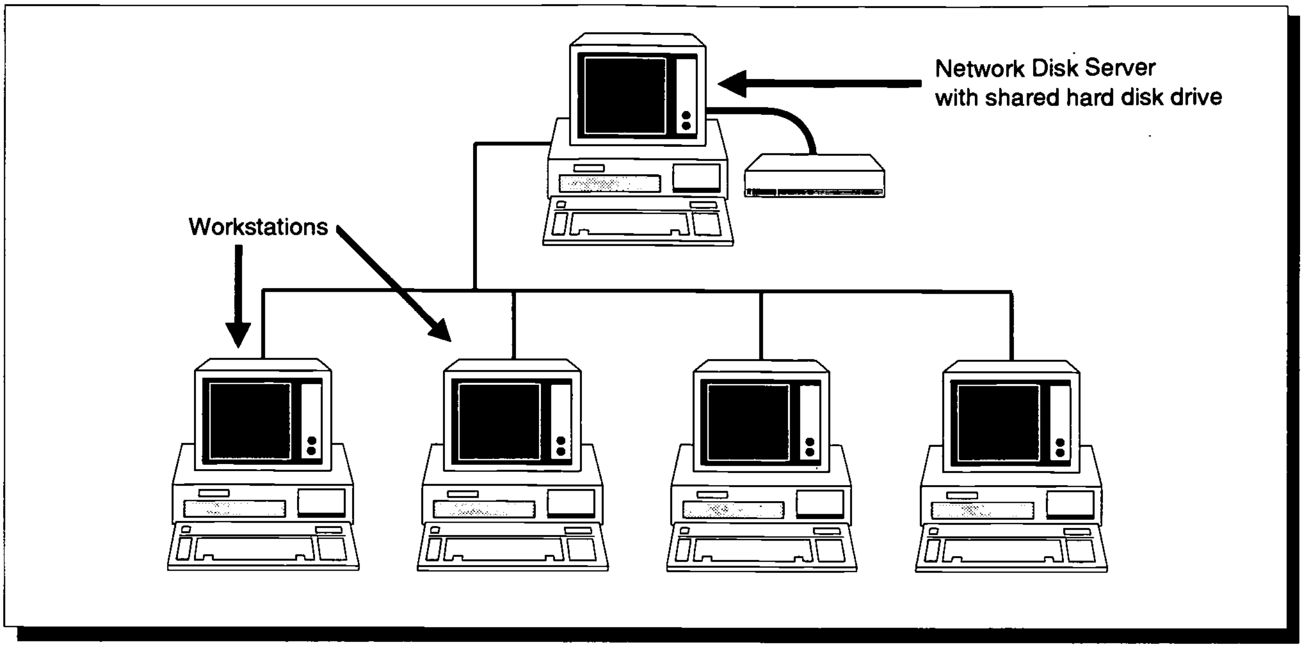

When the use of a hard disk drive is shared over a network, a computer is set aside with the single task of running the program that allows network users access to their space on the hard drive. This computer is called the server and it cannot be used as a workstation.

to make any password changes. If a user forgot his or her password, the administrator had to look it up in the directory. But the biggest task in maintaining the network was dealing with new users and requests for larger user spaces from existing users.

When a new user wanted to gain access to the network, the administrator had to create a user space in the unused area of the hard disk. If the entire hard disk was used up, the administrator had to delete someone first and then assign that space to the new user. When someone no longer needed their space on the network, their name and password were simply deleted from the directory and their user space freed up for other use. This created a "hole" on the hard disk – a free space between two other active user spaces. After a few months of use, these networks developed the "Swiss cheese" effect where there were holes all over the hard disk.

A request for expanded space from an existing user was by far the most difficult task for the administrator. The existing user space had to be deleted and then another larger space created somewhere else on the hard disk. The administrator had first to ensure that the user had made backup copies of all files stored in the existing user space before the space was deleted. Then came the difficult part – creating a larger space. There may have been lots of free space on the hard disk, but if it was spread all over the place in a number of holes, there could be a problem. Remember that user spaces had to be made up of contiguous sectors on the hard disk. Unless one of the holes was large enough to accommodate the user's expanded space request, the larger user space could not be created even if the total amount of free space in all of the holes was large enough. Now the administrator had to find a user space that was located right next to one of the holes and delete that user space to create a hole that was large enough for the expanded request. Then the user space that had been deleted had to be recreated somewhere else. It was not a pleasant experience! Often, requests for an expanded user space were met with a flat out "NO!"

To make this job easier, updated network software was written that gave the administrator the ability to compact the hard drive. Compacting physically moves all of the existing user spaces together and moves all free space into one large contiguous area on the hard disk. While it was time consuming and required that the network be shut down, compacting did provide a welcome solution to a major frustration for administrators. Today, some network software even does this job automatically, making it transparent to the administrator (that is, invisible). It is likely that this will be universally the case within a few years.

However, the automated server software being introduced today does not negate the need for planning. Networking is much more about people and their needs than it is about cables, computers, and network software. The importance of this network planning cannot be overstated. When network administrators began using disk share networks, they soon discovered the need for carefully mapping out how a network would be configured. Planning became the number one priority for the network administrator and became even more important when networking involved file sharing, a strategy that is described below.

File sharing

Disk share networks were a great step forward in the efficient use of computers in instructional labs and large offices. However, they were not without problems. A major inefficiency was discovered in the use of programs. Each user space on a disk share network is a world unto itself. Files saved by the user were theirs alone, and for the user's personal files, that's exactly the way they wanted it. But network administrators noticed that, in addition to personal data files, most users also used their network spaces to store copies of the programs they used frequently. Consequently, their requests for space on the network were quite large. Furthermore, if several users worked with the same program, they each had a copy of the program in their user space, thus wasting an enormous amount of space on the hard disk. It was also wasting a lot of time, especially in instructional settings where a number of student users would try to run the same program at the beginning of a class. The server had to go to each user space on the hard disk for each request for the program and send it to one computer, repeating the process for each computer requesting the program. The student who requested the program last had to wait a long time before they could begin to work.

Again the network software programmers went to work, and devised a system where the network server only stored one copy of a program that was to be used on the network. When multiple requests for the same program were received by the server, it only had to load the program from the hard disk once. Then the server could send the program out to the computers on the network much like the way a radio station broadcasts a song out to many individual radios. By having a single copy of the program in a public area, individual users needed much smaller personal user spaces while still being able to run the program from the network server. This is called file sharing. File sharing is an extension of disk sharing; thus, file sharing networks are also referred to as disk sharing or disk serving networks.

Network administrators next discovered that users were inadvertently saving their personal work in the space set aside for programs. They also found that some users were making illegal copies of programs for use on their home computers. To prevent these problems, file share networks were developed that made a distinction between public and private spaces. Public spaces were reserved for files such as programs that would be shared among many users. Network users could only read information from these spaces. They could not write information to that area. Furthermore, the network

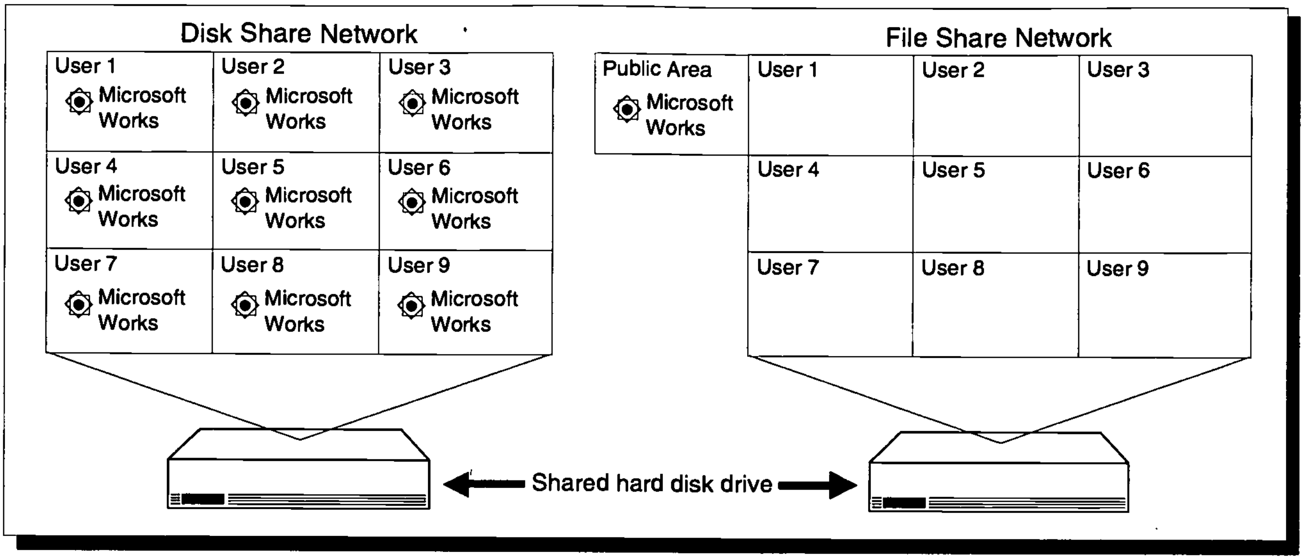

Private areas on a disk share network are a world unto themselves, accessible only by the single user who knows the password. This leads to a great inefficiency in disk storage because each user on the network may have a copy of exactly the same program. This is illustrated by the contents of the network hard disk on the left. A file share network adds the capacity for creating public areas that can be accessed simultaneously by all network users. This is illustrated by the contents of the network hard disk on the right. There only needs to be a single copy of a program file stored in the public area. More space can be used in the private areas for the files of individual users.

software would not allow users to make copies of programs in public spaces. Some networks gave these public areas an "execute only" designation. The only one who could write information to a public area was the network administrator, and even then he or she had to know the write-password to gain such access.

Administrators soon discovered that other files besides programs could be shared among users. They created public read-only spaces for files containing network news. Teachers began using public read-only spaces for items such as copies of notes, assignments, tests, dictionaries, or encyclopedias. However, the number of public areas soon became so great that users had difficulty finding just the public areas that were of interest to them. A new level of sophistication was added to network software that allowed the administrator to limit access to each public space to those users for whom it was appropriate.

This meant that users had to be grouped into categories with the same network needs. The administrator would create a category of user and give specific privileges to that category. Then, when an individual user was identified as being of that category, they automatically received all of the associated privileges. Users were generally unaware of the use of categories on a network because they were not a visible feature. The concept of categories or groups of users was important only to the network administrator, who had to devise a way to get the appropriate resources to many sets of users.

Private user spaces on a file share network were usually only accessible by one user and were password-controlled for security. Thus, the network reverted to a disk sharing strategy for private spaces while using a file sharing strategy for public areas. However, the distinction between public and private spaces has become less clear on many new network systems. Networking software for some microcomputers allows administrators and users themselves to specify whether a directory or subdirectory will

be publicly accessible or not. Some networks have evolved even further with the ability to specify whether individual files are public or private. This is the case in UNIX networks, and the networks found on many minicomputer and mainframe computer systems.

Furthermore, some networking software has evolved to the point that many traditional administrative tasks have become so automated that they are transparent even to the administrator. Network software such as Novell Netware allocates space on the server's hard disk dynamically, as needed, thus eliminating much of the planning and calculation associated with configuring a fixed volume size network. Today, Novell administrators need only concern themselves with buying a bigger hard disk for the server if the current one is getting full.

Summary

If you have only used a stand-alone computer, the ideas of print sharing, print spooling, disk sharing, and file sharing may take some time and effort to master. However, these concepts underlie the operation of all networks on all makes of computers. Therefore, you must take all the time you need to feel comfortable with these ideas. You may choose to skip the next chapter if your head is already spinning, but make sure you master the material in this one before moving on.

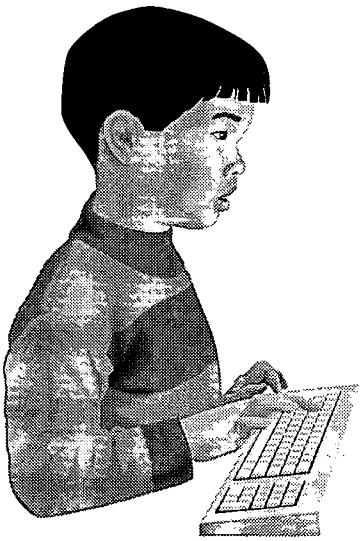

Chapter Two

Advanced networking concepts

Chapter 2

Advanced networking concepts

Topics covered in this chapter...

- *Electronic mail and conferencing*
- *Repeaters and star controllers*
- *Routers and bridges*
- *Gateways*
- *Backbones*
- *Network modems, telebridges, and remote access*

So you still want to know more...

Microcomputer networks can grow into very sophisticated systems. They can grow in physical size and they can grow in their capability to provide powerful features to network users. The growth in size, both in the number of computers joining the network and the distance between the machines, can present challenging technical problems. When two or more networks are joined together, the problems become even more challenging. In this chapter, we will look at a few of the more sophisticated features available on networks and explore some solutions to the problems associated with growing networks.

Electronic mail and conferencing

Electronic mail

With the development of file sharing networks for microcomputers came the addition of a powerful feature – electronic mail. Electronic mail is the ability to leave messages for other users when they are not around. The network software saves the message on the server's hard disk until the message is read. This was a logical extension of the network's use of the hard disk. A seemingly trivial addition, electronic mail has grown into the major reason many people have an interest in joining computer networks.

Conferencing

The use of electronic mail led to the development of electronic conferencing on some microcomputer networks. Conferencing is an extension of electronic mail that allows a user to place messages into a file in a publicly accessible area where any network user can read any message that any other user has placed there. Usually, a single file is dedicated to the discussion of a specific topic. All new messages are appended to the end of the file so that it is possible to read the entire discussion that has taken place on a particular subject. Thus, a continuing discussion on a subject can take place even when the participants cannot physically meet.

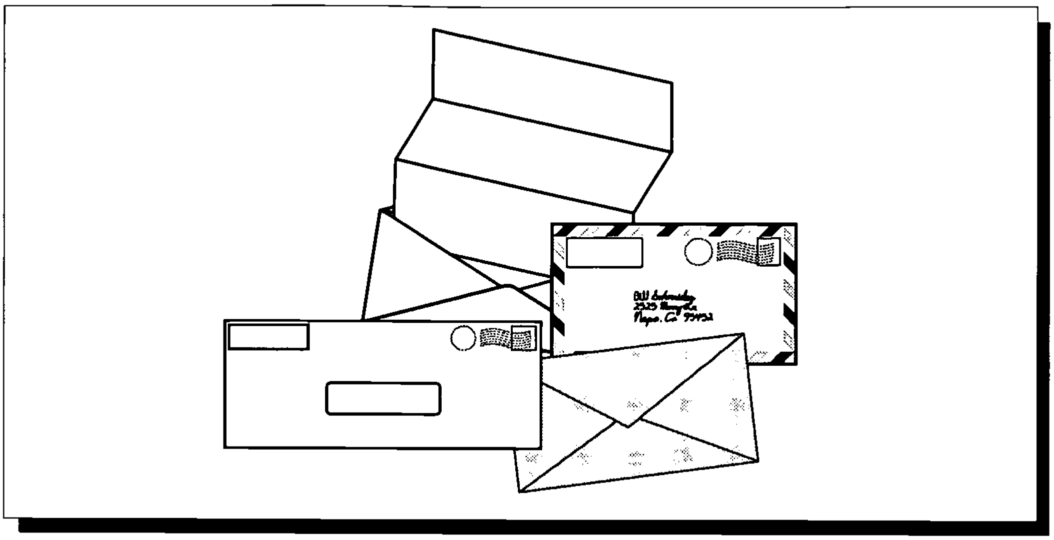

We use the Postal System to get messages to people we cannot meet with physically in the same place at the same time. Electronic mail is used for exactly the same purpose, with the added benefit of no postal charges and instantaneous delivery of your message virtually anywhere around the world.

Internetwork connections

With the dramatic rise in the number of microcomputer networks in the late 1980s, many schools and businesses found themselves having several networks in a building that were complete worlds unto themselves. While there is certainly nothing wrong with this scenario, many people began to see that connecting the networks together into a larger school-wide or office-wide network would provide useful extensions to existing network services. Electronic mail and electronic conferencing could involve more users, perhaps allowing everyone in an office or school to communicate with each other. File sharing could be widened to give more users access to programs or data files. High-cost peripheral equipment could be shared widely, reducing the need for purchasing additional equipment for separate networks. School-wide and office-wide networking would also allow users to access their private user space from any workstation in the building.

Joining separate networks together is called internetwork connection. Let's discuss four of the most common strategies to join networks together.

Repeaters and star controllers

A repeater is a device that amplifies network communications to enable signals to travel over greater cable lengths. Network signals fade due to resistance in the wires. After a certain distance, a signal becomes too weak to be understood by computers at the other end of the cabling. Signals may travel acceptably over the cabling inside two separate networks, but the combined cable length when the networks are connected may be too great for an unboosted signal. The repeater may be all that is needed to make the combined network operate effectively.

Repeaters have only two ports. Signals travel in one side of the repeater and are amplified before being sent out the other side. Multi-port repeaters sit at the center of a star network configuration. These devices are called star controllers; each port is

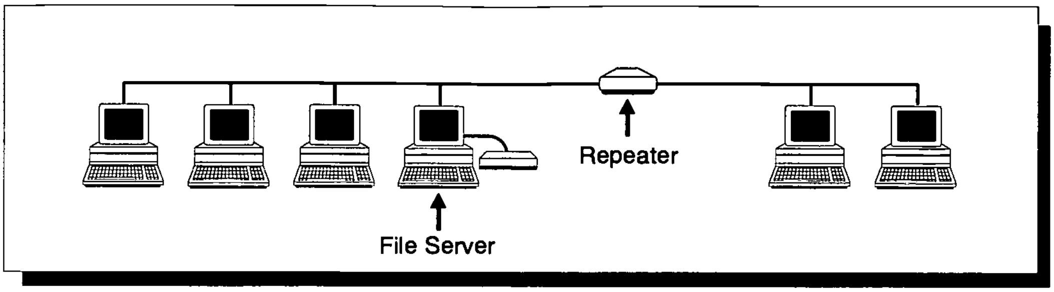

In this network, two computers are located a significant distance from the file server. A repeater is used to boost the strength of the signals travelling both ways on the network cabling to ensure that network signals do not fade before reaching the two computers that are separated from the rest of the network.

attached to an "arm" of the star. Star controllers create what is known as an active star topology, where network signals are boosted so they can travel at an acceptable strength anywhere on the star.

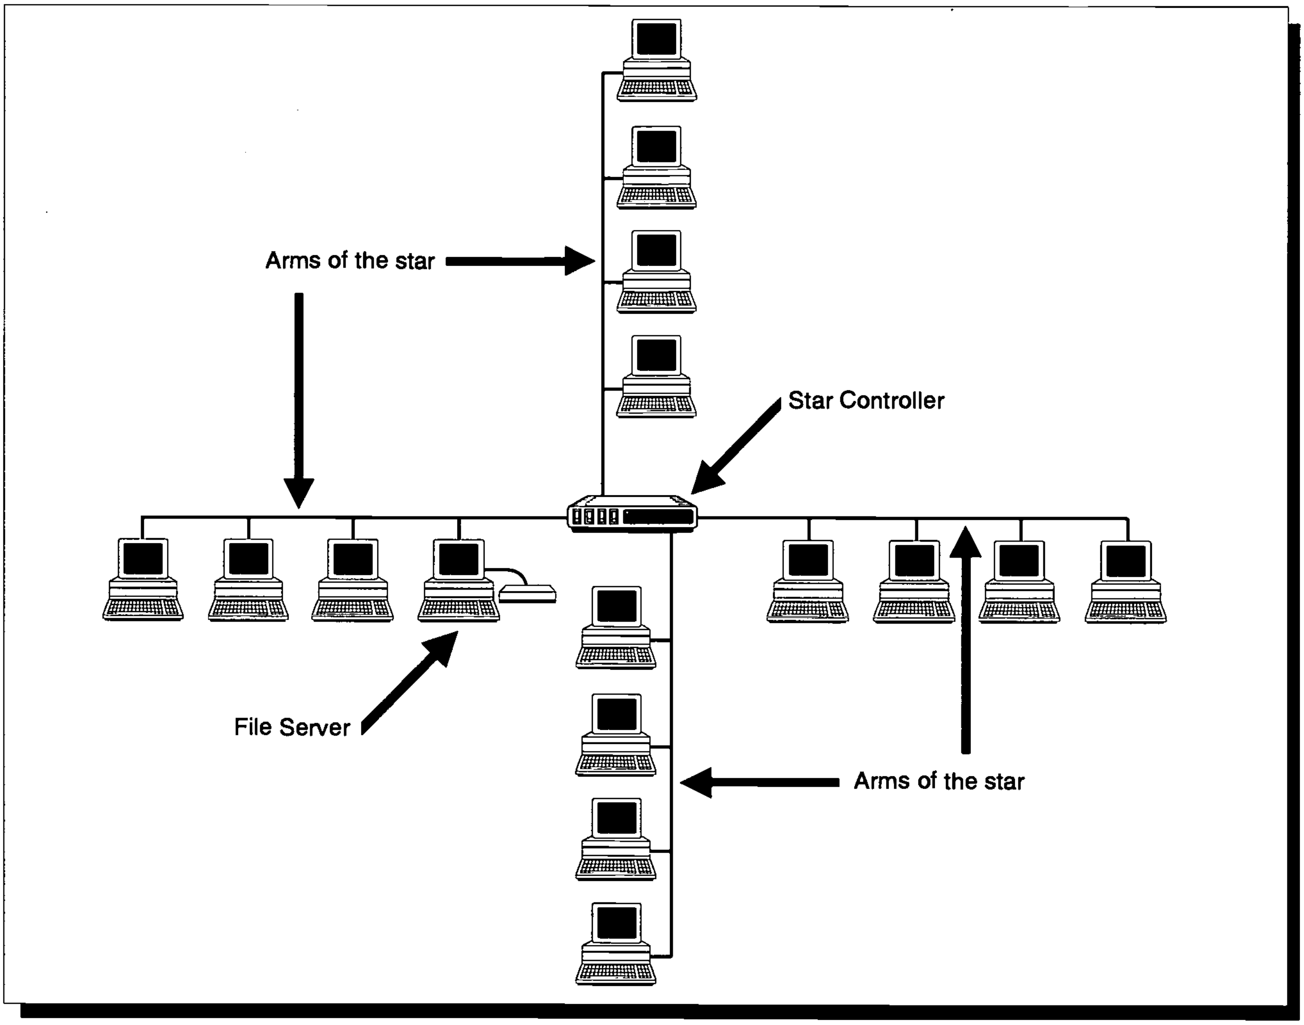

The cabling of this network is so long that signals fade before reaching their destination. By placing a star controller in the center of the cabling, a star configuration is created. Signals are boosted along each arm of the star by the star controller.

Routers/bridges

The basic concept of internetwork connections is relatively straightforward – run a cable between two or more local area networks so that all computers on all the networks can communicate with each other. The actual implementation of this idea, however, can be quite complex. Computers connected to a network operate much like telephones connected to a party line. When one person on a party line is talking, all of the others are listening and waiting their turn to talk. So, too, all the workstations on a network "listen" and wait for their turn to communicate while one computer is transmitting or receiving a message. This means that simply running a cable between two networks can greatly slow the speed of both because more computers are on the party line. Each computer has much longer to wait before getting a chance to send or receive a message.

To address this problem, hardware manufacturers developed a device called a router . A router is a box with two ports that is located between two networks. One port of the router is connected to one network and the other port is connected to the other network. Both networks must use the same communication strategy – for example, two AppleTalk networks for Macintosh computers or two Novell networks for IBM-compatible computers. Each network is given a unique identification number by the router. The router then monitors all messages being sent on both networks. If a message is being sent between two computers on the same network (a workstation loading a file from a server, for example), the router does not let that message pass to the other network's cabling. In this way, the router isolates local network communications and limits the number of computers that must wait and listen when any particular message is sent. The router only lets a message pass when it originates on one network but is addressed to a device on the other network (for example, when a user on one network prints to a laser printer on the other). By isolating local messages, routers create network zones on a larger interconnected network. Some people and some hardware manufacturers call this device a bridge. However, the distinctive feature of a router is that it maintains a table of all the network numbers that are interconnected instead of just the numbers of the two networks connected to the device, as is the case with a bridge. A bridge passes on all messages not intended for the local zone while a router passes on messages only if the target zone is in its table.

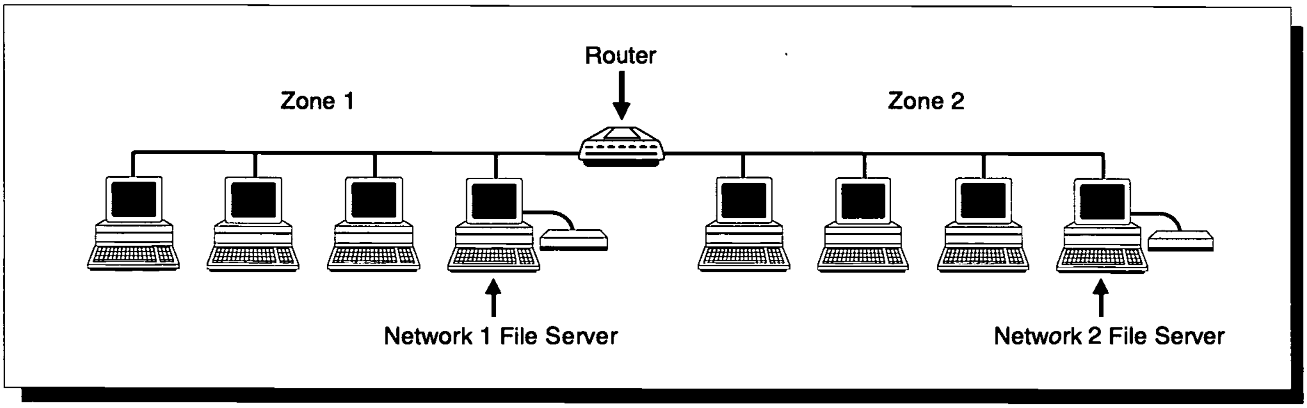

The basic rule for networks is that when one computer is sending a signal over the cabling, everyone else listens. This is similar to a party line in the telephone system. A router is a device that isolates signals that are intended for a local file server so that all computers do not have to wait while a single computer sends a signal. Here the router isolates signals intended for the servers in each zone. Only signals targeted for the server in the other zone pass by the router.

Gateways

When you want to connect two dissimilar networks, for example, an AppleTalk network for Macintosh computers and an Arcnet network for IBM-compatible computers, you need a different kind of router, known as a gateway. A gateway not only monitors messages being sent within the two networks, it also contains the instructions necessary to translate the signal so a message can pass properly from one network to the other. In this way, it is possible to connect different computer networks into a larger interconnected network even though the networks involve different makes of computers.

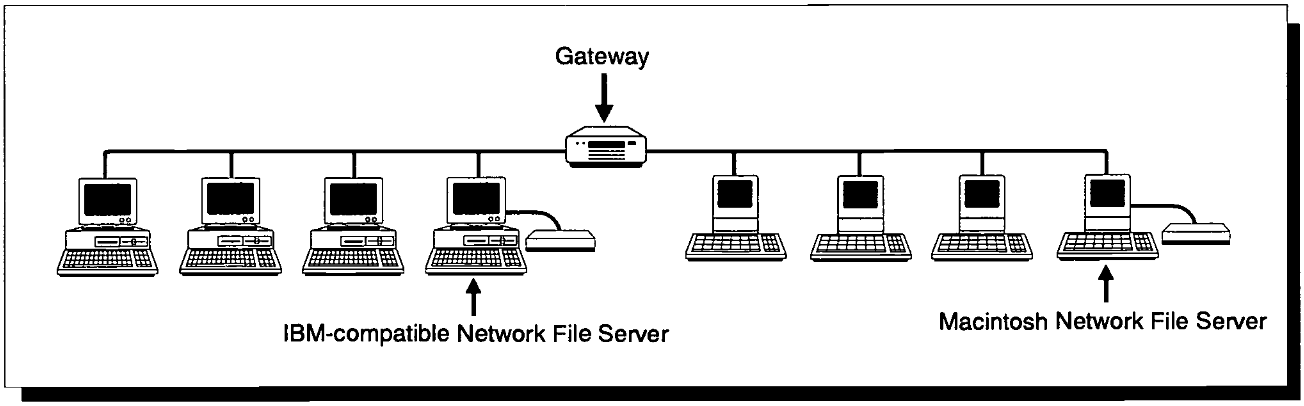

Gateways are routers with the additional capability of translating signals from one type of network to another. In this case, the gateway isolates network signals for both networks. When a Mac requests some electronic mail from the IBM-compatible file server, not only does the gateway allow the signal to pass, but it also translates the signal so the IBM-compatible file server can respond to the request. Then the gateway translates the response signal into a form the Mac can understand when the electronic mail is sent back.

The tasks that routers and gateways perform can be accomplished either by hardware or software. Hardware routers and gateways are usually small boxes into which you plug the cabling of two networks. Software routers and gateways run on a computer and the cabling of the two networks is plugged into the back of the machine. You may be able to choose whether to have the computer completely dedicated to the router/gateway tasks, or to use the router/gateway computer as a workstation. If you choose to use the computer as a workstation, then the router or gateway is said to be running in the background. We recommend that software routers and gateways be run on dedicated machines unless they are running in the background on the network server and that server is a powerful computer.

Backbones

Routers and gateways allow you to connect networks efficiently because they isolate network signals. However, the efficiency of these devices can quickly be lost if three or more networks are daisy-chained together with any combination of routers and/or gateways. A daisy-chain is created by connecting the ends of networks together so that a single cable runs from the first computer on the first network to the last computer on the last network in the chain. If a computer in the network at one end of the daisy-chain sends a message to a computer in the network at the other end of the daisy-chain, then all of the computers in each of the networks in the middle must listen and wait until that communication is finished. In this topology, a single computer's communication could cause significant delays to several networks that are not at all involved in that specific communication. It could also result in a lot of very frustrated network users. This topology is illustrated in the following diagram.

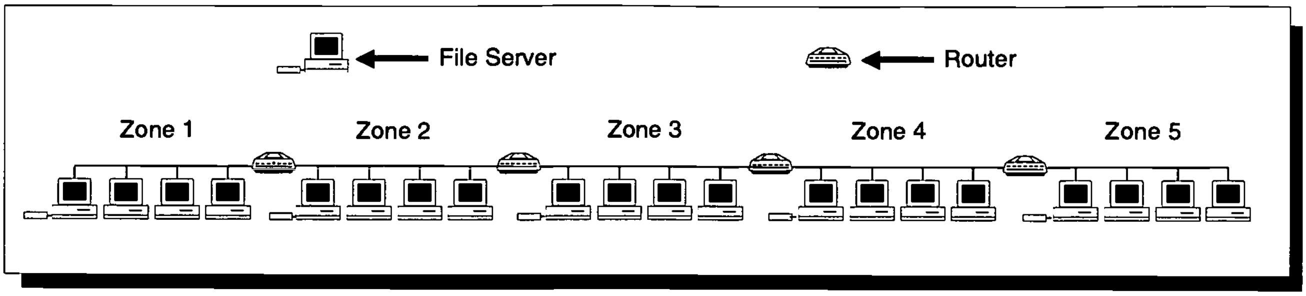

This diagram illustrates the wrong way to connect networks together. It is true that the routers will isolate the signals intended for the local server in each network, but a major problem will occur if a computer in network 1 makes a request for some information from the server in network 5. Each router will let the signal pass because it is targeted at a zone that is further down the line. That will mean that every computer on every network will stop and wait while that one communication takes place. Network users in zones 2 through 4 will not be pleased at the delays this configuration creates.

A method is needed for allowing any two networks on a large school-wide or office-wide network to communicate with each other without having that signal pass through the cabling of any of the other networks in the building. This can be accomplished by using a backbone topology for network cabling. A backbone is a cable that runs between the routers or gateways for all of the individual networks connected together, as shown in the following diagram.

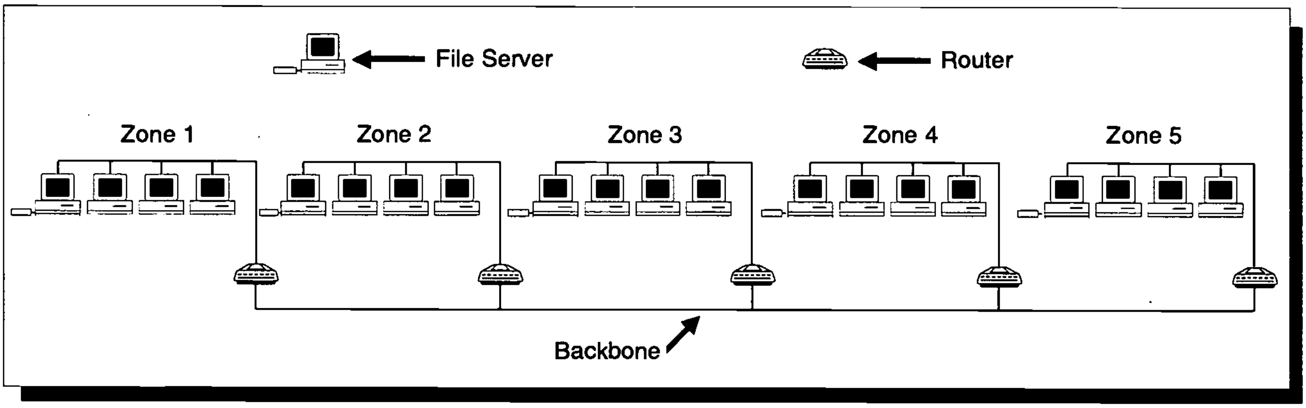

This diagram illustrates the proper way to connect a number of networks together. This configuration is called a backbone. Notice that no computers are connected to the backbone itself. Its sole purpose is to facilitate internetwork communications by allowing signals to travel along the backbone to the appropriate router.

Notice that this cabling configuration allows a message to travel between two different networks without passing through the cabling of any other network other than the two involved in the communication. Once a message passes through a router or gateway onto the backbone from the sending computer, only the router or gateway for the receiving network will let it pass the message onto another network's cabling. The other routers and/or gateways block the message.

The cabling for a backbone may run the length of a school or it may be as short as eight to ten feet. Short backbones are often located in or near the administrator's office for ease of trouble-shooting and reconfiguring the network. Having a short cable for the backbone can result in lengthy cable runs from the individual networks to the backbone. To ensure that signals do not fade, repeaters are often used in conjunction with the routers and/or gateways on the backbone. This configuration is illustrated on the next page.

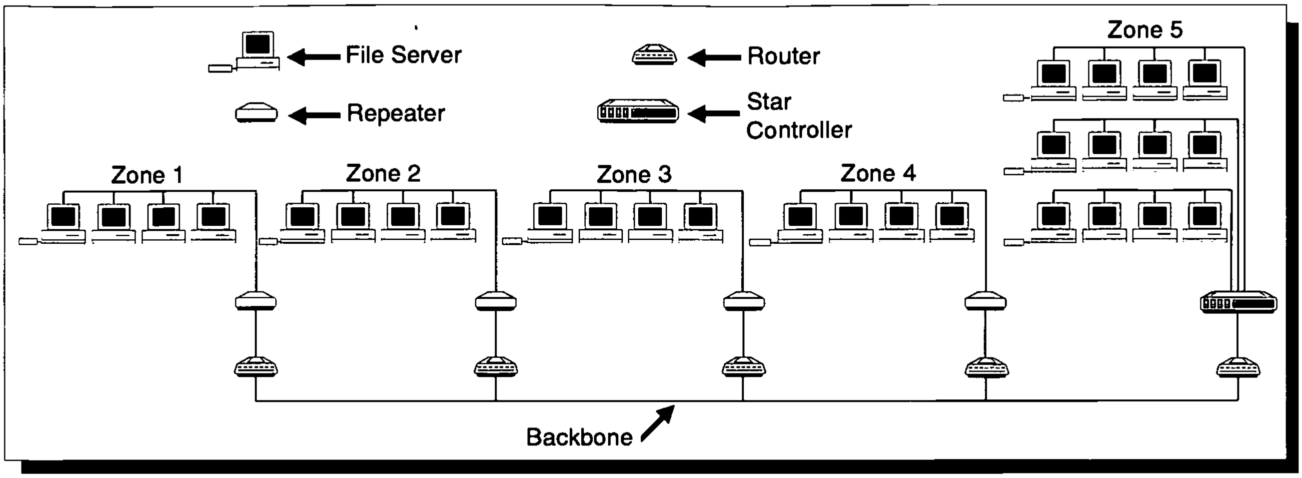

Here is a more complex connection of five networks using a backbone. The networks are not close to each other so the signals must be boosted from the backbone to the individual networks. Notice that the large network in zone 5 requires additional signal boosting via a star controller.

The speed of the signal travelling on the backbone is another consideration with this topology. A common strategy is to use cabling for the backbone that has a higher communication speed than the cabling for the individual networks connected to the backbone. This ensures the backbone does not cause unacceptable delays for internetwork communications due to the speed loss involved in passing a message through two routers/gateways and along the backbone. This configuration is commonly used when connecting several Macintosh networks together. Each individual network could use the relatively slow LocalTalk cabling system. But these LocalTalk networks would be connected to a high-speed backbone consisting of Ethernet cabling. The backbone would be so fast that requests for information from a server on another network would be met virtually as quickly as requests for information from the local server.

Remote connectivity

Remote connectivity refers to the ability to connect to a network from a remote location via the telephone system. As microcomputer networks grew in number and sophistication in the 1980s, several network developers added remote connectivity features. Two main ways exist to add this feature to a network. One way is to allow network users to access phone lines via a network modem. The second way is to allow users to log-on to a network from their computers at home (or wherever else they take their computers) using a telebridge for remote access. Let's discuss both.

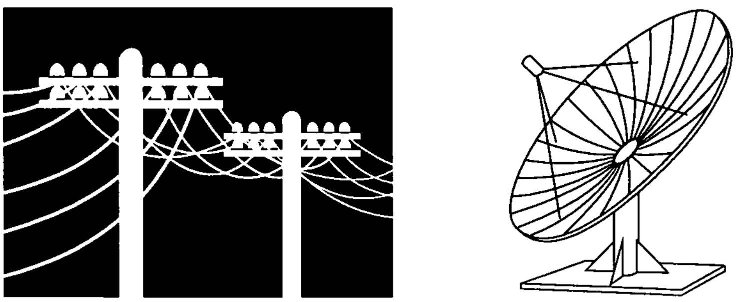

Using the phone system and satellite technology, you can connect to a network that is literally on the other side of the world.

Network modems

A modem is a device that converts the digital information stored in a computer into wave form so the information can be transmitted over a telephone line to a remotely located computer. Most modems are dedicated devices designed for the use of a stand-alone computer. However, providing modem access for all network users is too costly if a separate modem must be purchased for each workstation. To address this problem, modem manufacturers developed network modems that are designed for a multi-user environment. They are equipped with all of the instructions necessary to allow any workstation on the network to connect to the modem. The network modem could be a special piece of hardware that you purchase specifically for its network capabilities or it could be a normal stand-alone modem connected to a computer running special software that makes the modem networkable. Most network modems are connected to a single telephone line, meaning only one workstation can use the modem at a time. But the use of the device is shared among all network users. Multi-line network modems can handle several simultaneous connections to phone lines, but they are more expensive and arrangements must be made with the phone company for multiple phone lines to be run into the network modem.

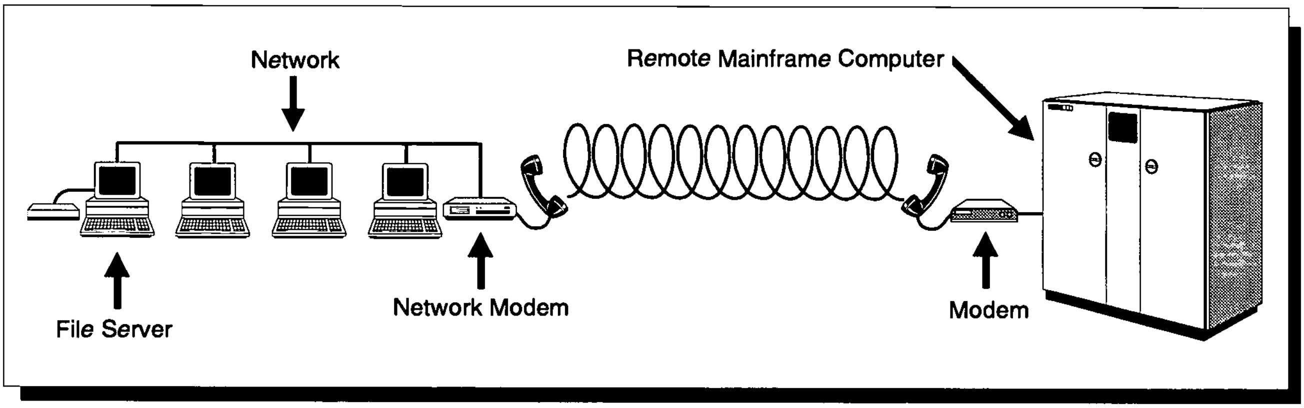

A network modem is a shared device that enables network workstations to access remotely located computers via the telephone system. Here the network modem is being used by one of the network workstations to connect to a mainframe computer to access large database files. Most microcomputer network modems are connected to a single phone line, limiting remote access to one workstation at any given time.

Telebridges and remote access

Networking has led to significant changes in the way that people do their work and in the way they communicate with each other. Using a network, a group of people can work together on the same files that are accessed in turn by each member of the group or, in some cases, by several or all members of the group simultaneously. The network also allows people to communicate with each other using electronic mail. All of this can take place without the work group members being physically together. This has added great flexibility to how work is done in modern businesses. Not long after the introduction of local area networking, users began to clamor for the ability to connect to their office networks from remote locations such as home, cars, and hotel rooms. Network hardware manufacturers responded by developing the telebridge, a device that allows a workstation to log-on to a network over a telephone line. Various telebridges are now available that can handle a single telephone connection or as many as twenty or more simultaneous remote connections to the network. The use of these devices has spawned an entirely new way to "go" to work, called "telecommuting."

The telebridge is used to log-on to the network from a remote computer. The telebridge allows the remote computer to become another node on the network; that is, the remote computer is treated exactly the same as any other workstation on the network. All information must travel back and forth between the remote workstation and the server via a modem. Although this configuration can be acceptable for accessing data files on the network, running programs is not a pleasant experience. Even high-speed modems are not fast enough to allow remote users to run programs from the server on their remote computers – it's just too slow. Thus, when logging on from a remote computer, only information files are exchanged between the remote workstation and the server. All programs are run from disk drives attached directly to the remote computer. A telebridge is illustrated below.

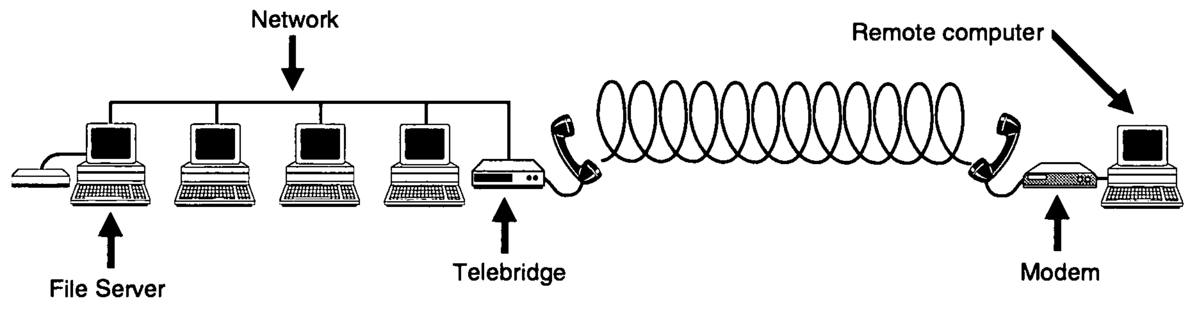

Telebridges allow people to connect to a network from remote locations. The idea is to make the remote computer act as if it were physically connected to the network so it can gain access to data on the file server. Notice that the telebridge allows the remote computer to become another node on the network via the telephone lines.

However, there is a second strategy for allowing remote users to access data and run programs from the network server at quite acceptable speeds. To make this happen, a workstation that is connected to the server by cabling is configured with both a modem and special software that allows remote users to log-on as if they were actually in the room using that workstation. The remote computer must also have special software to communicate with the network workstation via the modem. The special software simply sends screen updates over the phone lines to the remote user but all of the programs are being run on the workstation that is physically on the network. The remote user literally takes control of that machine from the remote location, sending commands via the modem. The amount of information for the commands and screen updates that must be sent over the phone lines is relatively small and so the remote user may not notice any significant loss in speed, even though they are working elsewhere. This configuration is called remote access and it is illustrated on the next page.

Businesses have been using telebridges and remote access for a number of years. Schools, on the other hand, are generally unaware of these devices and their potential. This technology paves the way for students to telecommute to school.

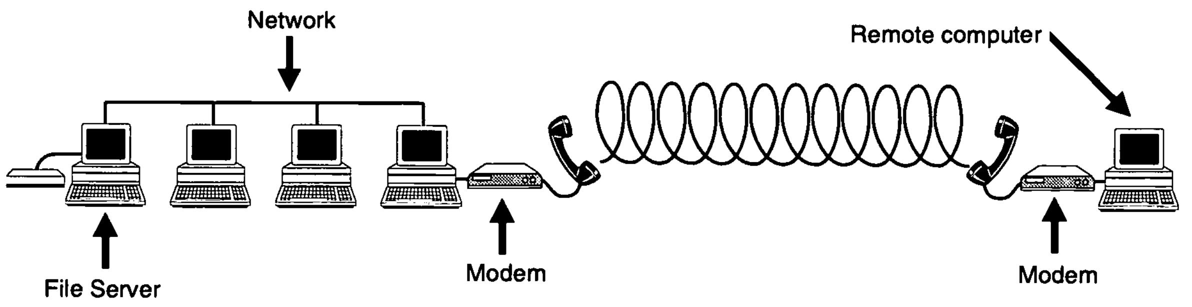

This is remote access. Here the remote computer and a network workstation use modems and special software to allow the remote computer to literally take over the control of the network workstation. All programs will run on the workstation. The remote computer will only receive screen updates via the phone connection.

A word to the novice networker

By now you should have a better idea of just what computer networking can do. If you are new to this field, then your head is probably spinning with all of this new information. Don't worry – you don't have to know all of the concepts in this chapter to get going on installing a computer network in your school. The concepts of electronic mail and conferencing, internetwork connections, and remote connectivity are not critical for creating and running classroom networks. However, it is important that you understand the concepts of print sharing, print spooling, disk sharing, and file sharing if you plan to construct a network for your classroom. Take a moment and do a mental check on those concepts right now. If any of them are a little "fuzzy," go back and reread those sections of the previous chapter before continuing on with Chapter Three.

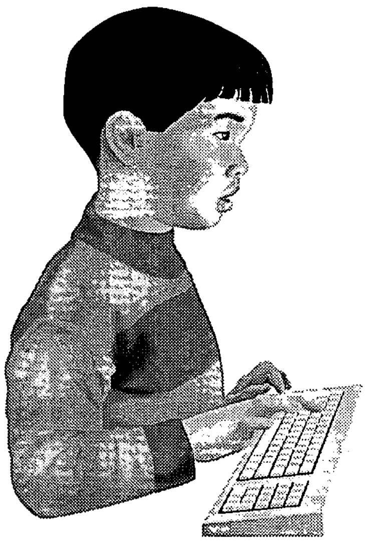

Chapter Three

Why network classroom computers?

Chapter 3

Why network classroom computers?

Topics covered in this chapter...

- *The historical development of computer use in schools*
- *Reasons for networking classroom computers*

Computers in schools

Microcomputers empower people so they can do whatever they do faster, more accurately, and with greater precision. They can even do tasks that were impossible without the computer's assistance. Modern life is dependent on computer technology. Thus, students in school today must be prepared for this electronic world. If schools are to have any hope of being relevant, they must have computers.

However, simply having computers in schools is not enough. Even using computers to do work or play games is not adequate preparation for the world students will face. The key is equipping students with the mental skills that will allow them to utilize the power of new technology to solve problems and to grasp the significance of the information that new technology provides.

Therefore, we want to state clearly at the beginning of the discussion regarding networking in schools that the focus of all technology used in education must not be on the technology, but on the reasons for using the technology. Simply reading this book and installing a network in a classroom will not be of any lasting value unless the teachers who use the network have a clear and coherent strategy for utilizing the technology for some educational goal. We wholeheartedly recommend that individual teachers and entire school staffs consider why computers should be used in schools, who should have access to the machines, and how the computers can be used to enhance student learning. The results of this discussion will go a long way toward defining the need for a network in your school and who will be connected to it.

With a clearly defined goal in mind, computers can be an effective tool for teaching. Furthermore, networks can make the job of using computers substantially easier. To get an idea of what a network can do for the classroom teacher, let's take a quick look back at the history of microcomputers in schools and identify the problems that teachers encountered when using various hardware configurations.

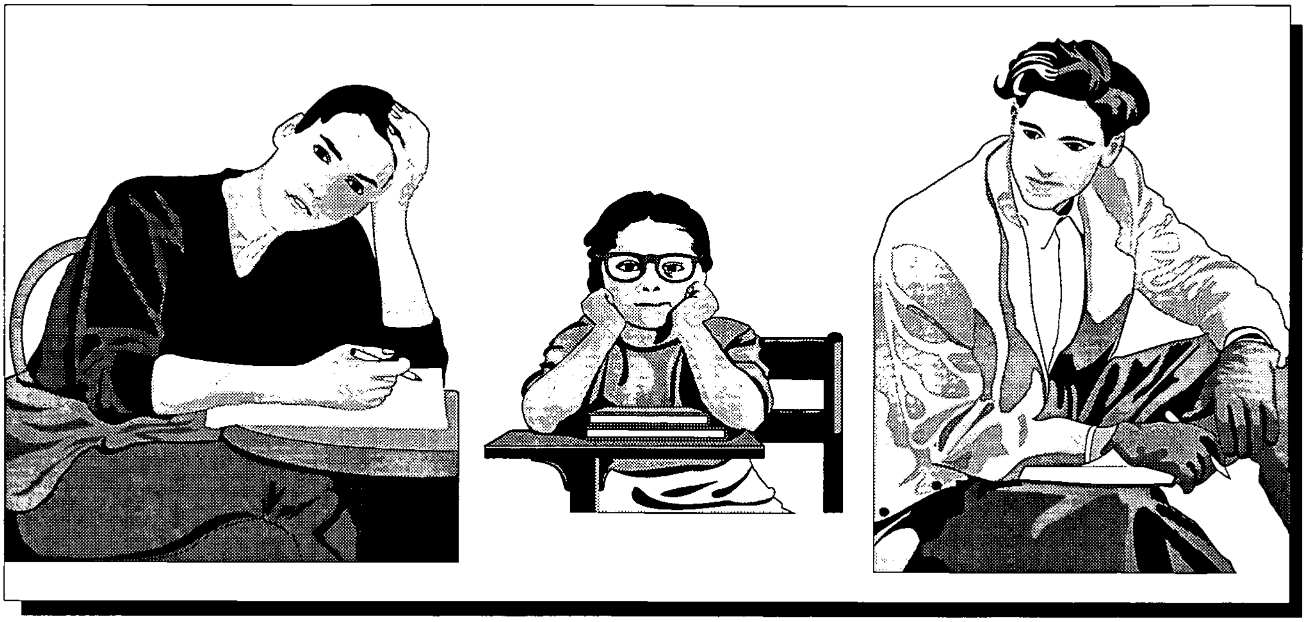

Kids have to be the focus – not technology. Without a clear understanding of how the technology will meet the needs of students and assist in the delivery of curriculum, no amount of hardware will make any significant difference in schools.

In the beginning

Many educators saw the potential for using microcomputers in classrooms shortly after personal computers hit the market in 1976. The technology was a little different in these early years than what many of you are familiar with today. The computers had relatively slow processors, very small memories, and cassette tapes for storing programs and data. The era was typified by decisions such as, "Should I buy a computer with 8K of memory or really go all out and get 16K?!!!"

There were few professionally written programs available, so you ended up writing them yourself (with varying degrees of success). Students were primarily involved in learning to program the computer in a language called BASIC. Programs and data were stored by connecting a cassette tape recorder to the computer. Most classrooms only had three or four tape recorders so they travelled around the room as required. Saving and loading information was a slow process. All in all, getting information into and out of the computer was a nightmare.

Printing was an adventure. Most classrooms had only a single printer (if they had a printer at all). The printer was dedicated to the use of a single computer. As a result, teachers developed a unique printing network system, known as "WalkNet," where a student who wanted to print something saved the file on a cassette and then walked over to the computer with the printer, loaded their file, and printed it. All in all, printing was no fun either.

As you can imagine, computer use during this period was limited to the few misguided teachers whose lives found new meaning in spending every waking hour cursing the day they ever heard the word "computer," and to the unfortunate students who ended up in their classes. You could usually spot the teachers who taught with this early microcomputer equipment by going into the staff room at lunch time – they were usually sitting all alone in the corner of the room, mumbling incoherently to themselves, and suffering a nervous twitch which caused their head to jerk to the side every few seconds. Just be thankful you don't have to deal with that technology in your classroom today.

Here is the face of a teacher who used microcomputers when they were first introduced. This teacher's face is now permanently like this – the experience of trying to use and maintain a lab of the first microcomputers has been indelibly etched on his life.

In the Middle Ages

By 1980, not only had the major personal computer manufacturers begun to produce disk drives for their computers, they actually all worked reliably! In the early to mid 1980s, second generation microcomputers equipped with faster processors, expanded memories, a floppy disk drive, and professionally written programs for business and education made computer use much more realistic for a great many more teachers and students.

You could actually get a computer to do what you wanted in this period of microcomputer development. The professionally written software tools and educational programs actually made the computer reasonably productive. Programs and data stored on five-and-one-quarter inch floppy diskettes loaded more quickly than those saved on cassette tapes. Five-and-one-quarter inch diskettes were later superseded by higher capacity three-and-one-half inch floppy disks in hard plastic shells. This newfound power led teachers to the idea of creating computer labs where entire groups of students could be taught at once.

However, while the development of computer labs for writing, or business studies, or computer assisted instruction, was a great idea, the actual day-to-day maintenance of such labs proved to be a headache. The first problem was maintaining class sets of all programs to be used in the lab. The information stored on disks is not permanent and damage can be caused by static electricity, electromagnetic field, or careless use, where the disk surface is touched by dirty hands or dropped on the floor and stepped on. Teachers found themselves continually recopying disks to keep the lab running. Teachers also found themselves having to collect the program disks at the end of a period and trying to figure out who did and didn't return their disks.

Student data disks were another major problem. The information saved on these disks was subject to all of the same hazards. In addition, students would often lose or forget their disks at home, resulting in wasted class time and possibly forcing them to redo lost work. If the teacher wanted all students to be working on the same project file, then class sets of data disks had to be maintained, subject to all of the problems already discussed.

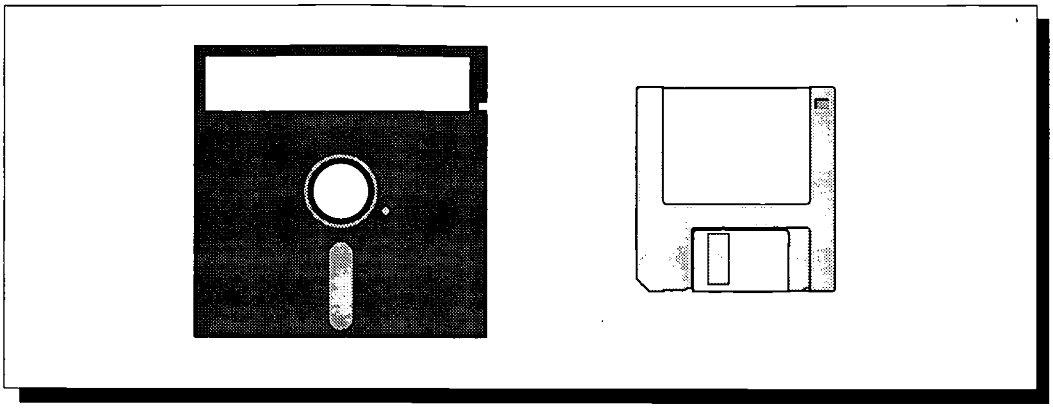

Using floppy disks was synonymous with computer use in the 1980s. Lab administrators were endlessly creating, distributing, collecting, recopying, and replacing floppy disks. Many computer labs today still use stand-alone computers equipped with floppy disk drives.

The disk drives themselves were also a major problem because they would break down or need adjusting due to continual use. Teachers in these labs soon discovered that only certain disks would work in certain drives as the disk drives slipped slowly out of the correct alignment or the correct rotation speed. Those responsible for keeping a lab running regularly found themselves either on the phone asking for a technician to come and fix a disk drive, or taking disk drives apart themselves and adjusting alignment and speed.

Printing was done primarily on a computer equipped with a dedicated printer. Thus, "WalkNet" printer networks were still the only way to get a file printed. As you can see, running a computer lab equipped with individual floppy disk drives in each computer was a lot more trouble than teachers anticipated.

The modern era of microcomputers

Today, virtually all computers come equipped with an internal hard disk drive. These drives are much faster and have substantially more capacity than the drives that use floppy disks. In addition, hard disk drives tend to be more reliable than floppy drives, requiring much less adjustment. Programs and data files are stored on the hard disk in each computer and load very quickly. Many computer labs in schools are now equipped with hard drives in each computer, so it would seem that maintaining them would be straightforward. Unfortunately, this is still not the case.

There are still management problems for the person responsible for a lab equipped with hard disk drives. Possibly the most critical problem concerns the unauthorized and illegal copying of the programs stored on the hard disks. In a lab where computers have only floppy disk drives, the teacher can collect all of the program disks and lock them away. But when the programs are stored on the hard disk inside each computer, there is no way to prevent students or others from coming into the lab and copying one of the programs, unless you turn off the power to the computers. Even if you turn off the power outside class time, students still get access to the computers during school hours. Besides, turning off the power to all the equipment just hurts all of the people who wanted to use the computers to do some legitimate, productive work. You can purchase software tools that will lock up your hard disk and prevent illegal copying, but the cost is prohibitive when you have to buy a copy for each hard disk in the lab. Furthermore, going around to each hard disk to install and configure the locking software is very time consuming. It is not a very realistic solution.

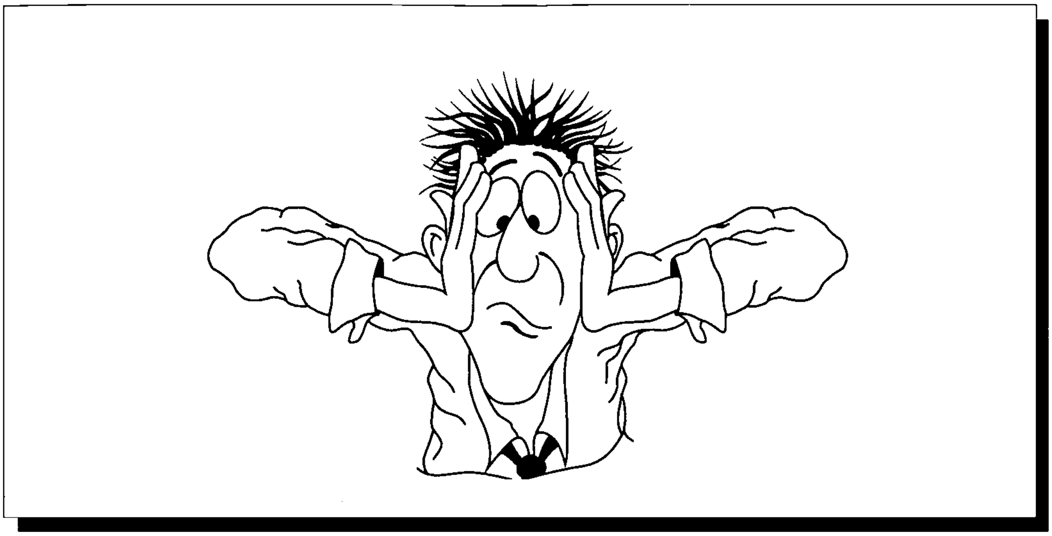

Here we see a teacher responsible for a lab of stand-alone computers equipped with hard disk drives. This person is showing the effects of trying to keep all of the hard drives running. He can't afford the expensive hard drive security software for each computer so he must try to keep up with all of the things his kids are saving on and deleting from each computer. It can't be done!

Another problem that occurs is the inadvertent reformatting of the hard disk. This can happen when a student starts the computer with a system disk in the floppy disk drive. The computer will then allow the student to reformat the hard disk without protest. In many cases this will be inadvertent, but mischievous little minds sometimes do it on purpose. However, the net effect is the same – all the information on the hard disk is lost. The person responsible for the lab now faces the rather unenviable task of recopying and reprotecting all of the programs on the hard disk. Students who had work saved on the hard disk face the prospect of redoing it.

Yet another problem with hard disks is the absolute hodge-podge of files that can end up being saved all over the place. So many files, in fact, that we have seen situations where someone has looked for over half an hour to find a file. Students who are new to the computer will often think they are saving their work on their personal floppy disk when they are actually saving it on the hard disk. We have seen students redo an entire file when they couldn't find it on their personal floppy disk, only to later find the file sitting on the hard disk. In addition, students who save their work on the hard disk, either on purpose or by accident, may find that other students have copied their files or erased them.

Shared files are also a problem. If the teacher wants all students to be working on the same file, then he or she must go around to each computer and copy the file onto each hard disk. Then the teacher has to hope that no one deletes or changes the file before all students have used it.

Finally, unless some form of print share network has been installed, hard drive-equipped computers must still use the WalkNet printing scheme. Hard disks are certainly better than the cassette tape recorders we started with, but they are no panacea for using computers in schools.

The reasons for networking classroom computers

Now that we have outlined some of the problems involved in using computer in schools, let's turn our focus to a solution – connecting classroom computers together into a network. No matter what type of computer equipment you are using, life will go much more smoothly when you install a network. Current networking software for both Macintosh and IBM and compatible computers provides many useful features that are ideally suited for education. We can see seven major reasons for installing a network in a classroom right now. Let's discuss these in detail.

Program security

Networks allow you to store copies of programs in locked, read-only spaces on the server. This allows many users to run the program but it prevents anyone from tampering with the program itself, because the network will not allow users to make any changes in a read-only area. What's more, current networking software will protect the program file from unauthorized copying. Some networks actually alter the program so that it will not run on a stand-alone computer. This means that even if someone were able to get a copy of the program from the network, it would be useless because the program simply won't run when the person tries on their own home computer.

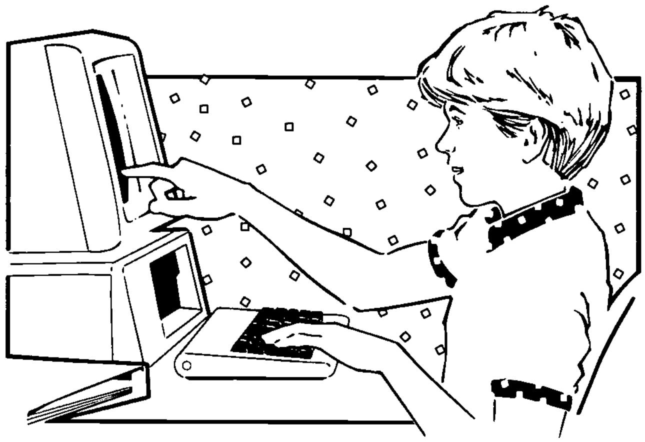

This is what we all want to see when computers are used in schools – kids using the technology to do productive work. Networks make computer use easier for both students and teachers while adding a number of new features that greatly increase the power of the computers.

Personal work files are much more secure

Students and teachers can save their personal work in private spaces on the network server's hard disk. Access to those spaces is password-protected so as long as people carefully guard their passwords, their work is safe from inadvertent or malicious copying and/or deletion. This means that all network users don't have to worry about the loss of information from damaged or lost floppy disks. Students do not have to worry about wasting class time because they forgot their floppy disk in their locker or left it at home.

Reduced need for floppy disks

Once personal work can be saved on the network, floppy disks are needed only for archiving old files. In fact, the only floppy disks that may be needed are the boot disks to start up the workstations on the network. After that, all information comes from the network. Some networks even remove the need for boot disks at the workstation by allowing computers on the network to get all of their start-up information directly from the network server. This is called remote booting, which we will explain in much greater detail in the second section of this book.

Shared files

Networking software enables the teacher to share data files as well as program files to all network users. The teacher can create a publicly accessible area on the server that will be used to store assignments, tests, clipart, and other files that students may need. The teacher only has to save a particular file once before all students can get to it. No more copying entire class sets of data disks or running around to all the hard disks in the lab just to give a single file to all students.

Shared printers

Alas, WalkNet has been superseded. No longer do students have to walk over to the computer with the printer. Networking allows printers to be shared. Even better, most current networking software has built-in print spooling capabilities that makes printing much more pleasant for students, especially when many try to print at the same time at the end of a period.

Shared resources

Not only will networks allow users to share printers, but any number of peripheral devices as well. This is particularly important in education because powerful information resources can be accessed by entire classes if the right devices can shared. CD-ROM drives, for example, can provide access to resources such as encyclopedias, historical timelines, or collections of high resolution images. Network modems can allow one or more students to access resources located outside the classroom. Using network modems, students can access card catalogues in libraries and universities, search through a multitude of databases stored on large computers all over the world, and get direct access to the wire services that newspapers, radio stations, and television stations use.

Easy maintenance

Networks provide the amazing prospect that teachers responsible for computer labs may actually once again have a personal life! Certainly, planning and installing a network requires a lot of work, but we hope this book will help you in making that a much more straightforward process. If a network has been well planned, it may continue to meet the needs of its users for several years before major changes to its configuration are required. The single major task that must be performed daily is the maintenance of boot disks for the workstations. But compared to the problems we outlined earlier in this chapter, that's a small price to pay to get rid of all the headaches associated with running a classroom full of computers.

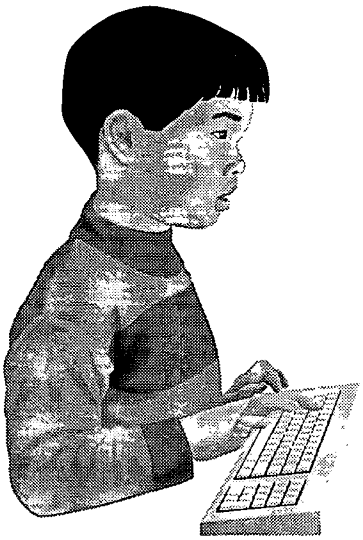

Chapter Four

Why network computers for school administration?

Chapter 4

Why network computers for school administration?

Topics covered in this chapter...

- *Administrative uses of computers*
- *The pros and cons of administrative networks*
- *Implementation strategies for administrative networks*

Computers and school administration

Computers handle large amounts of data very well. This capability led to the use of computer systems for administrative tasks at educational institutions. In the late 1960s and early 1970s, big and expensive mainframe and minicomputer systems were used by universities, colleges, and large trade schools. By the late 1970s, computer-based administrative systems were the norm for large educational institutions.

The microcomputer revolution

For the most part, elementary and secondary schools did not use computers for administrative purposes until the early 1980s (except for the few schools who sent their scheduling and/or report card information out for processing on mainframes). By then, inexpensive microcomputer technology had made it possible for K – 12 systems to afford computers for school administration. The first administrative software packages appeared on Apple II computers, with versions for IBM and Macintosh computers soon to follow. This software was designed specifically for the tasks involved in running an elementary or secondary school. While no administrative package did everything, these microcomputer-based administrative systems became very popular because they did most of the major tasks reasonably well, and certainly much more efficiently than manual methods.

The major tasks that administrative software can perform are:

- Storing student demographic information
- Scheduling students, including loading student course requests and building timetables
- Storing school-wide attendance records and generating daily attendance reports
- Storing student achievement records and generating reports for parents
- Generating reports for the Ministry or Department of Education for the province or state

This information can be gathered on paper and keyed in by secretaries, recorded on optical mark recognition cards and read by a card reader, or entered by teachers and stored on disks which are read by the main office system.

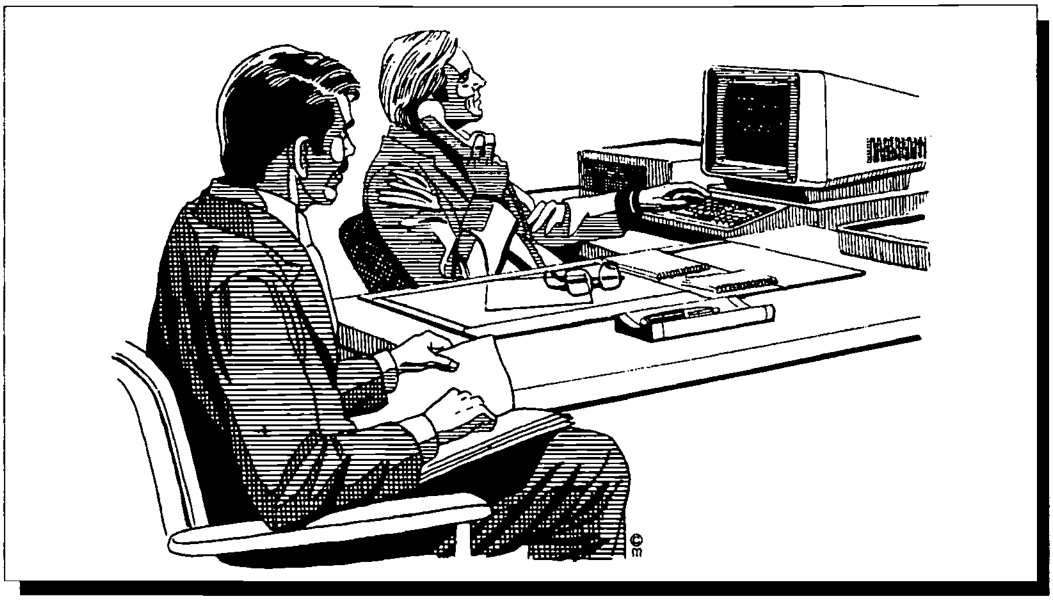

The processing speed and data storage capabilities of computers make them ideally suited to many of the administrative tasks required to keep schools running. But while computers make many administrative tasks easier, they also make the job of school administration more complex.

Initially, school administrators were the only ones who made direct use of the information stored in the computer, but it quickly became apparent that many others would benefit from direct access to this data. Counsellors, for example, often need quick access to a student's timetable, report card, or attendance information, when dealing with requests from parents or administrators. Thus, the idea of using a microcomputer-based file sharing network for administrative systems was introduced in the mid 1980s.

The pros and cons of administrative networks

The major benefits of administrative networks are:

- Many people can simultaneously view and edit student records. This gives principals, counsellors, and teachers the ability to use up-to-date information for their daily work.
- Data entry can be more efficient because it can be distributed to a number of people instead of being the responsibility of a single person. It also physically distributes the sources of data throughout the school. The burden of tedious tasks can be spread around to a number of people, thus making the job easier. Collection of report card marks would be a good example.
- The centralized storage of student information can minimize the re-entering of the same data in different places throughout the school. For example, the central office obviously needs to keep student information for contacting parents/guardians, mailing school-related material, and completing forms. But if the library wants to begin using a computer-based circulation system, much of the same information must be entered into the library computer. This is time consuming and costly. However, with a network, all computers connected to the central file server have access to all information stored there. Thus, the library circulation system can get student information it needs without requiring the re-entry of a single piece of data.

- Backing-up copies of all critical information files is centralized. The secretary responsible for data entry can make a backup copy of all information on the network simply by making a copy of the files on the server. The ability to centralize this task and build it into the daily responsibilities of a member of the clerical staff eliminates the potential for lost information due to the oversight of one or more of the network users.
- Access to productivity software for word processing, spreadsheets, or databases is shared. Like all networks, administrative networks can provide access to many different programs, while making the tasks of ensuring software security and updating to new versions much easier for the person responsible for the computers in the office.
- Access to expensive peripheral equipment such as laser printers is shared. Most schools realize that correspondence will look much better when printed on a laser printer, but the cost of several of these machines can't be justified. As a result, only one or two people have access to high quality printing, or the person connected to the laser printer is inundated with files to be printed. With a network, not only can the printer be shared by all, but the network will likely provide print spooling as well.

The major concerns of using administrative networks are:

- The security of the system. Once workstations throughout the school can access the administrative data, the central office could possibly lose control of who is viewing and editing information unless steps are taken to ensure only authorized access. Decentralized control requires that all users exercise appropriate judgment to maintain a secure system. Student access to sensitive data through remotely-located teacher workstations can be a major concern. Besides the obvious desire for students to change such information as their letter grades, poorly supervised or unsupervised remote workstations make the network vulnerable to the malicious altering or deleting of critically important data files.
- Input errors are much more difficult to prevent. A centralized system run on a single computer by a single well-trained secretary minimizes input errors. This is certainly not the case with a large network where data may be entered by a number of people, some of whom have only minimal typing and computer operation skills. A single operator also ensures that data is entered consistently. For example, while typing a person's name in all capitals is not an error, it may be inconsistent with the format of other names already entered into the system. You will be surprised at the number of different ways the same information can be entered when the administrative system is opened up to multiple users on a network.
- The cost of providing each administrator, counsellor, and teacher with a workstation, plus the cost of wiring the office and possibly the entire school for the network, can be considerable. These computers are not usually used for instructional purposes, meaning additional costs beyond those of the existing computers in the school.
- The cost of in-service training for network users can also be considerable.

Implementation strategies

There are two approaches to implementing an administrative network in a school. The first is to install an "administration only" network. By this we mean that the network will not be connected to any other instructional networks in the school. To begin with, most schools will connect these networks only to the offices of secretaries, administrators, and counsellors. As the need arises, the network may be extended to teachers, but only through specific computers designated for administrative purposes. One of the major advantages of this implementation strategy is that use is limited to a relatively small number of people. These people can be trained on the use of the computer system and on standards for data entry (on many of the networks only one or two secretaries input all of the data, while other users simply view the information). Another advantage is the security of the information on the network. If the network is limited to the secretaries, administrators, and counsellors, unauthorized student access can easily be prevented. New users can be trained on how to maintain network security. Furthermore, use of the network by new users can be phased in, when the need arises, when funding is available, and when the new users have been properly trained.

> **"Networks are more about people than about computers and cables."**

The second strategy for implementing an administrative network is to install a school-wide network on which both administrative and instructional programs will coexist. There are great advantages to this configuration. It allows student performance results to be integrated into the administrative system without re-entry. Computer-based testing results are transferred directly into the administrative system for record keeping and reporting. Attendance can be viewed and updated from anywhere in the school. Electronic mail can be extended to everyone in the school, including students, teachers, administrators, and counsellors. Booking of any and all resources in the school can be done via the computer. This strategy provides an excellent infrastructure for individualized instruction.

However, this kind of networking strategy entails a significant leap in sophistication and complexity, and requires a great deal of planning. Currently, this kind of networking is only in the early stages of development. Across North America, many innovative schools are involved in creating networks to support a more individualized, student-centered learning environment. Functional software to run integrated school-wide networks will be generally available in the next two to five years. Therefore, any long term network plans should consider the possibility of such a configuration, to ensure that the appropriate groundwork is laid ahead of time.

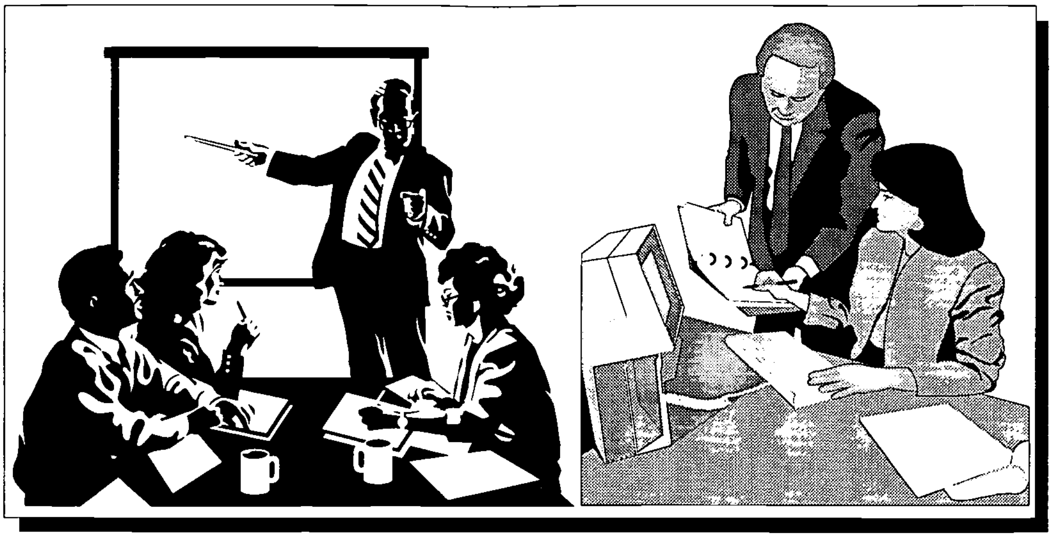

One of the most common mistakes school administrators make is to spend considerable amounts of money on the hardware and software for computerized administration packages, but not to budget for staff training. Training is the key component in the successful implementation of computers in school offices.

Regardless of the strategy used, you must remember that networks are more about people than about computers and cables. The success of any network is dependent on the proper training of its users. This is especially the case with administrative networks because the information stored there is absolutely critical to running the school. Therefore, adequate training must be provided for all those who will be using the network. Secretarial staff who will be spending a great deal of time, if not all of their time, using the network must be especially well-trained. This involves much more than a couple of one- or two-hour training sessions to get them going. Proper training should include an on-going plan for upgrading computer skills.

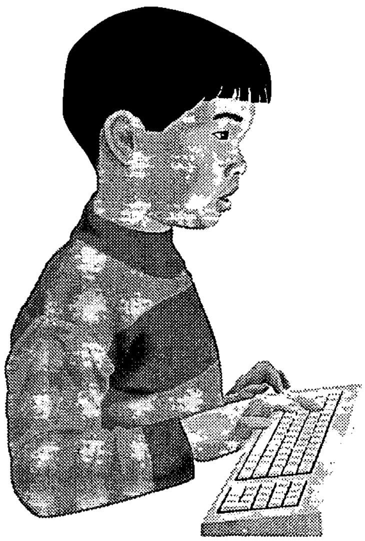

Chapter Five

Why consider school-wide networking

Chapter 5

Why consider school-wide networking?

Topics covered in this chapter...

- *Why school-wide networks?*
- *Uses for a school-wide network*

Networks in schools

In the mid 1980s, networks of microcomputers began to appear in schools. At first, networks were most likely to be located in the computer science and business education classrooms in secondary schools, and later in the writing classrooms in elementary schools. Teachers of these computer-based courses welcomed the installation of the network because it provided an efficient way to get programs and information to students without having to maintain class sets of floppy disks. After meeting with immediate success in specialty labs, networking soon spread to the school office, the library, and to other clusters of computers throughout the school. Today, a significant number, if not the majority, of computers purchased for education are destined for some sort of network.

The growth in the number of computers in schools combined with the growth in the popularity of networks has led to the idea of connecting all computers in a school to the same school-wide network. This configuration brings many new possibilities to the instructional and administrative use of computers in education. The hardware and software for such networking have both increased in sophistication, while at the same time decreasing in price to the point of making such a scheme a realistic option. The major purpose of this book is to equip you with the conceptual and practical skills to plan and implement school-wide networking in your school or school district.

Why school-wide networks?

A school-wide network can be defined as a network connecting computers throughout the school, providing communication between those computers, and giving each computer access to all of the resources on the network. There are three main reasons for using a network throughout a school. First, stand-alone computers are limited in the resources they can provide, especially information resources - they do not fulfil the promise of the technology. Second, networks are designed to provide security lacking in stand-alone computers. Third, the maintenance of a network is easier and takes less time than the maintenance of a large number of individual hard disk drives or floppy disk drives.

Access for everyone - any time, anywhere

A computer is much more than a tool for productivity tasks such as word processing - it is a point of access to information resources both inside and outside the school, many of which are not available or easily accessed by stand-alone computers. These resources might include the library catalogue, a selection of printers around the school, electronic mail inside and outside the school, databases located inside or outside the school (or country), information on CD-ROMs (such as encyclopedias), and a wide variety of productivity and educational programs. The goal of computer use in schools today should be to have each computer in a school capable of accessing many or all of the information resources in the school, and capable of communicating with all other computers in the school.

We already live in an information age, and the rapid growth in the amount of information we have experienced in the last thirty years will continue into the future at an astounding rate. Having access to up-to-date information is a key to success in the modern world. The use of a school-wide network creates a powerful infrastructure for bringing current information sources to students whenever needed in their classrooms or from home.

The implications of school-wide access to such resources are significant. Every computer in the school becomes an interdisciplinary device. With this kind of network, students can easily jump between information resources that have been compartmentalized in traditional schools, where students travel from one discipline in one classroom to another discipline in another classroom. With students accessing up-to-date information in a rapidly changing world, the role of the teacher changes from being the source of information to being one who equips students with the skills necessary to find information when they need it and helps students see the significance of the information they obtain. School-wide networks allow students to do much of their work when and where it is appropriate. For example, if a student has some free time in a Math class, he or she can easily work on their English essay because the computer in the Math classroom can access the school-wide network where the student's work is saved, and where all of the resources needed for the essay are stored as well. Later, if the student requires assistance from an English teacher, the student can walk over to the English classroom, access the essay from the computer there, and get the help that is needed.

Combine a school-wide network with the ability to connect to the network from home or any other remote location, throw in portable notebook and hand-held computers, and you have a very powerful infrastructure that can enable students to do their work any time, anywhere. Then provide an electronic mail feature, and you see that teach-

ers and students could be communicating about school work and students could be completing assignments when it is not possible for the teacher and student to be in the same room at the same time. This does not mean there will no longer be any social interaction among people, but it does mean that students could continue their education even when they are sick, injured, or away on an extended family vacation. Teachers and students also have a way to communicate when their schedules conflict.

Security and control

As soon as you consider a computer as a point of access to information resources, you must also consider the security and control of those resources. Security is the protection of programs and data from loss due to hardware failure or vandalism, and from the viewing, copying or changing of files by unauthorized people. It also includes the protection of copyrighted programs from illegal copying, commonly called software piracy. All of these issues are easier to address if the information is stored on one central file server than if it is stored on many individual hard disks, because the network software is designed to prevent inadvertent or malicious changes to important files.

Ease of maintenance

Networks provide the best solution to keeping the computers in a school "up and running." A network replaces hard disks in each computer with a larger hard disk in a file server. Other peripherals, such as printers and modems, can also be shared through the network. Consider a small school with thirty computers. Since the programs and information are all stored on one hard disk instead of thirty, they are easier to maintain. To update a program or database requires changing it on one hard disk, not thirty; to allow access to a modem requires only one phone line, not thirty; to increase storage space requires buying only one additional hard disk, not thirty. By reducing duplication a network can reduce the time and effort required to maintain the computers in a school.

Uses for a school-wide network

Networks for student learning

The variety of programs available for student use is growing rapidly. This includes many productivity programs such as word processors, desktop publishing programs, spreadsheets and databases. It also includes a multitude of educational resources such as simulation programs; Computer Assisted Instruction (CAI) programs where the computer presents curriculum material; Computer Managed Instruction (CMI) programs where the computer manages, schedules, and monitors learning activities; and learning games. Storing all that software on every machine in a school would not be practical. If there are a few hundred megabytes of data (a few hundred million characters) it would be prohibitive to have a disk drive that large on every computer, but it could be provided on only one server in a network, or a few servers connected to a school-wide network.

In some schools the variety of educational software in use is very limited: a word processor, a few favourite instructional packages, and possibly a few learning games. No one should expect that limited use to remain for long. Hardware and software capabilities are improving rapidly, and education is increasingly seen as a lucrative marketplace by developers. The quality and variety of educational software will increase steadily.

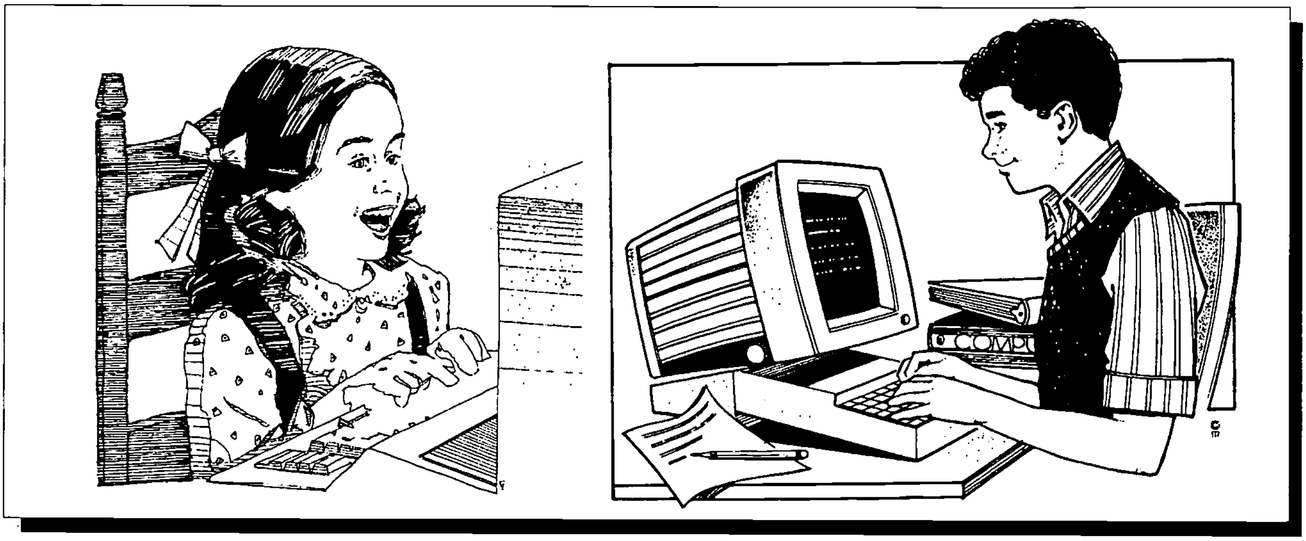

Computers are much more than sources of information. They are capable of presenting instructional material in exciting new ways that facilitate student learning and inter-curricular study.

At the same time, the drive for more individualization of student learning will make for an easier fit between computer use and the educational system. In fact, many educators believe that the only way schools can provide individualization is through technology. This requires extremely flexible computer facilities, which is exactly what school-wide networks provide.

Network access to the virtual world

School-wide networks make it possible for students to travel somewhere in the school as well as access remotely-located information resources without ever leaving their chairs. This has been termed "virtual travel" and is just one example of the virtual world that computers are opening up. Students will be able to travel down virtual hallways to resources in other parts of the school. But school-wide networking does not stop here. In time, network links will be made with other schools in the district. Then links will be made among these district-wide networks, and ultimately among country-wide and world-wide networks.

Sound like Star Trek? Think again. Many of these links are already in place and world-wide networks already exist. Many schools have installed networks that allow students to communicate with people around the world. In these schools, the students can participate in projects and discussions with students around the world and with scientists, writers, engineers, and virtually anyone else they wish to contact. In some states and in the province of British Columbia, plans have been developed (but not yet implemented) to extend such connections to every school in the state or province. The student can begin to experience the world through the network in ways otherwise impossible.

Thus a school-wide network can be viewed as the first step in creating an infrastructure which allows students to experience the world electronically - to go on virtual field trips, "being there" without really being there. Such language may be overly dramatic in describing the full screens of text we see in present communications, but other types of "virtual worlds" are also being developed through computers. In schools today we are adding still graphics, low quality video, and sound to encyclopedias. Within the next one to three years, this will improve to include full motion video

We must ensure that we are preparing students for their future, not our past. The future that children in school today will experience will be significantly different than the one their parents and teachers experienced. Technology will play an increasingly important part in people's lives. Using world-wide networks, it will be possible to interact with others around the globe. Implementing school-wide networks is the first step in making this a reality for kids in school.

in multi-media databases accessed from CD-ROM's or hard disks. The next step will be to supply this full motion video over the network. This will require faster networks than commonly used today, but those networks will be available and they will communicate with present, slower networks. Making use of high speed fiber optic cables within buildings and between buildings, these future networks will be a natural extension of the networks we are designing today.

A further extension of such virtual worlds is a computer which controls the visual and aural environment of the student, simulating a reality that exists only in the programs of the computer. This "virtual reality" can be presented to the student through a screen and sound source combined in a helmet and can be controlled by the student with input devices such as a special touch-sensitive glove. Through the helmet students see and hear exactly what they would if they were, for example, in the Louvre, and with the glove they determine where they look and travel throughout the galleries. While the possibilities for abuse of such technology could be disturbing, the possibilities for education are equally wonderful. Technology with such capabilities is

not far away, and will probably affect students presently in our school system. While today's networks are only in the early stages of providing this capability, they can be considered as the first step toward such virtual worlds.

Networks for school administration

As we have already discussed, the use of computers for school administration is becoming increasingly important. It is useful to note here that school-wide networks can be of great assistance for school administration. In fact, school-wide networks will allow schools to reorganize how instruction is delivered because the network will facilitate new ways of monitoring and reporting on student performance. By maintaining a database containing all of the records for all of the students in the school, the network can allow a teacher to view and update student performance information from any workstation in the building. This will allow schools to move away from the traditional whole-class approach to learning, to a more individualized system where complete learner profiles are maintained on school-wide networks.

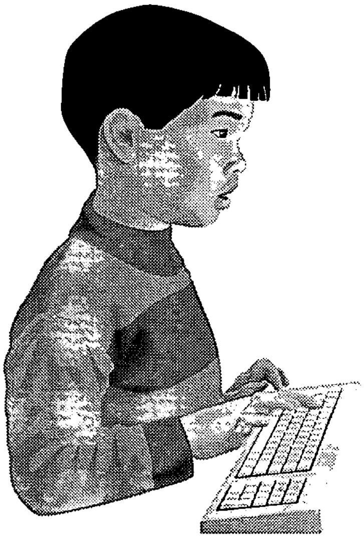

Chapter Six

How to plan a school-wide network

Chapter 6

How to plan a school-wide network

Topics covered in this chapter...

- *Successful school-wide networks*
- *Educational issues in school-wide networking*
- *Defining the needs of the users*
- *Required resources*

Keys to successful school-wide networks

The first key to successful use of computer networks is planning, and as the venture becomes more complex the planning becomes even more critical. A school may get by with a poorly planned network in one lab, but the faults become glaringly apparent when the network is expanded throughout the school.

A plan for networking a school has two parts. There should be a long-term (3 - 5 year) vision for the school as a whole as well as a short-term (1 - 2 year) plan for a particular network or lab within the school. The short-term plan is necessary to fill immediate needs, but the results should also fit into the long-term vision for the whole school.

Involve the school staff

The design of a school-wide network should have considerable input from the entire school staff. It is not acceptable for the computing facilities within a school to be designed by a single keen teacher, a school administrator, or (worst of all) a person from outside the school.

Someone from outside the school usually assists the staff in the planning process. That person may be a district computer coordinator or a technical consultant. The consultant can actually write the plan and take care of most of the technical details, but must remember that staff involvement is critical to success.

Planning will be much easier if the staff has a realistic vision. Thus it is worth added effort educating the staff, bringing them to see the possibilities as well as the limitations of networking, and working towards consensus in their vision. The school staff's input into the planning process and decision-making is only as valuable as their knowledge permits.

Planning a network is largely a matter of determining the needs of the users. The goal is to let the educators make the educational decisions, while the consultant ensures they are not confronted by excessive technical jargon. Some of the educational issues to be considered are addressed in this chapter, while later chapters deal with the technical issues.

Many schools implementing school-wide networks make the mistake of focusing on the technology. Networks affect many people and to be successful, all of those affected must be included in the planning. In a school, this definitely means involving the teaching staff. It may take longer, but the eventual success of the network will justify the effort.

Developing a long-term vision

The first requirement for planning a school-wide network is a future orientation. The school staff must envision the use of computers within the school at some point in the future; at least three years and more if possible. The vision should be written on paper, and then should be reviewed and updated annually.

That vision should have a far-reaching and global perspective, not limited by present difficulties: technical, financial, political or administrative. Nor should the vision concentrate solely on the present network. The long-term vision should extend beyond any particular technology, funding crisis or political obstruction.

The focus should be on the job to be done, not the technology. Even experts cannot forecast what computers will look like in five years, but most educators can discuss what they would like to be able to do in their classrooms and what type of information flow that would require. This stage of planning should be largely divorced from the technology, setting a general direction that the technology should serve.

Long-term planning deals with possibilities, not actualities. You are planning your network so that it can adapt to future developments in hardware and software, and to changes in the needs of the users. This means planning for flexibility.

Planning for flexibility

A long-term plan should be capable of growing to accommodate the long-term vision of the staff and flexible enough to adapt to changes in the technical and educational environments. Planning to allow future flexibility while meeting present needs is difficult, as it involves predicting the future. An experienced planner will realize that predictions of change, when provided by school staff, are often conservative and sometimes misdirected, and will attempt to provide even more flexibility than seems necessary. Thus the planner will interview staff members and include their visions, but be highly skeptical of statements like "Oh, we will never..."

While some limitations on flexibility are caused by the need for cost-effective solutions, other limitations are caused by poor educational and technical decisions. Predictions of general trends in technology and pricing should be considered in planning, even though such predictions are difficult to make. For instance, if a solution that appears optimal is new and expensive but likely to be much less expensive in the future, using a less capable, less expensive solution may be sensible for the near term with the expectation of replacing it in the future. On the other hand, spending a bit more today may open up possibilities in the future - possibilities that can be predicted and planned for. Planning a school-wide network is not a simple task, but the rewards of careful planning can be great for both students and staff.

Short-term planning

With a long-term vision for the school in hand, compatible short-term plans can be made to meet present needs. For instance, the long-term vision may include computers in every classroom networked together with a computer lab. It may also include telephones in every classroom, voice mail as well as electronic mail, and various connections outside the school. The short-term plan may be to install a single computer lab, but the vision for the future will influence the way that lab is designed.

Much of the planning and many of the decisions in a school are made by a small part of the staff, such as a department, a grade level, or a single administrator or teacher. These individual and departmental plans must be compatible with the long-term vision for the school. This should be possible, if the whole staff has been involved in developing that vision.

Use the work of others

Unfortunately, a vision cannot be completely developed until you know what is possible. Planning is an incremental process: meeting one need, thereby allowing new possibilities, creating another need. The trick is being able to prepare for a few of these cycles of trial and error. (The person who can truly see into the future of either education or technology has more lucrative pursuits than planning school networks - the lecture circuit awaits!)

To remove one's network planning at least partially from the realm of trial and error (it is never possible to do so completely), take advantage of those who have gone before. There is likely to be considerable expertise in schools in your area, and many examples of what to do and what not to do. This research is seldom published or even documented, but the time and expense of finding and listening to those who have already installed school-wide networks will amply be repaid.

Plan for the support required for school-wide networks

Networks are never trouble-free; the hardware does break down and network users always require support to gain maximum benefit. There are three types of support: educational support for teachers and students, administrative support for office staff and administrators, and technical support for hardware and software. Each must be planned. These three types of support will be considered along with other issues in the following chapters.

Educational issues in school-wide networking

It is all too easy to concentrate on the technical issues while planning computer facilities in a school, but it is the educational issues which will determine how, and how well, the computers are used. It does not matter what type of computers are purchased or what type of network cabling is installed if the system does not meet the

needs of the teachers and students. In other words, curriculum and administration first - computers, software and networking second.

The first educational decisions have to do with how the educational process will occur in the school, and how technology can assist that process. Other decisions which must be made by the school staff involve a definition (or redefinition) of roles. The staff must also decide what resources and support they will need to make effective use of the technology, although outside assistance may be needed to help make these decisions.

Technology and the educational process

Computer technology is making it possible for the educational process to happen in new ways. Computers will greatly alter the traditional practices that teachers and administrators have used for many years, just as computer technology is radically changing processes throughout society. The ways in which administrators manage, teachers teach and students learn will define their needs for technology and will also be shaped by what technology offers.

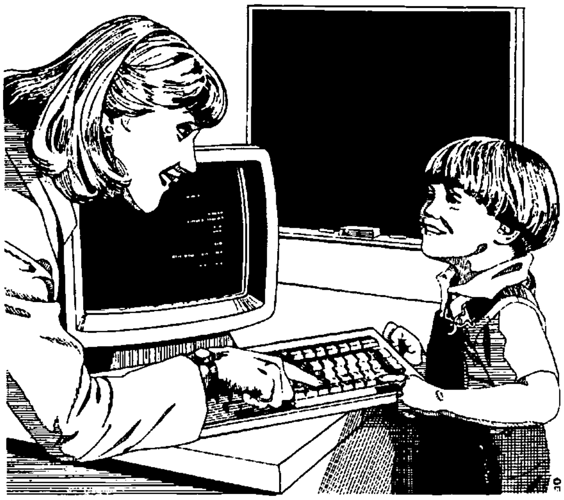

Use of computer technology is a fact of life in the modern world. Equipping school children with computer skills is an important part of their education. But computers have the potential to do significantly more for education through their ability to create challenging new interactive learning environments. This potential is greatly increased when individual computers are connected to a school-wide network.

The learning process

Technology is not free of biases, and neither are computer networks. Networks favor some learning and teaching styles over others. In particular, networks have a bias against the traditional classroom in which the door of the classroom is closed, the only information available is within the classroom, and the teacher is in complete control of that information. Although a network can be used in such an environment, the fit is not natural.

School organization

Networks may also have hidden biases in the area of school organization, or at least may bring the staff's vision within the realm of possibility. Whole classes working on the same task, small groups of students working cooperatively, students working on individual projects, students working with teachers, students working without teach-

ers, students working at home, students working in local businesses, local businesses working in school: these are just some of the possible arrangements that a computer network could support. A school staff may have a vision of a school where the organization of learning resources is a radical departure from the current organization of the school. If not, they should at least be made aware of the possibility of such alternative arrangements. Wherever possible, the design of the computer network should have the flexibility to support an evolving school organization.

A long-range plan which envisions computers being used only within today's structure will likely be too limited. In developing a vision of the future, the staff should be encouraged to consider technologies only now being developed and the opportunities they will offer. They should also be encouraged to consider the ways they would like their school to evolve in the future.

Defining the needs of the users

If a network is to meet the needs of its intended users, it is critical that the network planner has accurately defined the needs of those users: students, teachers and administrators.

Defining student needs

How will students be using computers? First, if you believe that students will only use computers in specific computer-oriented classes, you have no immediate need for a school-wide network and should put this book on the shelf for now. If students will be using computers throughout the curriculum, what will that include? Students who only use computers as tools to create products will have quite different needs from those who also use computers to present curricular material as well as find and process information. If they also use computers to manage their time and resources and to communicate, their needs will be different still. In our opinion, students will require access to computing resources in every room in the school (except the cafeteria and washrooms - for the present). Within those rooms, computer use will vary from an infrequently used add-on to an integrated part of the learning process. In some rooms one computer may suffice, while in others there will be one computer for each student. Determining where each room in the school fits in that broad spectrum is the task of the school staff.

If students also need access to the school's resources from their homes or from their own computers which they bring to school, the plans have one more requirement. The network will have to be able to connect to their computers and enable them to expand their learning beyond the school walls.

Not all computers will necessarily meet all student needs. There may be some computers used for general purpose tasks while others are dedicated to specific tasks, such as music, art, video production, or the exploration of virtual worlds.

One of the decisions which the school will have to make is where students will store their information. In some cases, students need to store their work on file servers. In other cases they can store their work on floppy diskettes. While the floppy diskette approach is less convenient for the students and may cause problems if students forget or lose their diskettes or if the diskettes are damaged, it is more convenient for the network administrator. The servers require less disk space, there is less work involved in monitoring the use of the network, and there is less traffic over the network. This decision will depend on the age of the students and what they are

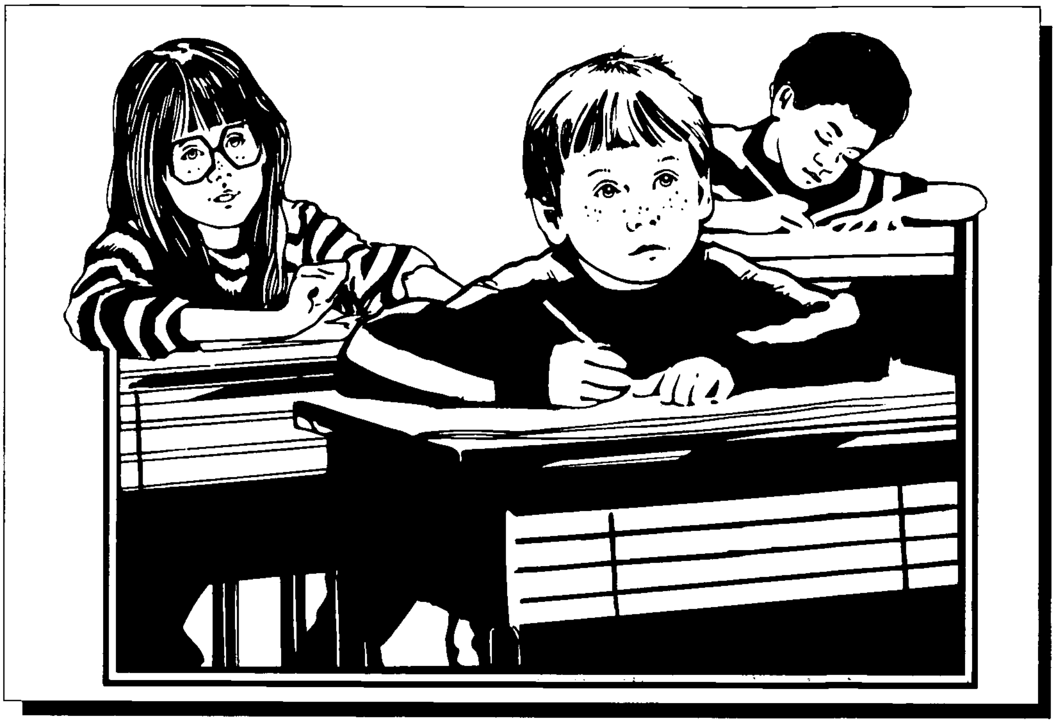

The success of any network depends on careful planning. Considering the needs of network users is the key. In a school, students will use the network more than any other group so the network must be configured to facilitate student use of network workstations.

doing with the computers. While technical factors need to be considered in such a decision, the issue is largely an educational one – how will information be handled within the school and who is responsible for it?

Defining teacher needs

How will teachers be using computers? The obvious uses are as administrative and production tools in areas such as word processing, desktop publishing, electronic mail, conferencing, finding information, planning, evaluation and reporting. The need is ubiquitous; teachers always need access to available resources .

They need access from every room they teach and work in. They probably also need access from their homes. In some situations they can share facilities with students; in others they must have their own machines. In some situations they will purchase their own computers; in others the responsibility will be the school's. Again there are many decisions which must involve the school staff.

The definition of roles within the school

A number of people must be involved to meet the needs of the people in a school. Confusion, duplication of effort, and even hurt feelings can be avoided by defining the roles those people will play. Will only a few teachers be involved with students using computers or will all staff have technology integrated into their curricula? Will staff use computers in their classroom or take their classes to a computer lab, or both? Who is responsible for the security of student work? What are the roles of parents, local businesses and other educational institutions? These and other similar questions should be considered as part of the planning process. In addition, there are some specific roles which will be affected.

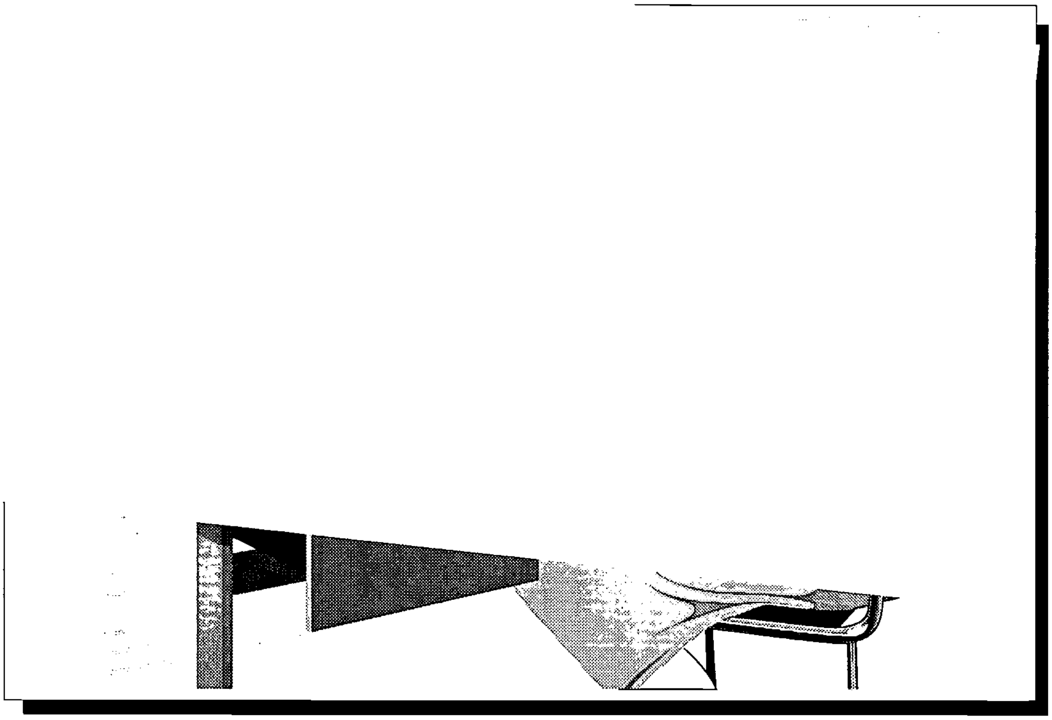

The other major user group will be teachers. Properly configured, a school-wide network can make a teacher's job easier. However, an ill-conceived network plan can make teaching a nightmare. Consequently, a careful consideration of teacher needs, including computer familiarization and training on network use, is critical to the success of the network.

The danger of the resident superstar

Historically, the implementation of computer technology in schools has often depended on keen, enthusiastic teachers willing to invest volunteer time and labour. Yet even though the contributions of such individuals must be appreciated, technology use within the school must not be overly dependent on one person, the "resident superstar."

One teacher should not be teaching computers to all students while the other teachers are not learning to do the same. One person should not install a complex network in a school that no one else in the district could support. In either case, the school will be in trouble when that person leaves.

The keen teachers must be persuaded to bring their colleagues along and encourage them to learn about technology. If a complex network is going to be installed, a backup source of support must be available, if not at the school then at least in the district. A school should not get in the habit of relying on a single person putting many extra hours into supporting computer systems, as that person will eventually burn out or move away, leaving the school and the students in a difficult situation.

Roles within the school must be defined to reduce dependence on any one individual, but this must be done in a way that does not discourage the keen teachers, or put roadblocks in their path. Sufficient time must be set aside to allow these teachers to share their expertise with their colleagues. The school system has huge amounts of untapped expertise within its ranks, and all teachers need opportunity and encouragement to be lifelong learners.

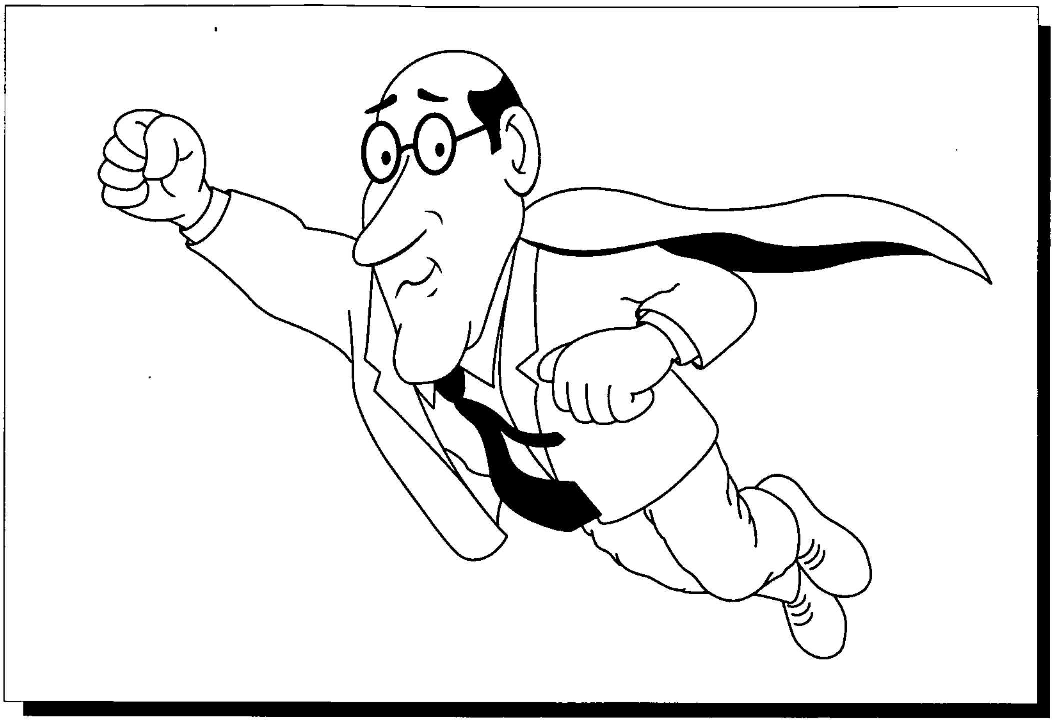

It's a bird, it's a plane, no, it's the resident computer superstar! We know it sounds funny, but many schools are too dependent on the computer skills of a single person. This scenario must be avoided, to prevent burn-out of the teacher with computer expertise and to prevent the collapse of the technology program if that person becomes ill or moves on to another job.

The role of the library and librarian

The role within the school that may be the most affected by technology is that of the librarian. What will the role of the librarian be in 3 - 5 years? What forms will information take in the library and in the school? Will the library be one isolated room or will it be a resource available throughout the school? The answers to these questions indicate a rapidly changing role for school librarians. If any educator is undergoing a "paradigm shift," it must be the librarian.

Computer technology opens up new possibilities for storing and accessing information, and the librarian who thinks only of paper is on the path toward extinction. Even the management of books, which will hopefully always be with us, will change when the resource catalogue is not a series of cards but rather a searchable database accessible from every room in the school. This is reality in an expanding number of schools, and the extension of the library catalogue to students' homes is not far away. Since both librarians and computer networks are managers of information and facilitators of its use, an analysis of the school librarian's role is a critical part of planning for a school-wide network.

The role of the network administrator

Every network needs one or more administrators. The person(s) responsible for this role must be designated. Their tasks include adding software, managing student accounts, helping teachers solve small problems, and answering questions. If a technical support person is available, this job might best fit into that role; or it may better rest with a teacher given the position of computer resource teacher. This decision must be made at the school level, as there do not seem to be two schools in the world with the same definition of these roles, but having one individual with final responsibility for the operation of the network seems to work better.

The role of the computer resource teacher

While the administrative work may be done by technical support staff, there is still a need for a school "computer resource teacher" to help educators make effective use of the computers. This person should be heavily involved in staff development and in-service activities, as well as in helping individual teachers find the best ways of using technology to improve their instruction. Planning, budget development and purchasing may also be involved. There can be considerable overlap between tasks done by a computer resource teacher and tasks done by technical support staff, the librarian and the school administrators, so cooperation and clear communication are critical. While the same person may fill the role of network administrator and computer resource teacher, one is essentially a technical role while the other is educational. Both are needed.

Many schools already have computer resource teachers, but they are often expected to fill this role in addition to a full teaching load. This role will become increasingly important in the future and will require time allotted to the necessary tasks. One of the roles of a computer resource teacher is to help teachers change their classroom practices. Some believe that teachers can change their practices with little time, effort or support. The opposite is, in fact, true; change is slow and requires continued support. In particular, educators expected to integrate technology need support from another educator, preferably one based in the same school. This position should have formal recognition, a role definition, and an allotment of time.

Required resources

More than computer hardware is required to provide a useful tool for people in a school, whether they are students, teachers or administrators. Software and other resources such as facilities, furniture, and supplies are needed as well.

Software

Every plan for computer use must take into account funding for software. Without appropriate software, hardware funding will be wasted. Software funding must accompany every computer purchase. Purchasing a computer without software is like buying a car without budgeting for gasoline - no matter how great the deal, the car is still useless. While purchasing software for a network is usually less expensive than purchasing software for an equivalent number of stand-alone computers, the costs are still large and must be planned for.

Software piracy is a serious concern to schools. Stating policies against copyright infringement without providing funding for software is hypocritical; the computers will be used and that takes software, legal or otherwise. Thus software funding must be included in all budget planning. The serious nature of software piracy and its possible consequences must also be understood by all administrators, teachers, teaching assistants, and students.

Physical facilities

Making significant changes to educational practice or to the philosophy guiding the educational system usually requires changes in physical facilities. The costs of such renovations can be substantial, so a school or district must plan them in advance and schedule them over a number of years. Installing school-wide networks in every school in a school district is an efficient approach to technology, but it is also a major project requiring the commitment of a wide range of personnel over an extended period of time.

Furniture and supplies

Plans must also provide for appropriate furniture if the potential of the technology is to be realized. It is all too common to see newly purchased computers placed on old tables of the wrong height, or to see students sitting on uncomfortable chairs. Such situations not only set poor examples to students and hamper the learning environment, but also may lead to injuries and health problems.

If access to outside resources is part of the plan, additional telephone lines may be required. Such access may also require long distance phone charges and subscription fees to information services. These operating costs can be considerable and must be planned in advance.

A variety of supplies are also required for effective computer use. These supplies will be purchased in any event, but operations will be smoother if they are planned in advance. Printers require paper and ink cartridges or ribbons. Diskettes will be required even when a network is in place. There should be funding available for small items that were not included in major orders or that require replacement. This might include printer cables, extra network cables and connectors, mice, mouse balls and mouse pads.

Section Two

An introduction to the technical side of networking

An important note:

Networking is an ever-changing field. To keep this section current, we have added chapter sixteen, The Never-Ending Chapter, at the end of the book. It contains new information and new recommendations that may be different than what we originally wrote in chapters seven through fifteen. While the majority of the material in this section remains valid today, The Never-Ending Chapter will help you see how advances in networking are changing the planning, installation, and maintenance of networks in schools.

We strongly recommend that you read The Never-Ending Chapter before you read anything else in Section Two.

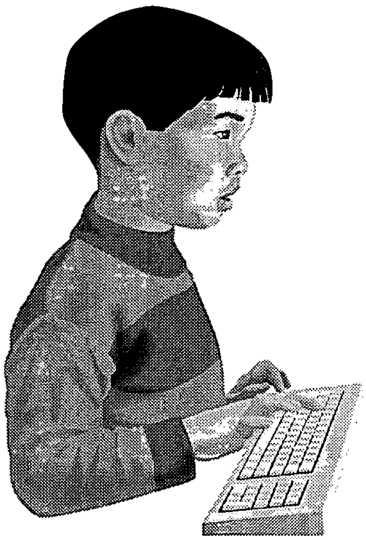

Chapter Seven

Technical details shared by IBM-compatible and Macintosh computers

Chapter 7

Technical details shared by IBM-compatible and Macintosh computers

Topics covered in this chapter...

- *What network administrators do*
- *A basic procedure for installing a network*
- *Understanding operating systems*
- *Applications on a network*
- *Network server software*
- *Communication strategies*
- *Cabling issues*
- *Network topologies*
- *Security issues*
- *Connectivity issues*
- *Facilities and support issues*

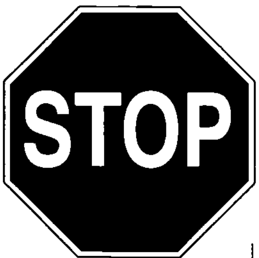

In Chapter One, you were introduced to the basics of computer networking. If you skipped that chapter and are *not* an experienced network administrator, go back and read it now. Chapter One explains many of the terms and underlying concepts required to understand the material we will cover here. The following assumes knowledge of the basics.

This chapter addresses many of the basic technical issues which are common to both IBM-compatible and Macintosh computers. There are many issues which could have been included, but we have tried to limit the discussion to those topics which are critical for someone just starting out with school networks. The material falls into two broad categories – network administration and general technical details.

An important note before we begin...

The material we will cover in this chapter is technical; there's just no way around it. So remember, if your head starts spinning from all of the technical details, it's time to put the book down and let the dust settle a bit before reading any more.

Network administration

An overview of the job

Most of the popular networks available today for both Macintosh and IBM-compatible computers are file share networks with print spooling capabilities. They are usually intended, at least initially, for local use only. That means that they are not interconnected with other networks, and they have no remote connectivity features. If you are unsure of what this means, go directly to Chapter One, do not pass go!

Hopefully, the discussion in the first five chapters will have given you a good idea of the basic features of this kind of network. More importantly, you should have gained an appreciation for the kind of planning that is required to create an effective network. You should also be aware of the potential for growth in the size of a network, the potential for interconnection with other networks, and the potential for adding new features to your network to meet the needs of your users.

Three important points for network administrators

People are more important than cables

It is all too easy to become so focused on the technical details of purchasing, installing, and maintaining computer network hardware that you lose sight of the real reason that your network exists. Computing, especially network computing, is technical and the endeavor naturally attracts those who are good with equipment. There is nothing wrong with this, but we have observed that many of those who can work wonders with equipment have difficulty working with people. We want to set the record straight, right from the start – people are far more important than the technical aspects of a network. For the overall success of your network, it is important that network administrators soberly assess their own strengths and weaknesses. If dealing with things is a lot more appealing than dealing with people, we suggest you consider teaming up with someone whose people skills can augment your technical skills.

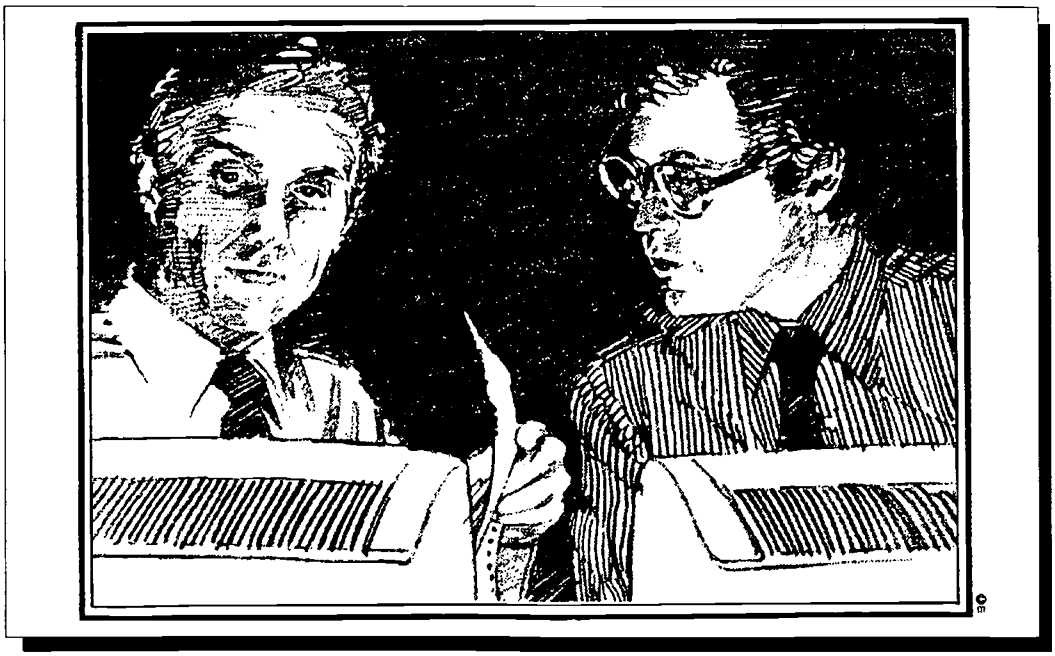

Don't ever lose sight of the people who will use the network you are creating. Computers are simply tools that help people be more productive.

You are a servant

We have encountered many network administrators who see themselves as the most important person in the building, at least from the perspective of computer use. While this may be true, the underlying attitude may cause many problems if the network administrator does not perceive his or her role as serving the needs of others. This is a critical point. If you see your position as meeting the needs of others who will use the networks you create, the chances are very good you will create useful networks that help people do their work, even if you are relatively new to the endeavor.

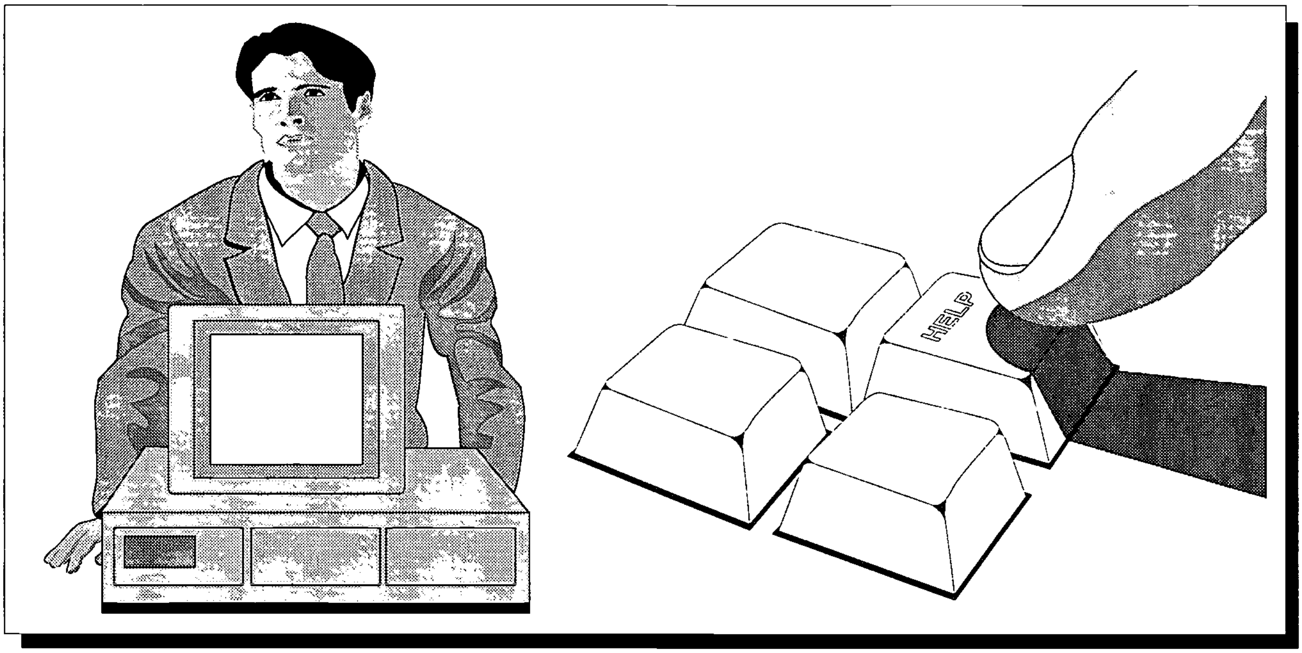

These illustrations capture the essence of the two major roles that network administrators fulfill. First, the administrator is the person with the technical skills who works in the background to make the computers run as intended. Second, the administrator is the one who provides assistance to network users when they are in need of help.

Planning is the key

You should already have picked up our emphasis on planning from the first five chapters of this book. However, just in case it didn't get through, ***planning is the key to successful networks!***

Aspiring network administrators must realize that your decisions will affect the ability of those using your network to perform their work. If a network is well planned, it will make productive work on the computer as easy as possible. If the network is poorly planned, you may be responsible for upset and frustrated users who are prevented from doing their work due to barriers they encounter as a result of your arbitrary decisions.

To further complicate the task of planning, both Macintosh and IBM-compatible networking solutions offer more than one choice of cabling configuration and network server software. Administrators must carefully consider the options. Decisions regarding the kind of cabling, for example, can have significant long-term effects on the ability of a network to serve the needs of its users. Different server software packages offer different features that may or may not be appropriate for your particular needs. Thus, decisions about cabling and server software must be part of the planning process. You must carefully assess the features that a particular network

"Planning is the key to successful networks"

provides and match them with the short- and long-term needs of your users. We will not discuss the criteria for making these decisions here because they are different for the Macintosh and the IBM-compatible computers. Instead, this material is covered in the following chapters which focus on specific computer hardware.

You may now be thinking, "How can I make the right decisions when I have never created or run a network before!" The answer is, of course, that if you are new to networking, you can't possibly know all there is to consider when you create your first network. Consequently, don't expect your first attempt to be perfect. In fact, you should plan on re-configuring the software for your network after you have had some experience with how the network is being used. Eventually, you may even have to re-configure the wiring for your network. We want to reiterate that the goal of this document is not to make your first network perfect, but to make it flexible. This flexibility will allow you to adapt the network to meet the changing needs you will undoubtedly encounter over the life span of any particular configuration. Sooner or later, the network will have to be redone. Hopefully, we can help you delay the need to totally re-do your network. Do your best to plan your network according to the guidelines set out in this book, get all the help you can, and then go with what you have projected. Experience will teach you the rest.

Basic procedure for installing networks

Network components and their function

As we have discovered, most microcomputer networks today are file sharing systems with print spooling and electronic mail capabilities. In this section we will identify the three main components of such a network and briefly describe their functions: the network server, the workstations, and the cabling.

The network server

The network server is a computer that is almost always dedicated to the single task of running the network. The computer acting as the server cannot usually be used as a workstation – it is completely dedicated to running the network server software. The server must have at least one hard disk drive installed internally or connected externally. Network servers may be connected to several disk drives to add more space for providing network services. These services include:

1. File access services:
   - file share access to programs stored in publicly accessible areas
   - file share access to information files in publicly accessible areas
   - password controlled, single user access to private user spaces
2. A print spooler for faster release of workstations when printing
3. Storing and routing electronic mail on the network

Additional equipment may be connected to the server to provide more services on a network. Two of the most common additions are CD-ROM drives for shared access to encyclopedias and other multimedia resources, and network modems for shared access to telecommunications.

Workstations

Workstations are computers that are physically connected to the server by cables. These are the computers that people will use to do their work. In addition to the cabling, each workstation will have some sort of interface and/or connector to attach it to the network cabling.

Network cabling

Network cabling would seem to be a simple matter of running a cable along the wall from computer to computer. Unfortunately, nothing could be further from the truth. There are many different cabling strategies for physically connecting a network. These strategies are known as topologies; we will cover several of them later in this book when we discuss specific hardware configurations. You also have many options for the type of wire you use for cabling the network. Again, we will discuss the various options in subsequent chapters.

General procedures for network installation

If you are creating a network for the first time, you may find it helpful to have an overall blueprint to follow. These steps are applicable to virtually every network. There will be variation, of course, because different networks for different types of computers will emphasize one step over another, and some networks may do some of the steps automatically. Each step will require that you do much more than what is listed here, but this list will give you a general plan to follow.

General procedures for network installation

1. Network planning:
   - Determine long-range and short-term network plans:
     - Interview prospective network users.
     - Determine the programs to be used now.
     - Project future use of network services from the interviews with network users and from research you have done into the major trends in computing. You will probably need outside help to get a reasonable picture of where computing is headed.
   - Detailed short term network plan:
     - Read-only public spaces for programs. Note: you will have to consider how the number of users accessing a program will be limited to the number of legal copies you own of each program you will install on the network. You must also consider how you will protect programs from being illegally copied by users.
     - Read-only public spaces for information files available to all users or specific groups.

- Private storage areas for network users:
  - Access privileges.
  - Automatic log-on features.
- Decide on the number of printers.
- Print spooler space on the server's hard disk if applicable.
- Decide on the program to be used for electronic mail if this feature is not built into the network server software. Also, consider the space that mail files will take on the server's hard disk.
- Booting of workstations – will it be by floppy disk, hard disk, or remote boot?

2. Physical network set up:
   - Computers (workstations).
   - Cabling:
     - Wiring.
     - Wiring closets/racks.
     - Hubs.
     - Backbones.
   - Connectors/interface for each workstation.

3. Prepare network server:
   - Format hard disk.
   - Install network server software.

4. Configure the network:
   - Set operating parameters for server software.

5. Prepare workstation boot disks:
   - Included here, since a boot disk may be needed for step 6.
   - Each boot diskette will contain system boot files and network drivers (called "Inits" by Macintosh users).
   - On some networks these files will be created on the server to allow workstations to "remote boot" without a diskette. We will discuss this in the next section.

6. Run the manager/administrator software (Note: there is wide variation in the approaches and capabilities of the administration programs for the various networks. Some of the following steps may be automated on some networks and transparent to the administrator):
   - Create groups or categories of users.
   - Set access privileges for the categories of users.
   - Create public read-only network spaces (volumes, partitions, or network disks).
   - Create private user spaces (volumes, partitions, or network disks)
   - Set automatic log-on features for network users.

7. Install applications:
   - Copy programs to the server hard disk.
   - Protect network copies from illegal copying.

8. Run and maintain the network:
   - User passwords.
   - Install files in the public information areas.
   - Install new programs or upgraded versions of existing programs.

Understanding operating systems

Network administrators must have a better understanding of how a computer system functions than the average computer user. This is because you will be responsible for installing and upgrading programs on the file server so that they will run on all workstations on the network. You may also be responsible for configuring boot disks that contain system files and additional files for communicating with the network. This section will give the novice administrator the required understanding of computer operating systems.

How does it do that?

How is it that a computer always does the same things when you turn on the power? The computer is like any other device, such as a washing machine or a dishwasher, in that it follows a set of predetermined steps. Wired-in circuitry determines what the device can do and the order in which it is done. If you want to change the order, you have to re-wire the machine.

The circuitry that determines what a computer can do is found on the chips inside. These chips not only determine what the computer does when you turn on the power, but they also contain the instructions for how the computer will carry out its internal calculations, get information from the keyboard, and control the screen. Together, these wired-in instructions that give the computer its basic functions are the core of the computer's "operating system" and are often called the Basic Input Output System (BIOS).

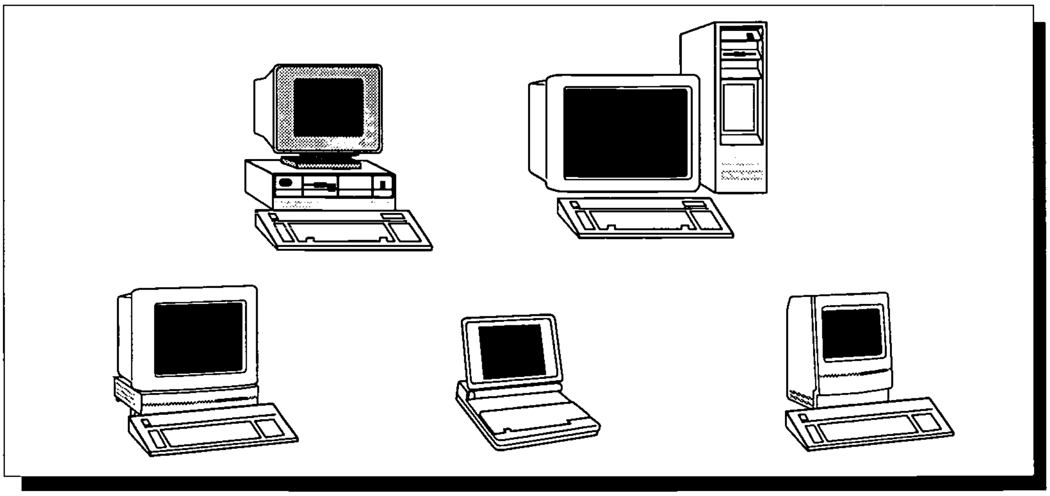

No matter what make of machine they have, most people just use their computer and don't worry about how it provides its various features. However, as a network administrator, you must learn more about the underlying operating system of a computer than the average computer user.

This built-in part of the operating system does not contain instructions on how to control the peripheral equipment that can be connected. Printers, plotters, disk drives, tape drives, scanners, card readers, mice, etc., all require that additional instructions be given to the computer before they will work properly. These additional instructions usually must be loaded into the computer's memory from a hard disk or floppy diskette as part of the start-up process. This is just a bit magical - loading the instructions about operating the disk drive from the disk drive itself. This process of loading the Disk Operating System (DOS) is called "booting the computer."

What does it mean to boot a computer?

The term, "booting a computer," comes from the old expression, "pulling yourself up by your bootstraps." The idea is that you get yourself going by literally reaching down, grabbing your boots, and pulling yourself up without assistance from anyone. This is very much what a computer does when you turn on the power. Simply by routing electrical current, a computer must go from an inactive state to a fully operational machine capable of performing complex logical and mathematical tasks – it has to "pull itself up by its bootstraps." Booting a computer means to start a computer and make the operating system completely functional.

From the description above, you might think that booting a computer would simply involve turning on the power so that the wired-in instructions on the chips could begin to function. This is certainly the case with many electronic devices like calculators, watches, microwave ovens, etc. You simply turn them on and they are completely functional. But booting a computer is not so straightforward. Have you ever noticed that after you turn a computer on, it only goes so far in the boot process and then waits to get information from the disk drive before continuing? In fact, if there is no hard disk drive, and no floppy diskette in the drive, the computer stalls in an endless loop of trying to get information from the disk drive, or the machine reverts to some sort of technical screen display. Why is this so?

The answer lies in the fact that the computer industry undergoes continual change. To keep computers capable of running the latest ideas in the computing world over their useful life, manufacturers must be continually updating the operating system for their machines. If the operating system for a computer were completely wired in, then computer owners would have to get chips replaced inside their equipment every time there was an update. With updates occurring at an average rate of once every year to eighteen months, it would be an expensive and frustrating experience keeping your computer current. Owners would either have to physically bring their computers to a technician to install the new chips or they would have to pay to have a technician come and perform the work wherever the computer was located. Neither option is very attractive.

To get around this problem, computers are programmed to get a significant part of their operating system from a disk and store those instructions in memory. This allows the manufacturer to send updates to computer owners on disks. No physical changes are required, and shipment is fast and economical. Thus, a boot disk, or bootable disk, is a disk that contains the operating system files the computer is looking for. This is why some disks which can be used to save personal files cannot be used to start up the computer – they don't contain the system files the computer needs. This is also why boot disks often have so little free space on them for saving personal work. In some cases the system files are so large that they may fill up the entire disk.

The problem of two floppy disks at each workstation (one with the operating system for booting the computer and one used strictly for saving personal files) led to the idea of "remote booting." To eliminate the boot disk at the workstation, a copy of the operating system files required for booting was stored on the server's hard drive. The server software was then programmed to "listen" for workstations being turned on. As soon as this occurred, the server would send a copy of the operating system files down the cabling to the computer that was just turned on. The end result was that the user now only required their own data disk because the network took care of booting the workstation. It also eliminated the annoying shuffling of disks in and out of the workstation's disk drive when saving a file. In fact, remote boot capabilities can totally eliminate floppy disks from network computers, if you have enough disk space on the server and choose to operate that way.

Operating systems and network administration

Network administrators must be aware of this strategy for loading a computer's operating system because they will be responsible for providing network users with boot capability for their computers, either on floppy disks or on the server. An administrator must also expect frequent updates to the operating system. This will mean updating the network server as well as users' boot disks. It will also mean deciding whether utilizing a certain update is worth the hassle.

To decide whether to update, you should be aware of a standard method computer manufacturers and software developers use to convey the significance of a particular update. A numbering system is used to differentiate between major upgrades and minor updates. While acceptance of this numbering system is not universal, it is used widely enough by major software producers that it provides an excellent yardstick for measuring the significance of updates. This method of numbering updates is used for both system software and application software. The numbering system is as follows:

System or Program Version Number	Interpretation
1.0	This is the first version of an operating system or application program.
1.01	An increase in the hundredths column of the version number indicates a minor update. This fixes some minor programming bugs in the previous version.
1.1	An increase in the tenths column of the version number indicates that some new features have been added to the software and/or some major programming bugs have been corrected.
2.0	An increase in the integer portion of the version number indicates a major upgrade. This may involve a massive restructuring in how the operating system or program looks and operates. Certainly a number of new features have been added.

The first thing a network administrator should do when updated disks are received is to clearly label them with the version number. It is easy to become confused as to which files are the new ones and which are old . To avoid this confusion, the version information should be written on the label of each new disk as it is received. If possible, also add the version number to the name of each of the files on the disk. For example, if the operating system is updated from version 7.38 to 7.39 and the master disks contain a file called, "System", then you should add the version number so that file name is now "System 7.39". If the file names cannot be changed it is a good idea to add the version number to the directory or folder they are stored in and record the size and date of the files for future reference.

The next step is to read the documentation accompanying the new master disks and determine whether the update will have any significant impact on your network. If the update corrects problems associated with an operating system feature that you don't use, then you can ignore this update and wait for something more significant. However, if the update fixes problems with the handling of the computer's internal memory, then this update may impact virtually all network users. In this case, deciding to update the server and all network boot disks is worth the effort. If you are using the remote boot capability, updating is much easier, because a single change to the network copy of the operating system files will immediately update all users.

One final consideration when making a decision to upgrade or not is that every update contains new errors (bugs). The more significant the update, the more likely that the errors will cause problems. It is our recommendation to not install any version of a program whose version number does not contain a decimal (example: version 1.0, 2.0, 3.0, 4.0, 5.0, etc.) until you have reliable verification from several people that the software works. The first version of a major upgrade likely contains many errors. As a general rule, wait for the version numbered with a decimal (example: 1.01, 2.1, 3.2, 4.05, 5.3, etc.). These subsequent versions almost always follow shortly after the release of a major upgrade and contain fixes for the errors that users have discovered in the first release of the new version.

Network updates and network administration

Computers must be given instructions to do everything you expect them to do. Thus, if you want a computer to communicate with a network, you must provide it with additional instructions. These instructions are contained in files that are added to the boot disk. These files are known by various names on different systems, such as "Inits" and "drivers."

When the producer of network software updates the programs that run the network, they will often send out two different kinds of files. The first is an updated network server program. The second is an updated version of the file or files for the network boot disks. Usually, an updated server program requires that all boot disks on the network contain the new Inits or drivers. If you are using floppy disks to boot workstations, this will mean recalling all boot disks, copying on the new files, and returning the boot disks to the users. It is critical that the files and the disks be labelled with the version number of the update to avoid confusion when the next update occurs.

Another cause for updating boot disks is the addition of new hardware on the network. Manufacturers of network adapters commonly send a new driver file along with their new hardware. The efficiency of the driver is often a large factor in the performance of a network, and ensuring that the most up-to-date drivers are installed is an important job for a network administrator. If a newly released computer is used for either the network server or for a workstation, you will likely have to get a new network driver for that machine. Since these drivers change so frequently, networking companies commonly make the files available on electronic bulletin boards. These bulletin boards can be accessed over the phone lines using a modem. Keeping the drivers for your network updated is one job that is well worth the effort.

Applications on a network

And there's more updating to do...

On file sharing networks, the administrator is responsible not only for keeping the operating system up to date, but also for keeping the application programs up to date. The administrator is the only one able to put new copies of programs into public areas on the server. Administrators often underestimate the importance of this task. They spend so much time planning and creating a network that they often feel the critical part of their job is done when the network is up and running. Users of a network see things from a different perspective, and successful network administrators are sensitive to users' needs.

While a great deal of forethought and effort goes into the creation of a network, most network users are much more concerned with what an administrator does to maintain the application programs than they are with the administrator's other work. Users are

People use computers and networks to be more productive. Thus, network users are much more concerned with the programs they use than they are with the technical workings of the network. Keeping up with the developments in the software industry and installing the latest version of popular programs is an important aspect of the network administrator's job.

not really concerned with what the network does as long as they can run the programs they need. Once a program is running on their computer, they are usually oblivious to the existence of the network that gave them the program. In fact, users will often be much more concerned with the features of their application programs than with network or operating system features. When updates occur to their favorite application, users will want immediate access to the new and/or updated features. The bottom line is this: administrators must remember the mindset of the user and focus on updating applications whenever it will make the lives of users easier.

Seeing a network from the users' point of view has four main ramifications for a network administrator. First, administrators must plan for the updates to programs that will certainly occur over the life span of any network configuration. You must see a network as a dynamic and evolving system rather than a static configuration that is set up and then forgotten. Almost all updates require more space than their predecessors. When you create a network, always allow a significant amount of extra room for program updates.

Second, administrators must be aware of the advances taking place in software development that might have an impact of their network. Knowledge of new developments will allow the administrator to bring new tools to the attention of network users. This is an especially important aspect of network administration because many users are not sufficiently computer literate to perceive the importance of new devel-

opments. Network administrators must be continually checking the pulse of the fast moving field of computing. Read whatever you can get your hands on – magazines, journals, books, and other publications that contain news briefs, opinions or other information regarding trends and timelines in the development of new software. Successful administrators see reading as a critical part of their job. You will be forewarned of the imminent release of new versions of software and you will be able to see how the new versions stack up against other programs and overall trends in software development.

Third, new software must always be tested for suitability to a network environment and then, and only then, installed for use. Many programs will work well on a stand-alone computer, but will not run properly when shared over a network. Always test the software yourself before announcing to your users that they will be able to use a new version of a program. Some updates will deal with new hardware that has been brought to the market – new printers, for example. Here is where the Macintosh and the IBM-compatible computers running Windows differ quite radically from IBM-compatible computers running DOS. On the Macintosh and with Windows, new printer drivers or drivers for any other peripheral device are part of the operating environment. DOS, however, handles printers and other peripherals from within each application. Thus, a new printer might require updated drivers for each of the applications you run on the network.

Some programs will not run for more than one user at a time even if they are placed in a public area on the network. Programs should be tested with two or three concurrent users before you declare them fit for general network use. Some developers have multi-user versions of their programs; if a particular program is needed by network users, but you run into the single-user problem, contact the developer to see if a multi-user version of the program exists.

Fourth, when new versions are installed, network users must be informed and trained. This means that you will need a reasonable level of proficiency with the program yourself. You should also consider the timing of updates that will result in significant changes for network users. Changing a much-used word processor just before year-end when teachers are writing reports and final exams and students are completing term papers is not a good idea – for the users or the administrator. To avoid a deluge of emotionally charged individuals descending upon you as the perpetrator of the ultimate insensitivity, pick a moderate- or low-use time on the network and publicize your plan for updating well before the deed is done. If there will be problems, the users will let you know. Furthermore, you can schedule training sessions for the new version, if necessary.

General technical details

Network server software

The three network server software packages most commonly used in schools are Waterloo MacJanet and AppleShare for Macintosh computers, and Novell Netware for IBM-compatible computers. There are a variety of other networks available, including System 7 for Macintosh computers, Microsoft LAN Manager for IBM compatibles, and TOPS for both. There are also a number of networks for other types of computers, such as UNIX-based networks.

For the person using a computer, the printer is critical because it is the device that produces the final copy of their work. Unfortunately for the network administrator, there is no other single piece of equipment that requires so much time and effort to install and maintain. Keeping printers running properly means continually updating printer driver files on the network.

MacJanet, AppleShare, and Netware have centralized file servers which can store data and programs. With these systems the file server is dedicated to that one task; it cannot be used as a workstation. Other network server software such as Apple's System 7, Lantastic, and TOPS are designed as peer-to-peer systems, with no centralized file server. Peer-to-peer systems allow each workstation to access hard disks on other workstations. These are seldom used for large student networks although they may be excellent solutions for smaller administrative networks.

MacJanet is designed for schools, but AppleShare and Netware are designed for business (most businesses using IBM compatibles use Netware and most businesses using Macintosh computers use AppleShare). Since schools are so different than offices, many schools add another layer to Novell Netware to make it fit the needs of the classroom. That additional layer is ICLAS, from IBM, short for IBM Classroom LAN Administration System.

Network communication strategies

The communication strategy or access method of a network is the way messages are sent between the computers. These messages could involve a workstation saving data on the file server or the server sending a program or data to a workstation. Understanding the basics of communication strategies will help you compare the variety of available networks. This topic is the subject of many books, so the following is an incomplete simplification.

Messages sent over network cabling are made up of packets. A packet is a group of characters and/or numbers with a specific fixed size. The size of the packet is determined by the engineers who design the network hardware. If a computer needs to load a large file from a server, it will receive the file in the form of hundreds of packets. The two main ways of getting messages from one computer to another without interfering with other users are by "collision avoidance" and by "token passing." In either case, some sort of traffic cop is required to regulate the flow of packets, or "traffic." The communication strategy defines the role of this traffic cop.

For those interested in technical details, we are using the term "collision avoidance" to mean all strategies called CSMA - Carrier Sense, Multiple Access. Early networks used a strategy called "collision detection" whereas modern networks use an improved strategy called "collision avoidance." We believe the difference is immaterial for the purposes of this book, and will use the one term for both.

In a system which employs collision avoidance, a computer wishing to send information to another computer is able to determine if the network is already busy. If not, the packet of information is sent. If another computer sends a packet of information at the same time, a collision occurs and the network is not able to send the message. The computer waits a short time and then tries again. Since the times for a message to travel across the network are measured in thousandths or even millionths of a second, this does not necessarily mean the network is slow. An analogy here would be a highway, with each computer on the network being like a driveway leading onto the highway. When a car wants to go from a driveway out onto the highway, the driver looks to see if there is traffic coming. If there is, the car stays in the driveway and waits a short while before looking again. The advantage of this strategy is that if there is no traffic, the car can go onto the highway immediately. The disadvantage is that if the highway is very busy, the car may wait a long time before getting onto the highway. Networks employing collision avoidance tend to live up to their speed ratings when there is little traffic on the network (few computers and small programs), but are much slower when there is a lot of traffic (many computers and large programs). These networks include LocalTalk, Ethernet and IBM Baseband.

The other way of allowing messages to be sent between computers is by transmitting an electrical signal or "token" to each computer in turn. When a computer has the token it can send its packet of data over the network. It then it passes the token to the next computer which in turn has its opportunity to send a message. Since this token is passed around the network in a very short time, a token-passing network can also be fast. An analogy here would be a game in which children sit in a circle and pass a microphone around the circle. Only the child with the microphone can speak, and for a limited amount of time before the microphone is passed to the next child. Since this is a more controlled communication strategy, token passing networks do not slow down as much as do collision detecting networks when the traffic on the network increases. These networks include Arcnet, IBM Token Ring and FDDI.

Network cabling

There are four types of cable commonly used to connect computers in a network. These are unshielded twisted pair (UTP), shielded twisted pair (STP), coaxial (coax) and fiber optic. Unfortunately, there are subcategories within each, so one unshielded twisted pair cable is not necessarily like another, and not all coaxial cables are alike.

Unshielded twisted pair (UTP)

Unshielded twisted pair cable is like the thin wire used to connect telephones, except the wires are twisted in a spiral to reduce electrical interference. The major benefits of unshielded wire can be low cost, small size and ease of installation. The detriments are susceptibility to electrical interference and usefulness for only limited distances (usually less than 100 meters).

One reason for the popularity of unshielded cable in business offices is that many office buildings are already wired with this cable for telephones. These buildings generally have 8 wires running to each telephone set, but the telephone only uses two. Some networks, such as PhoneNet (a variant of LocalTalk), can readily use the other wires. Others, such as Ethernet, may be able to use the wires already installed in buildings if they are tested and found to be clear of electrical interference. Since this wiring is so common, the expertise required to install it is readily available. While not many schools have telephone wiring already run to each classroom, those that do can take advantage of this expertise.

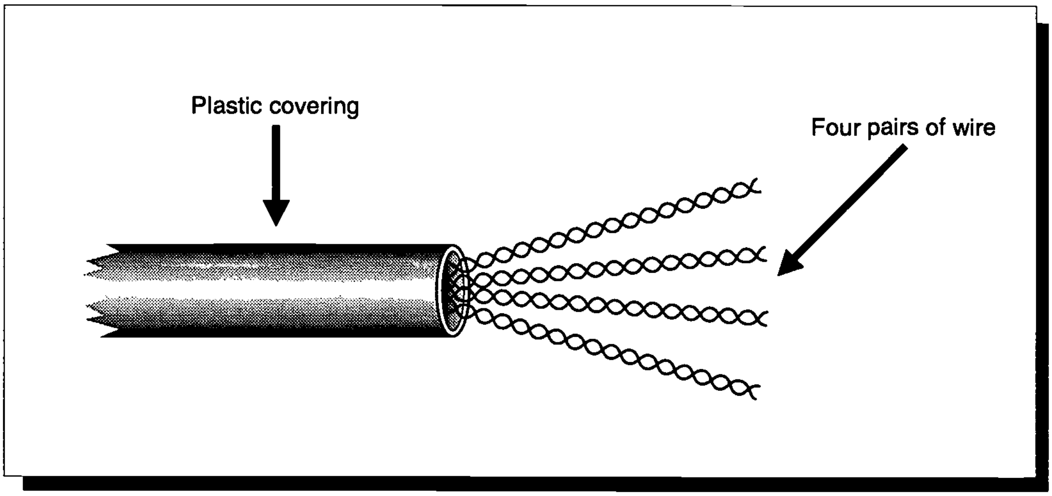

Unshielded Twisted Pair cabling (UTP). This cabling is the cheapest you can install, but it is also the most susceptible to interference. UTP is becoming a popular cabling solution for networks in schools. Newer versions of this wire support high speed networking.

There are special types of unshielded cable which are better than ordinary telephone cable for computer use. For instance, some special cable increases the allowable distance by changing the characteristics of the wire and its insulation. There is also cable especially designed for use with higher speed networks. Three common categories of unshielded twisted pair wire are:

- Category 3 wire - ordinary voice-grade phone wire that can handle speeds up to 4 megabits per second.
- Category 4 wire - designed for speeds up to 16 megabits per second.
- Category 5 wire - designed for speeds up to 100 megabits per second.

At the time of writing, Category 5 wire has only been on the market for a few months. However, the jump in performance with this kind of wiring is significant and we recommend that you install it wherever possible. Remember, just because there is already telephone wiring in a building does not mean it is appropriate for a computer network. Successfully using old wiring for networks is a black art requiring an expert.

The connectors normally used for unshielded cable are small and inexpensive. The normal modular connector used to plug in a telephone is the most common. It has 4 wires and is called an RJ11 connector. Similar connectors designed for computers have 6 wires (RJ12) and 8 wires (RJ45).

Since unshielded twisted pair can be used for so many other types of communications, there is a temptation to use a single cable for multiple purposes; for instance, using an 8-wire cable for a computer network and also a telephone line. This works in some situations, but it is frowned upon by the computer networking industry, especially for higher speed networks. Unless you have expert installers, avoid multiple uses of one cable.

Shielded twisted pair (STP)

Shielded twisted pair is similar to unshielded twisted pair cable in that it has pairs of wires in the center which carry the network signal. A metal foil is wrapped around those wires which acts as an electrical shield for the inner wire. The signal being carried by the network is thus well protected from electrical interference. Shielded cable is more expensive and bulkier than unshielded cable and more difficult to install, but is less susceptible to electrical interference and can in some cases be used over longer distances.

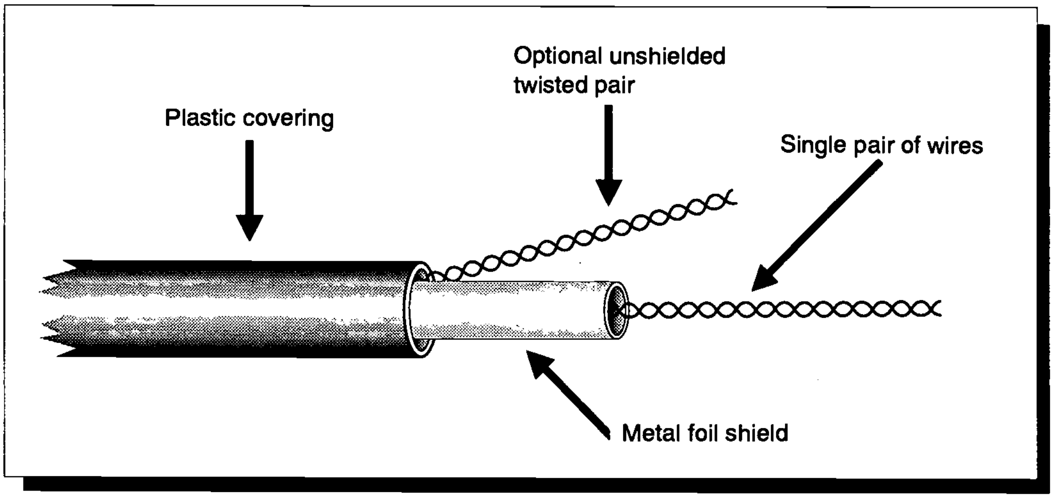

Shielded Twisted Pair cabling (STP). This cabling has the added benefit of a metal foil shield to protect the network signals from interference. This wire is an excellent solution where you have no choice but to run network cabling close to electrical equipment or other sources of electromagnetic interference.

The connectors for shielded cable (designed by IBM) are large, expensive, and reliable. A shielded cable commonly includes extra pairs of unshielded telephone wire so it can serve dual purposes. Connectors are also available to carry a video signal over shielded cable along with a network signal. This is possible because the cable shielding allows higher frequency signals, such as those used for video. Shielded twisted pair cable offers more potential for future growth than unshielded cable; as networks get faster this cable will more likely support the higher speeds.

Coaxial cable.

Coaxial cable is like the wire used to connect a TV to a VCR or to a cable TV system. Given the name coaxial because it has one wire running inside a tube of braided wire (along the same axis), the tube on the outside acts as an electrical shield just like the foil in shielded twisted pair cable. Coaxial cable can be used for long distances (300-600 meters) but is quite bulky and sometimes difficult to install. A system using coaxial cable is generally intermediate in price between systems using unshielded twisted pair cable and systems using shielded twisted pair cable.

There are quite a few varieties of coaxial cable, and many types of connectors. Coaxial cable also has the capacity to carry a lot of information at high speed, as evidenced by its use in cable TV distribution; thus it should theoretically be able to carry network signals at very high speed. However, unshielded twisted pair or optic fiber are currently receiving the strongest recommendations.

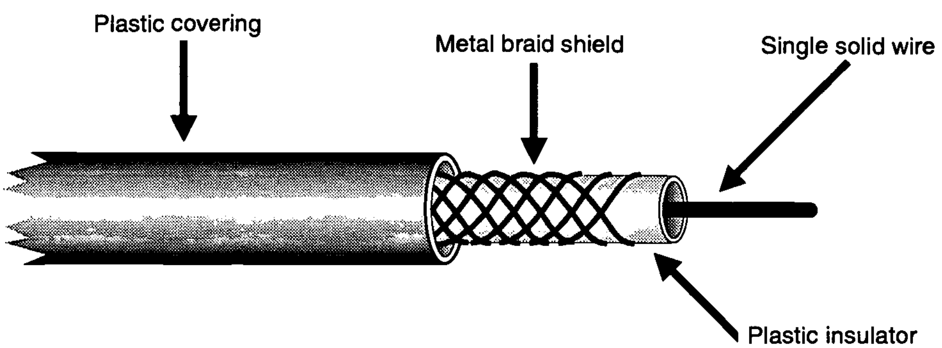

Coaxial Cable. This is the wiring you use to connect your TV or VCR to cable TV services. While the cabling is bulky and relatively expensive, it can carry network signals over long distances – up to 600 meters.

Fiber optic cable

Fiber optic cable is totally immune to electrical interference and can be used for distances up to many kilometers. It is, however, still an expensive option for LAN and school-wide networking. The expense is primarily in the cost of connectors, rather than the cable itself. Fiber optic cabling is new in microcomputer networking. There are not a lot of people with expertise in installing and using fiber optic cable. In the next few years the cost of fiber optic cable should decrease and the number of qualified installers should increase. Any work with this cable requires care and expertise, but work with fiber optic cable requires especially trained personnel, and the connections are time consuming and expensive.

The great benefit of fiber optic cable is its extremely high capacity for carrying electrical signals. Of all the cable types listed here, fiber optic cable gives the greatest expandability for the future. It can carry a wide variety of signals at once, such as TV, computer network, telephone and others. Comparing the cost of one fiber optic cable and associated connections to the cost of many cables for the required tasks will usually determine if fiber optic cable presents a feasible alternative.

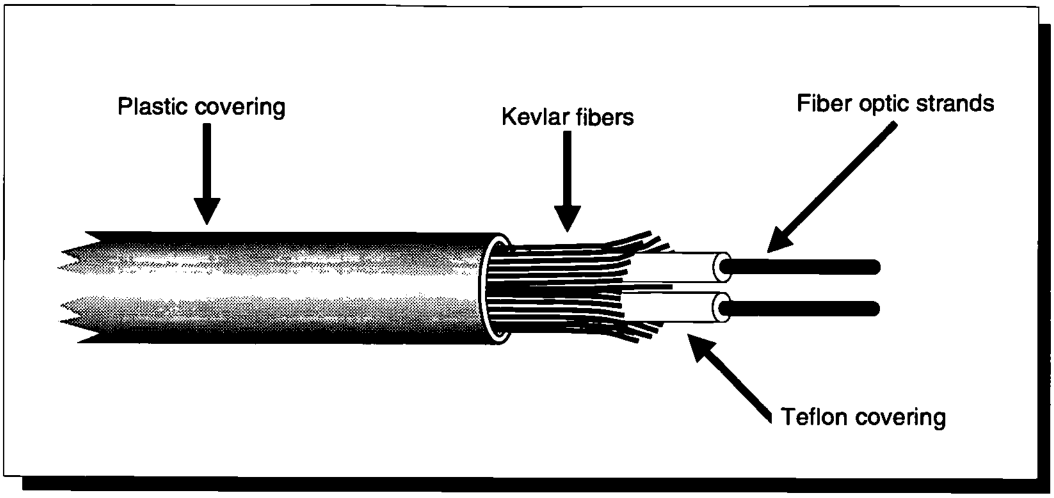

Fiber optic cabling. This cabling is very expensive, but capable of carrying multiple signals simultaneously. This cable is not generally used for networks in schools.

Network topologies

The topology of a network is the physical layout of the cables connecting the computer workstations and the file server. The choice of network topology has a substantial impact on ease of installation, cost, and reliability. It is usually determined by the choice of networking hardware. The most common topologies are the bus, the star, the distributed star, and various combinations.

In any of the following network examples, the file server is connected exactly the same as a workstation. If multiple file servers exist on a single network, a workstation can use resources from one or several.

Bus (Trunkline)

A bus or trunkline is simply a line of computers connected together with a single uninterrupted cable. This cable is the "bus", with each computer connected to it with a short wire and some type of connection device. A terminating resistor is usually required at each end of the bus to stop the signal from reflecting and moving back down the line.

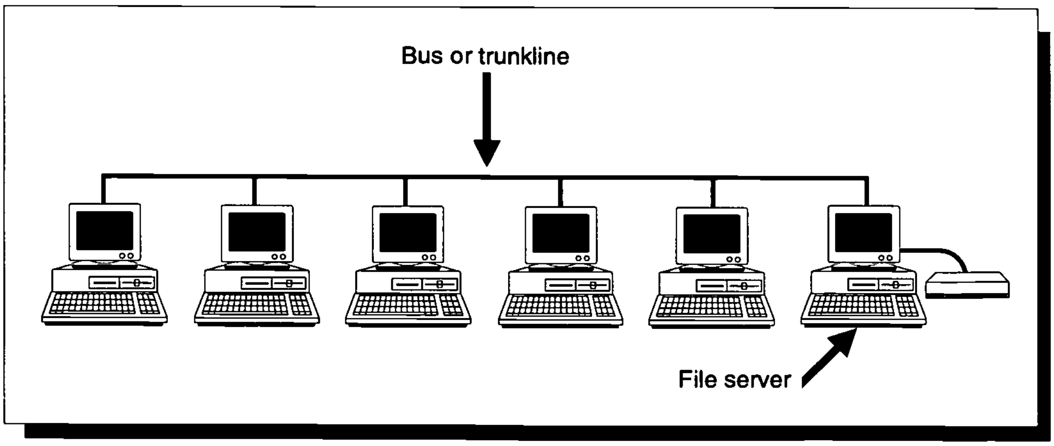

The bus or trunkline topology is very common in networking because it allows network signals to continue to reach all computers even if there is a break in the cabling running out to a computer or there is an electronic problem with one of the workstations.

A bus is usually simple to install in a lab. It uses the least cable of any system and is usually the least expensive. However, it is more complex to install between classrooms, as the bus cable must be routed from the floor or ceiling to the workstation, back into the floor or ceiling, and on to the next room. This is especially difficult if the bus is made of a thick, bulky type of cable. It also makes relocation of workstations quite difficult. If the cable is broken at any point, all the computers on the bus are off the network. If there are obvious locations in the line where each computer connects to the bus (as in the case of Ethernet), those connections are a tempting target for curious students, and thus a likely point of failure. A student playing with a connector can cause havoc with all the computers down the line. If the other computers are in different rooms, the student may not know what trouble has been caused, and the network administrator will likely never find the cause of the problem.

Daisy-chain

A daisy-chain is similar to a bus, except the electrical signal is routed through each computer as it moves along the line. Thus each computer has two connections, one for the cable with the incoming signal, and another to connect to the next cable in the chain. The disadvantage of the bus is made worse with the daisy-chain; since the signal must go through each computer, a malfunction in the network interface of any

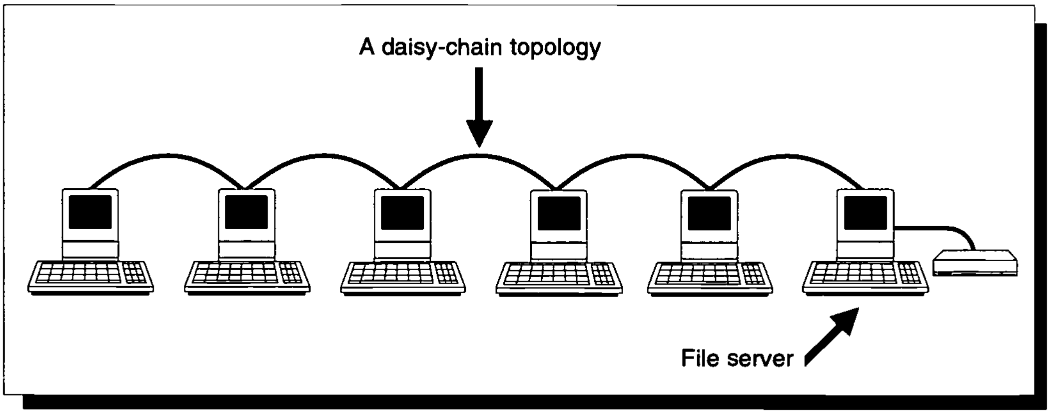

A daisy-chain is often the easiest network to install. However, a break in the network cabling or an electronic problem in one of the computers will cause all computers that are further down the chain to lose touch with the network server.

computer will stop the signal from reaching all the computers along the line. Since it is much more common for a complex network interface to fail than for a cable to break, the chances of network failure are much greater.

Usually the last computer in the chain requires some sort of terminating resistor on the unused connection, serving the same purpose as the terminating resistor at the end of a bus. The difference between a bus and a daisy chain is important for networking Macintosh computers and is considered further in the sections on Macintosh technical issues.

Star

A star topology uses individual cables to connect each computer with a central connection box, called a hub. See below for other names for this device. The star topology uses more cable and is more expensive than a bus. The advantage is that a separate cable connects to each computer, and if one cable breaks, only a single computer should be affected. In practice, this is not always the case, and some star-wired networks, such as Arcnet, will allow one computer to affect others.

A star topology which does not allow one workstation to affect another is the basis of a "structured wiring system." This is a new approach to cabling networks in business environments that improves reliability by ensuring that the failure of one workstation will not affect other workstations. A structured wiring system means that each workstation on a network has its own cable that runs from the computer to a wiring rack. No other computers will use that cable. Although it might seem like a lot of needless work to run all of these individual wires instead of running a single bus or trunkline, the extra work and additional cabling cost is well worth the effort when it comes to a smooth running network where problems are quickly located. For large networks the addition of "network management" allows the network administrator to see what is happening at each workstation and determine who and what is causing problems on the network. An optional module can be added to the wiring rack which has a processor that monitors and reports on all network activity. With network management, problems can be diagnosed and repaired quickly.

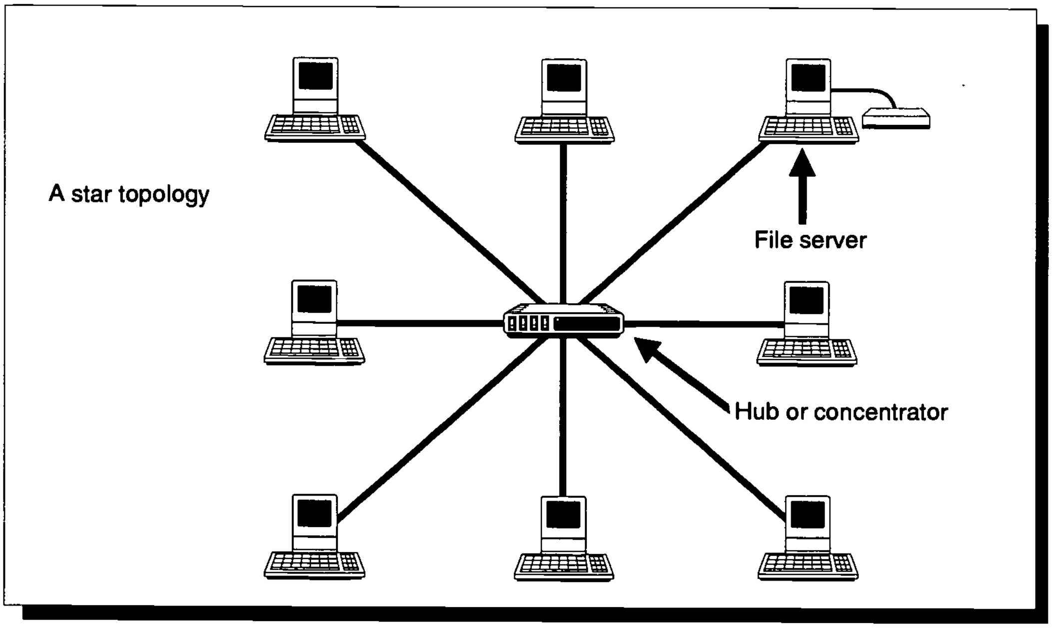

The star topology is the configuration most often used for high speed networks. It lends itself well to a structured wiring system where each computer is on its own cable run.

Most types of network hardware allow stars to be connected together to form a "distributed star", as shown in the diagram at the top of the next page. There will be a limit to the number of stars that can be connected together, usually stated as a maximum number of hubs through which a signal can travel in its journey from one workstation to another.

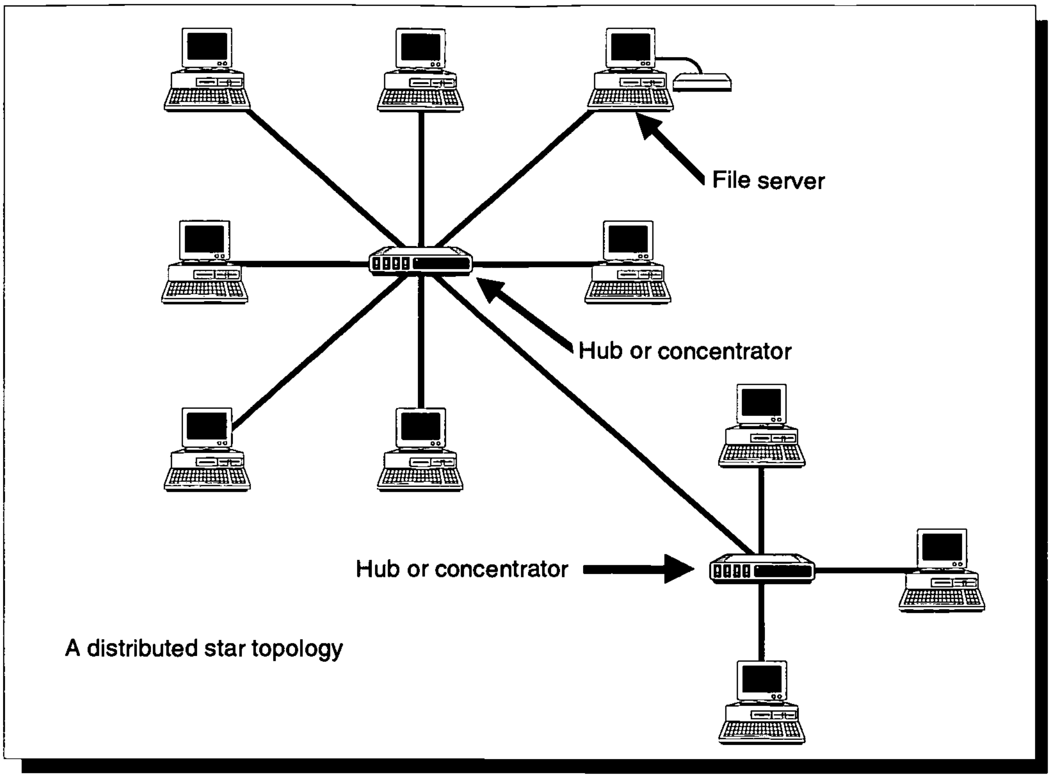

This is one network consisting of two stars connected together. This topology is ideal for buildings where groups of network users are on different floors or in different parts of the facility. The maximum distance between the hubs is determined by the type of cabling being used.

Backbone

Some networking hardware allows a combination of the above topologies, so it is possible to have a number of stars or distributed stars connected together over a "backbone", like the backbone of a skeleton with ribs branching off. The backbone can be used to provide connections between multiple networks (using multiple file servers) or to allow the network to expand throughout a large building. If the backbone is connected only to routers (note: some file servers also act as network routers), it allows a user on one network to access data in another server without having to travel through the cabling of the other network. The signal just travels from server to server along the backbone. This feature is especially critical if the networks attached to the backbone are very slow. In the following discussion of network performance you will see that using a backbone to break a large network into smaller networks is a good planning technique. This technique is used with Ethernet networks in the examples in Chapters Thirteen and Fourteen. A Token Ring is also popular as a backbone with IBM-compatible computers, as shown in Chapter Ten.

On university campuses, the backbone is usually a fiber optic cable connecting the various networks in buildings throughout the campus. The building-wide networks may be buses or stars. As the network expands, the complexity increases, but the goals are always to maximize performance and reliability and minimize maintenance and administrative tasks.

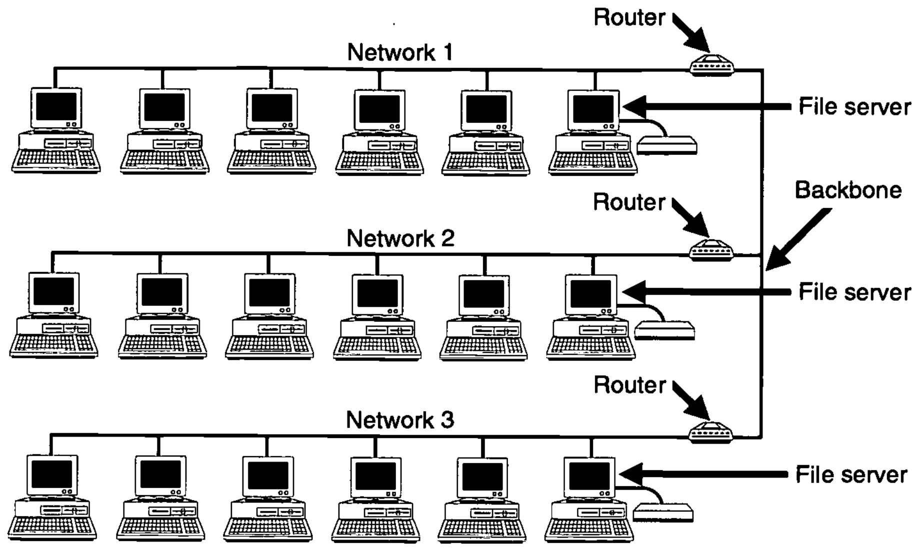

The backbone allows network signals to travel from one network to another without tying up all of the computers on all the inter-connected networks. To isolate internetwork signals to just the two networks involved, you must combine the backbone topology with the use of routers. The routers are located at the end of a network, just before the cabling connects to the backbone. The routers do not let signals pass onto their network unless a signal is specifically targeted for that network.

Hybrid networks

In a large building the connection of every machine to a single point may use too much cable or require cable lengths beyond what is permitted for a particular type of cable. Another problem with larger networks is the cost of running a high speed backbone to all computers. To address these problems, a network may be set up that combines two or more topologies. A common strategy is to set up star topologies in different areas of a building to meet the specific networking needs of the people located there. Then, to create a building-wide network, a high speed backbone is run between the stars. There are many possibilities when it comes to combining network topologies, as long as the overall building-wide network has been properly planned.

Hubs to hook the wires together

The simplest networks are made with wires running from computer to computer in a daisy chain or bus topology. Larger networks require a device to connect the wires from individual workstations together into a star topology. The generic term we use for these devices is "hub." The two tasks accomplished by all hubs is to amplify and re-time electrical signals to make up for the distances travelled to the workstations. Other very desirable properties of some hubs are the ability to isolate the signals from the workstations to ensure that one workstation does not influence another, and to remove a workstation from the network if it is malfunctioning.

While "hub" is valuable as a general term, each type of wiring hardware has a different name for its own device which accomplishes this same task. A complex LocalTalk network of Macintoshes may use a hub called a "star controller." IBM Token Ring networks use a sophisticated device called a "multistation access unit (MAU)" and IBM Baseband networks use a simpler device called an "expander." An Arcnet network

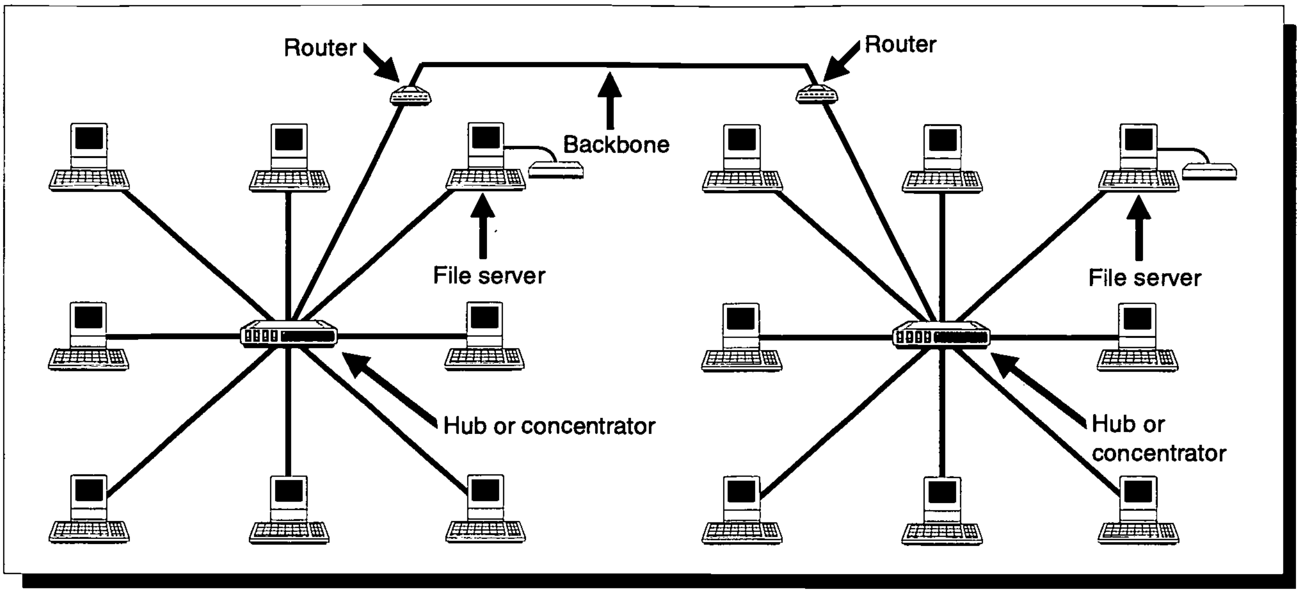

This is a hybrid network where two star networks are connected by a backbone. Again, routers would be used in conjunction with the backbone to isolate those signals that were intended for a local file server. Usually the backbone uses a high speed cabling and hardware to facilitate internetwork communications.

may use two different types of hubs. The central hub, called an "active hub," does amplify and re-time signals, while a secondary hub, called a "passive hub," does not. Thus an Arcnet passive hub is just a connection box which does not isolate workstations and is to be avoided whenever possible.

Some Ethernet networks have a star topology and thus require hubs. The names of these devices are varied and changeable. In advertising for Macintosh users they are called star controllers. In advertising for users of IBM compatibles they are called either hubs or "concentrators." To add even more confusion, the true meaning of concentrator is a very sophisticated unit which accepted modules that act as hubs for either Ethernet, Token Ring or FDDI. In summary, then, we will use the term "hub" to mean any connection box in a network, while we will use the term appropriate to the particular network hardware when speaking of specific networks.

Network performance: speed and flexibility

The ultimate test of a network is how "transparent" it is to the user. That is, how noticeable is the delay while a workstation waits for a response from the network. If there is no delay, the network is not noticed and is therefore called transparent. With an ideal network a user would not notice any delay at any time and thus would not know if a network or a local hard disk is being used. This ideal is achievable for some applications with some networks. Other networks are far from this ideal.

The best way to compare the speed of two networks is to run the same program on each and use a stopwatch to measure the time to complete a task. The ultimate test is to ask an entire class of 30 students to load a program at the same time, and then to measure the time until the last computer is loaded. This time may be a few seconds on a fast network with a small program - in some cases it can be even faster than using a local hard disk. On the other hand, a slow network and a large program can take 5 to 10 minutes to load.

Unfortunately, tests such as these require the networks to be already in place, so a more abstract way of comparing network speed is needed. The computer industry measures the speed of a network as the rate at which data can be sent between computers, in megabits per second (Mbps). These often-quoted numbers do not always give a direct indication of the actual speed of the network for the user, but are useful as a rough comparison. There are other factors as well to be considered than just the raw speed of the cabling. One is the communication strategy being used by the system. For example, token-passing network systems can handle far more traffic without bogging down than collision avoidance systems. If you will be using high speed cabling, four factors to consider when evaluating the speed of the network are:

- the communication strategy being used,
- the capabilities of the network drivers installed on workstation boot disks,
- the speed of information access of the hard disk in the server, and
- the speed of the processor in the computer acting as the server.

As a network grows and more computers are sending packets of information along it, performance will suffer. At some point there will be so much traffic on the network that its speed will be unacceptable to the users. Unfortunately packets do not just go directly from source to destination - each packet may be sent all around the network before it reaches its destination. This is especially troublesome on a network which uses collision detection, as all the packets travelling around the network make collisions more likely.

One of the most effective ways to improve the performance of network is to break it up into two smaller networks and link them with a router. The router is intelligent and only passes on to the other network those packets of information which need to go there. This is the reason for the popularity of backbones. The backbone connects smaller networks through routers. In a network using Novell Netware, the router is included in the file server. Thus the use of a backbone breaks a large network into several smaller networks and improves the performance of all of them. For examples of using backbones in this way, see Chapters Thirteen and Fourteen.

Security issues

Since a school-wide network is designed to store programs and data for the entire school, it must be secure. This means security from equipment failure as well as from users reading or destroying data belonging to others. Security is achieved in three ways: limiting what users can access, using passwords to protect data from other users, and protecting data from physical failure with backups.

Levels of access

Network server software such as MacJanet and Netware have security features built in. These programs allow for different categories of users that have different rights to access the file server. To provide security, each user is given a different "account" with a unique "user ID" and a password. As long as any user "signs in" with their own ID and password they will have only the rights they require.

This means the server software can allow users certain operations with files in specific directories or folders. For example, a user may be allowed to read a file, create a file, write to a file, delete a file, rename a file, or create a new subdirectory or folder. Not all types of server software can control all of these operations. MacJanet and AppleShare, for example, only have two access levels: read-only and read/write.

The categories of users could be network administrator, teacher and student. The access rights they are given could be something like this:

Type of File	Administrator	Teacher	Student
Admin program	Read & write	No access	No access
Teacher program	Read & write	Read only	No access
Student program	Read & write	Read only	Read only
Teacher data	Read & write	Read & write	No access
Student data	Read & write	No access	Read & write

By setting up different rights for different users, the operating system can provide very good security for all users. In this example, only the administrator can make changes to programs. Users cannot access each other's data, except in special public areas.

The amount of control the network administrator has over the rights of the users is determined by the server software, and each type of server software has a slightly different approach to security. Server software systems differ in the amount of responsibility they give the network administrator and teachers for managing student work files. MacJanet and AppleShare give administrators and teachers access to all student data and passwords. Novell Netware, on the other hand, does not allow a teacher to read or change a student's password. It normally does not allow a teacher to read or change a student's data, unless that data is placed in a special area designed for the exchange of information. Netware also has the option of marking a file as "execute only," meaning it cannot be copied from the network, even by the administrator. This feature is designed to avoid software piracy.

Another difference is the amount of control the administrator has over specific parts of the server's hard disk. With MacJanet you can only change the access rights to a predetermined and non-expandable part of the hard disk. This has serious drawbacks, forcing the network planner to calculate carefully the amount of disk space required for each purpose and limiting flexibility once the network is set up. Older versions of Netware improve on this situation by allowing access rights to be set for individual subdirectories, which can change in size as the needs of the users change. Later versions of Netware and most server software based on the UNIX operating system go even further, allowing access rights to be set for individual files.

Some programs such as electronic mail add extra security by encrypting (hiding them in a code) critical information. Thus, even the supervisor cannot read another user's mail. Passwords may also be encrypted. With such methods, microcomputer networks are used to hold critical information in business and can safely be trusted in education, if proper procedures are followed.

Passwords

Virtually all network security problems are related to passwords. If a student is allowed to learn a teacher's password, that student could do anything the teacher could, including deleting all the teacher's work. If a student is allowed to know an administrator's password, that student could do anything the administrator could, which is almost anything - a potentially disastrous situation! Thus, maintaining security on a network is largely a matter of teaching teachers (and especially network administrators) to keep their passwords secret and to change them on a regular basis.

Backup

Any data stored on computer should have a backup copy in case of equipment failure or other data damage. There are two approaches to backing up data on a network. The first is built on the fact that it is much easier to back up data stored on a file server than to back up a large number of individual hard disks. The network administrator then backs up the entire network. This approach is almost always used with administrative networks.

The second approach to network backup is to make data the responsibility of the user, and have the network administrator only back up the programs. This reduces the work of the network administrator and is often used with student networks. It does not, however, ensure that all data is properly backed up, and may lead to data loss by those who have not backed up their personal data.

The most common file server backup procedure involves a tape drive connected to a network workstation. The backup software allows the administrator to choose what data is backed up, and will usually work unattended. Thus, a backup can be done during the night when there are no users on the network. Low-cost tape backup units hold up to 250 megabytes, while more expensive units can hold gigabytes of data.

The frequency of backup is dependent on the kind of information stored on the network. Programs only need to be backed up when they are modified. Administrative data, such as student records, library catalogues and financial information, should be backed up daily. Student data is often backed up less frequently. However, as students become more dependent on the use of computers, regular backup will become increasingly important, whether that task is done by the student or the network administrator.

Connectivity issues

One of the resources which can be made available to users of a network is connection to other computers located outside the school. This can be done by dialing out on a modem, by allowing remote users to call into the network, or by setting up links with other networks.

Modems on the network

The most common way for microcomputer users to connect to other computers is with a modem, dialing out from the network on a normal telephone line. Any computer on a network can have a modem and telephone line dedicated to its own use. If there are many users on the network, however, a better way is to make one workstation, equipped with one or more modems and phone lines, available to all. The modem workstation does not require anything special; this is a good use for an old computer that no longer meets the needs of users. Special modem sharing software is used, and the workstation with the modem may be dedicated to the task or set up to do other work as well. Alternately, a device called a netmodem can be shared over the network. A netmodem makes one or more modems automatically available to all.

Sharing modems for dialing out from the network can be a source of real savings. If there are 50 users on the network who occasionally use a modem, two or three network modems can probably meet their needs. This saves the cost of many modems; more importantly, it also saves the monthly charges for many telephone lines.

Remote access

As teachers and students become more dependent on the use of computer networks, they will demand access from their homes. Providing that access for a few users at any time is not difficult. A workstation can be set up with a modem, dedicated phone line and remote access software. The remote caller can then call in and take control of the workstation. This is easy and quite acceptable for most uses. The speed is reasonable because the software does not send programs and data over the phone line, only changes to the screen - the computer at home is simply monitoring what is being done on the network computer. The security of the network is not compromised if proper procedures are followed, since the remote caller must sign on with an ID and password before gaining access.

As the number of remote users grows, however, the cost of dedicating a network workstation to each remote user becomes unacceptable. There are a variety of solutions, but none are optimal. One approach uses a device called a telebridge, which in effect connects the remote workstation directly to the network. Since the telephone connection is as much as a thousand times slower than a normal network connection, this is useful for exchanging files but not for running programs remotely. As faster modems and telephone links become available, these solutions will improve.

Network to network links

Schools in the province of British Columbia are starting to connect their Local Area Networks (LANs) into larger Wide Area Networks (WANs). Some districts have connected networks in the schools with a network in the board office to form a district-wide network. Such district-wide networks are often used for electronic mail and for shared access to administrative data, other databases and accounting information. Depending on the types of computers and LANs used in the district, there are a variety of ways of creating district-wide networks. However, the various district-wide networks may not be able to communicate with each other.

To get past this potential problem, "protocols" have been developed which act as "international languages." Protocols determine how information is sent and received on a network, and are defined by the engineers who design the network hardware. Unless two computers are transmitting and receiving using the same protocols, it's like two people talking different languages - no communication takes place. These network protocols are supported on a wide range of hardware, and allow dissimilar computers to communicate with each other, including IBM-compatible and Macintosh microcomputers, UNIX-based computers, minicomputers and large mainframes.

The most commonly used protocol of this type is called Transmission Control Protocol/Internet Protocol (TCP/IP). All computers can understand this protocol as long as they have system drivers, Inits, or communication programs containing the instructions necessary to send and receive messages in this format. TCP/IP includes at least three ways for computers to communicate. The first, called File Transfer Protocol (FTP), allows the user of one type of computer to log-onto a different type of computer, view files stored on that machine, and exchange files with it. The second, called Simple Mail Transfer Protocol (SMTP), allows two dissimilar computers to exchange electronic mail. The third, called Telnet, allows one computer to simulate a terminal on the other computer. If the network in your school has a TCP/IP link with a mainframe computer at some distant university, you can use FTP to see what files are available at the university and download them for your own use; the electronic mail program on your network can use SMTP to exchange electronic mail with the distant mainframe; and you can use Telnet to run its programs.

Many schools are forming TCP/IP links with a network called the "Internet." The Internet is a "network of networks" connecting computers at schools, universities and research facilities worldwide. It provides electronic mail and conferencing with educators and students worldwide. The major cost of using the Internet is the local connection. Once a connection is made, the cost is the same whether you are communicating with a computer in the same city, across the country, or around the world.

The details of connecting local networks into wide area networks are far beyond the scope of this document, but there are several other publications that deal specifically with this kind of networking. To venture further, you will need to read up on this topic and probably will need some expert help, because there are many complications, especially when the networks have different operating systems. There are also significant capital and operating costs. The operating costs can be very high if leased phone lines are used, especially in widely dispersed rural districts. Most districts will require financial and technical support to make links beyond the district. The British Columbia Ministry of Education, as well as many other provincial and state agencies are developing plans to support this capability.

Facilities issues

The largest school facilities issue in planning and installing a school-wide network is probably the installation of cables throughout the building. Decisions will have to be made with regard to the routing of cables, the location of wiring racks, and the type of cables used.

Running cables through buildings

A corollary of Murphy's Law should be: "when network cabling is installed in a school the only sure thing is that changes will be needed sooner than expected." The layout of computer labs will be changed, more computers will be needed, and computers will be wanted in different places. Proper network planning can help defeat Murphy in two ways. First, designing a network after consultation with all staff and planning for flexibility and future growth will reduce the number of changes demanded. Second, putting extra thought into cable installation can make it easier and less expensive to make additions.

Cables in new buildings are normally run through conduit (aluminum pipe) or along some type of cable trays or ring supports. Conduit provides support and protection for the cable and insulates it from other sources of electrical interference. To change or add to cable in small conduit is difficult, however. Thus the idea that a school is well designed because it has lots of conduit is not necessarily correct. To provide flexibility for the future, it may be best to use oversized conduit in some places and avoid its use in others. Cable in the walls usually requires conduit, but cable running horizontally in ceilings or under floors can use more easily expandible support such as trays or rings.

Planning for flexibility is even more difficult in older buildings. A plan should be made for the eventual cabling of every classroom, and any wiring project should be a part of that plan. If not, the ceilings will start to resemble a plate of spaghetti, the cable will eventually have to be discarded and the school rewired with an organized plan. Such unnecessary expense is a direct result of poor or nonexistent planning.

Wiring closets and racks

Most network configurations require wiring hubs or other connections in one or more locations. Business offices have rooms called wiring closets dedicated to this purpose. In a school, it is usually possible to reduce the space required for this purpose to a wall-mounted rack.

A network wired as a star will require a number of hubs. These are usually mountable in a standard 19-inch wide wiring rack. The rack may require an A.C. power outlet and should have a lockable cover.

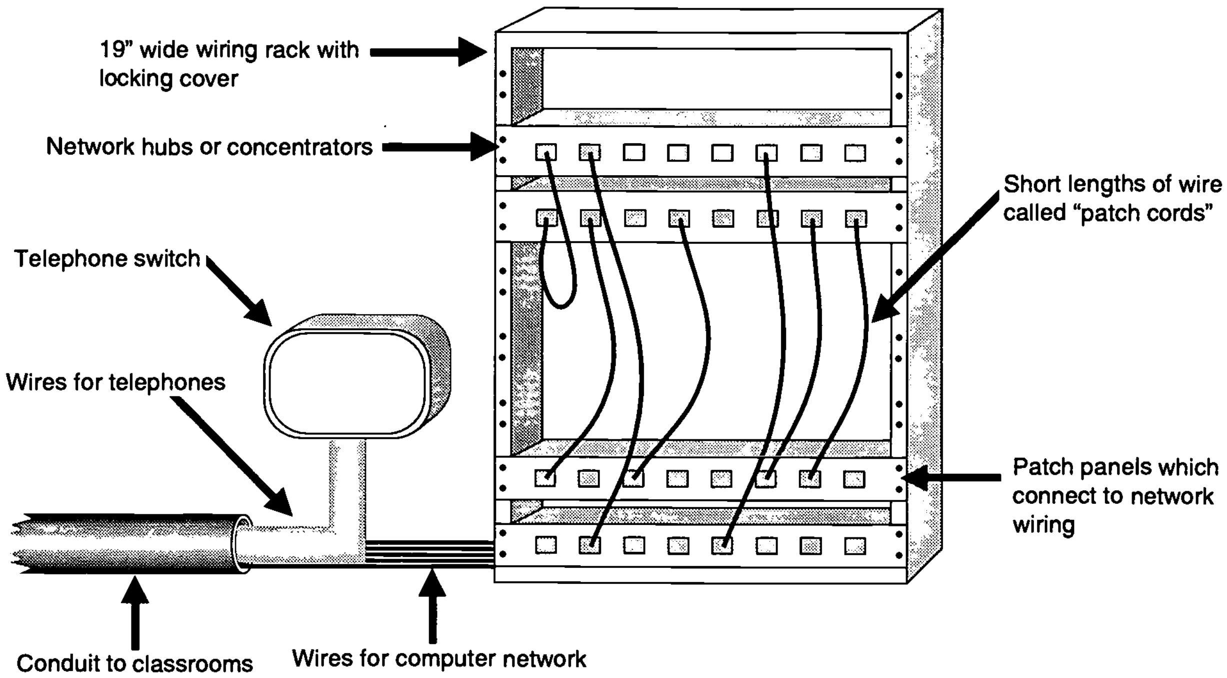

Here is the configuration of a typical wiring closet. The purpose of such an arrangement is to create a structured wiring system where each computer is on its own wiring run from this central location. All of the wires to the individual workstations on the network come back to the patch panels shown at the bottom of the wiring rack. Patch cords are used to connect the wires to the network hubs or concentrators. These devices connect the wires together to create the network.

There are significant advantages in placing the telephone switch system beside the network wiring rack. A telephone line can be run to each classroom along with the network cable to provide a telephone to each teacher. The diagram above shows how the various components may be located. The cost of running such lines is minimized if they are installed at the same time as the network cabling. However, if you are working with the telephone company when installing your network cabling, remember that the phone company is excellent at wiring, but may not be the best source for assistance in network planning. If you are planning any kind of link between your network cabling and the telephone wiring in your building, don't go to the phone company first. Get help from a company with networking expertise, then call the phone company when you have a good idea of what you want.

Plenum vs non-plenum cable

Special plenum-rated cable must be used if it is to be installed through a space used as a heating system air return (a plenum) but not in a conduit. This cable is designed not to give off toxic fumes when burned and not to carry fire along the cable. This type of cable can be twice as expensive as ordinary cable.

Support issues

After planning, technical support is the second requirement for successful long-term network use. Two possible sources are staff hired for this purpose by the district, and hardware and software vendors. The former are required for proper support of the computers in the school; the latter can be a valuable source of information for implementing new technologies.

Technical support staff

Technical support includes installation, troubleshooting, and fixing problems. Actual repair of broken computers may or may not be included. Many school districts leave technical support up to teachers, but this is an inefficient and expensive approach to the problem. These tasks are best done, for less cost, by those with technical training. Therefore, school districts should set a goal of hiring one or more technicians to handle the technical aspects of computer and network installation and maintenance. Technical support staff should have close communication with educational support staff. Thus technicians should plan their work in conjunction with, or be under the supervision of, the computer coordinator (on a district level) or the computer resource teacher (on a school level).

Working with vendors

A close working relationship with hardware and software vendors is a valuable resource. This is especially true for districts with limited in-house technical expertise and support staff. Almost all equipment vendors have technical support telephone lines. Most are free; others are available as part of a purchase or for an extra fee. Having a readily available source of information and advice can make otherwise impossible tasks simple and can lead to a significant saving of staff time and associated costs. The best technical support telephone lines put you directly in contact with technicians and engineers, not sales people. The greatest benefit is obtained if the person making the call has technical training.

Any hardware purchase decision should include a careful consideration of the service and warranty offered. Computer repairs are extremely expensive; within a few years of purchase, replacement becomes more cost effective than repair. If a longer warranty can be negotiated, it is worth some extra cost. Of course, a warranty is useless if the vendor is no longer in business, so the stability and size of the vendor are important considerations. Further savings can be realized by negotiating on-site service when you buy computers, especially if your vendor is not close to your school. Developing a good relationship with a competent and reliable vendor is a wise course for any school district.

IBM-compatible and Macintosh technical issues

A large majority of computers purchased by schools at this time are either IBM, IBM-compatible or Apple Macintosh. While the issues considered to this point of the book pertain to both platforms, there are significant differences.

The first difference is that all Macintosh computers have built-in LocalTalk networking. This makes it easy to decide what type of network to use, and a large majority of Macintosh networks use LocalTalk. Most Macintosh users do not have to face the decisions that IBM-compatible users must make about networking hardware, cabling or topology. At first the task is much easier for Macintosh users, but as the networks grow, this simplicity may be lost. Macintosh users may want to move away from LocalTalk to a faster network. Chapters Eleven, Twelve, and Thirteen are required reading for Macintosh users facing that decision.

Users of IBM-compatible computers do not have the advantage of built-in networking and must face all the complex decisions at the beginning - as you will see in the next chapter. The cost of adding networking hardware is also a significant restraint for those networking IBM-compatible computers. What may make up for the cost and complexity of IBM networking is the resulting speed of the network. Applications will be loaded in seconds rather than minutes, and large networks will be able to function without local hard disks.

Another difference between networking Macintosh and IBM-compatible computers is the server software. The graphical user interface of Macintosh software makes installation of a Macintosh network easier. On the other hand, the server software on an IBM network eliminates or automates some processes that are done manually on a Macintosh network. So IBM network administrators need not be concerned with technical issues such as the sizes of private and public areas, or with tasks such as updating large numbers of boot disks.

You will see these differences in the next few chapters. Some issues are very complex for users of IBM-compatible computers and hardly arise at all for Macintosh users. Others are very complex for Macintosh users and can be ignored by IBM users. In the final chapter we attempt to put all of these issues together by designing a network for both IBM-compatible and Macintosh users. We believe that both types of computers can be effectively used by educators and students. The differences are already becoming blurred and we expect future computers to combine the advantages of both.

Where to find technical help

If you are new to networking you may have found this chapter confusing and may find the next chapters overwhelming. You will need to go somewhere for help. While this book will not make you a technical expert, it should help you understand the questions that need to be asked and evaluate the answers you receive.

Trust your vendor?

The school or district can either research these issues enough to make the proper decision, or put their faith in a vendor to do this for them. In reality, you will probably end up doing a little of both. While a good relationship with a vendor is important, leaving critical decisions to a vendor without researching them yourself is dangerous. Vendors do not generally understand the needs of education and must have their own profit as the final consideration. In addition, putting too much faith in the vendor creates dependence and puts the school at risk of losing control of decision making. So the wise network designer takes the advice of the vendor but knows enough to apply it to the educational environment.

Don't have too many vendors

Compatibility is an issue that must never be taken for granted. Will the designated network work with the computers you already have and also with the computers you buy next year? Will one brand of network card function in a network with other brands of network card, or with hubs of a different brand? Unfortunately the answer is often NO! Once again, making sure your network will function properly requires a lot of research. You will have fewer compatibility problems if you have a single networking vendor. Purchasing every component from a different vendor is a very dangerous path to take.

Issue a request for proposal?

A popular approach to network design is to ask a number of vendors to make all the decisions - by issuing a "request for proposal" (RFP) - and then choosing one you like. This is a useful and often enlightening step, but it is difficult to know what to do with the results of a RFP if you have not already done the homework. The solutions proposed by various vendors will be impossible to compare fairly.

If you are going to issue a RFP, you are usually better off making some of the decisions yourself first, at least tentatively. This will help to define your needs, limit the latitude given to vendors, and make it possible to compare their proposals. You should also realize that the results of your first RFP will probably become the standard you follow for years.

Don't forget about other educators!

While there are many sources of answers to technical questions, the only sources with experience in the same environment are other educators. Networks operating in schools face quite different problems from networks operating in business offices, and a school somewhere in the province or state has likely already faced the same problem you are trying to overcome. We have found that one of the easiest ways of increasing our productivity is developing contacts with computer coordinators and technical support staff in other schools and districts. A phone call to another school district often provides a much clearer answer to a technical question, and you will be sure the answer is not motivated by a profit opportunity.

Now that you know other sources of help, dive into the technical information for your computer (Chapter Eight for IBM-compatible, Chapter Eleven for Macintosh), then look at the examples. Chances are good that one of the examples will be similar to your situation.

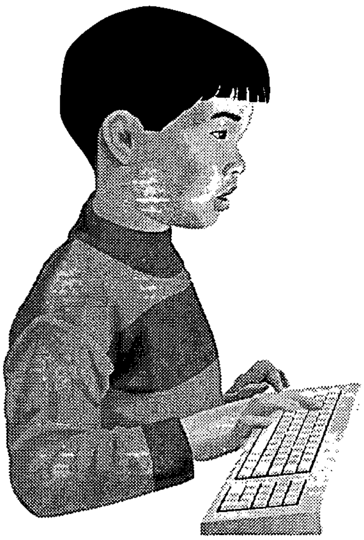

Chapter Eight

Technical issues in IBM-compatible networking

Chapter 8

Technical issues in IBM-compatible networking

Topics covered in this chapter...

- *Decisions to be made for IBM-compatible networking*
- *Quick and simple advice*
- *More complete technical details including:*
  - *Server software*
  - *Types of networking hardware*
  - *Choosing a cable type and topology*
  - *Choosing workstation software*

Introduction

This chapter considers the technical decisions which must be made by those who network IBM-compatible computers. At the beginning of the chapter is a description of those decisions and, for those who are wary of the technical details, some advice from the authors. After reading that advice you may want to read only certain sections of this chapter and then move on to the examples in Chapters Nine and Ten. From that point on, you should have a good idea of the questions to ask your vendor. For those more interested in the technical details, a thorough reading of this chapter, along with discussions with your vendor, should enable you to choose the best approach for your school.

The IBM PC

The microcomputer phenomenon began in 1976 when Radio Shack brought the first pre-assembled microcomputer to the market. In 1977, Apple Computer introduced both the Apple and the Apple II computer models. By 1979, there were a number of manufacturers of microcomputers and it was possible to buy a computer equipped with a reliable disk drive, a printer, and basic business software for $5000. Microcomputers were changing businesses overnight. Clearly, this was not a passing fad; microcomputers represented a huge market.

IBM had been the major world-wide producer of computers for well over a decade. In fact, the name IBM was synonymous with computers. But the computers that IBM produced were large mainframe and mini-computers that cost tens of thousands to several million dollars. This focus caused IBM to misjudge the market potential of microcomputers. Although they had possessed the technology for microcomputers since the early 1970s, IBM was a very late entrant into the microcomputer market when it introduced the IBM PC in 1981.

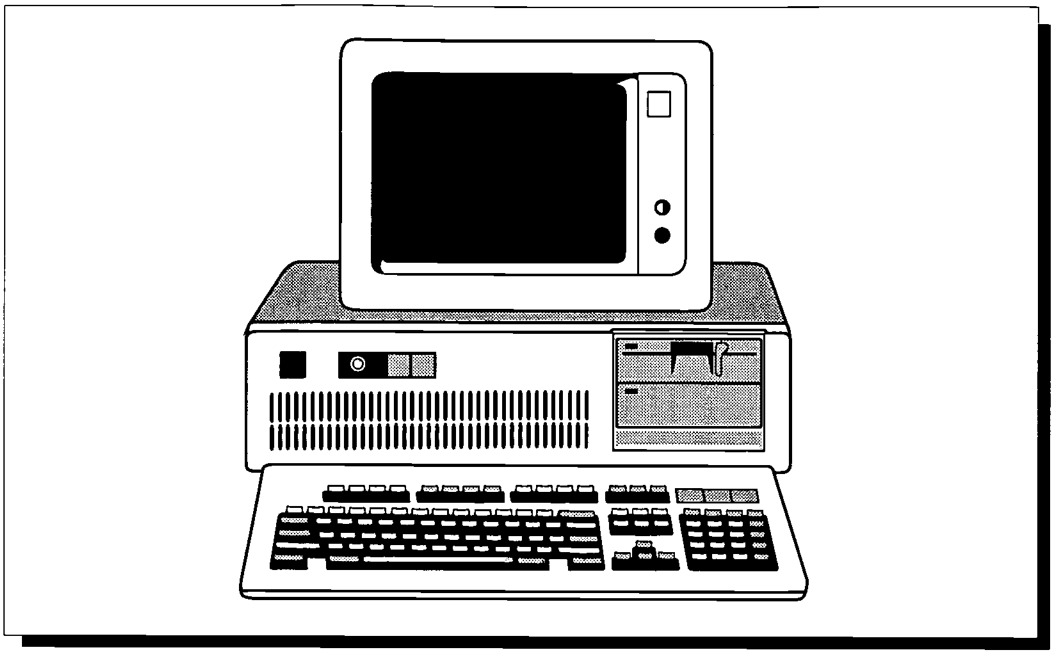

The original IBM PC, introduced in 1981, quickly became a world-wide standard. IBM looked poised to become as dominant in the microcomputer market as it had been in the mainframe and mini-computer markets. However, some critical errors in judgement allowed other manufacturers to clone the PC, and IBM lost control of the market it created.

The fact that you could buy a computer with the IBM name for under $5000 was enough for business and the general public to make the IBM PC the standard machine to buy. IBM's entry into the microcomputer market forced most other manufacturers to close up shop because they just couldn't compete with the IBM name. In the early 1980s, it appeared that IBM machines would completely dominate the microcomputer world. However, even though it had introduced a line of microcomputers, IBM continued to misjudge the microcomputer market. It simply couldn't shake its large computer mentality.

The people at IBM continued to see the future of computers in the mainframe and mini-computer market. As a result, they didn't put their normal effort into the production of their PC. IBM did not take steps to preserve the microcomputer market for itself as it had done in the mainframe and mini-computer markets. It made two critical errors which had long term impacts. First, it used off-the-shelf electronic components that were readily available to any computer manufacturer. Secondly, IBM allowed someone from outside the company to write the disk operating system (DOS) for the machine. That person was Bill Gates, who ran a relatively small company called Microsoft. Gates was able to release a look-a-like version of the original PC-DOS under the name Microsoft DOS, MS-DOS for short. These two errors in judgment allowed a number of other computer manufacturers to produce computers that could do virtually everything that the IBM computers could do, at a lower cost. These computers came to be known as IBM-compatible computers and they began selling more quickly than IBM's machines. Although IBM has tried to regain control of the market, it has been unsuccessful.

The net effect of the emergence of the IBM compatibles has been to create virtual jungle of computer products. The purchase and maintenance of IBM-compatible computer equipment has become a complex task for those responsible in both business and education. Making networking decisions in this environment can be a real challenge. There are so many different network software packages, cabling

systems, interface cards, routers, star controllers, gateways, etc. that the number of decisions can be overwhelming. While we cannot hope to give an exhaustive treatment of networking IBM-compatible computers, we do hope to help you make some effective decisions regarding the configuration of the network for your school.

IBM-compatible networks: too many flavors of ice cream?

The four decisions which must be made in the design of any network are:

1. Server software
2. Networking hardware
3. Cabling
4. Workstation software

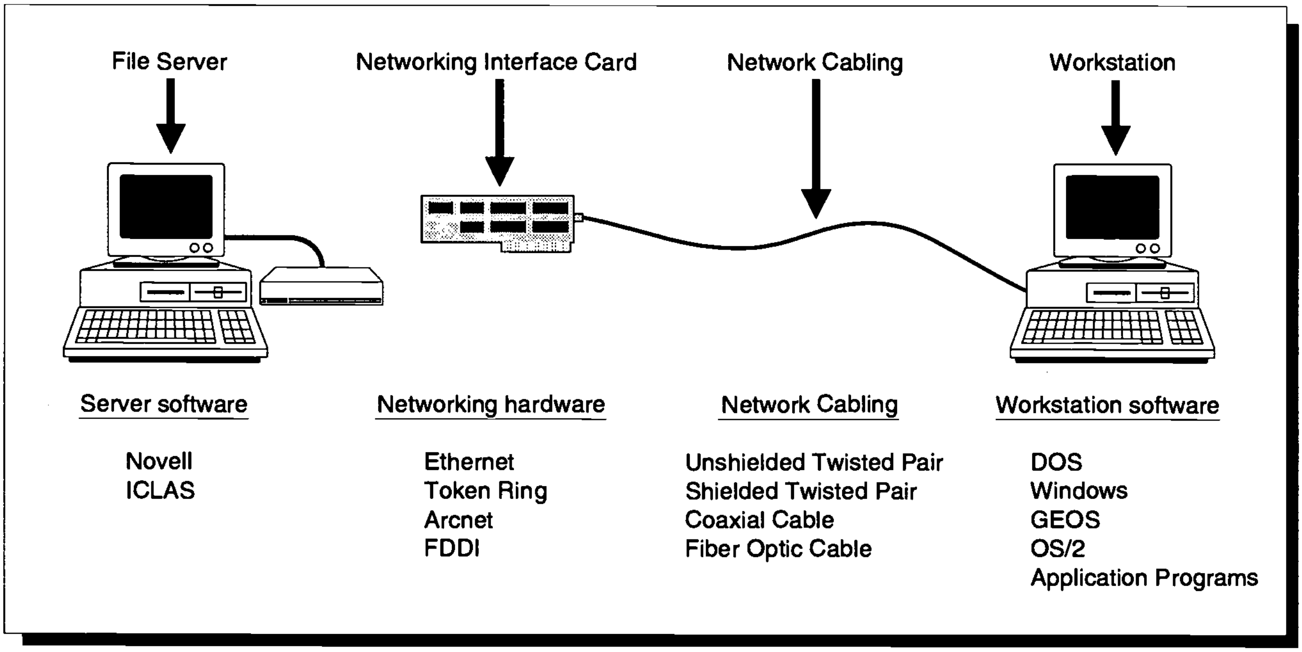

Here are the four major decisions to be made when designing any network. The various options for networking IBM-compatible computers are listed below each heading.

All of these decisions should be made on the basis of the long-term planning that has already been done. To ensure compatibility, these four decisions must be seen as interrelated, together forming a single choice of a networking system.

Just like making choices in a gourmet ice cream shop, choosing standards and components for networking IBM-compatible computers is confusing and difficult. There are so many flavors, some with ridiculous names, others with ridiculous prices. This abundance of choice is one of the dangers in the world of IBM-compatible computers.

When picking a flavor at the ice cream shop, no single flavor suits everyone. Each person has to pick the flavor that's best for herself or himself. Likewise, when you are designing a network of IBM-compatible computers, no single configuration is best for everyone. Each school must choose a design that best meets the needs of its users and fits the available budget. You have succeeded if you create a network that meets your needs today and has some capacity for expansion in the future. But you must create it, for few pre-packaged solutions exist and every decision you make will open up a new set of issues.

How we organized this chapter

Recognizing that many of our readers are new to IBM-compatible networking, we decided to give the novices our unabashed recommendations for setting up a network. This will allow the uninitiated to get on with the business of dealing with vendors and technicians. In addition, we hope that our quick and easy advice will allow them to talk intelligently about the network they want without spending inordinate amounts of time poring over unintelligible technical writing.

For those demented few who actually want to understand the technical background of IBM-compatible networking, we have included a section that contains more complete details. We hope this information will help you begin customizing your networks.

Quick and simple advice

The best way to use this chapter is to read it through, talk to your vendors and other sources of technical support, and then make all the decisions yourself. If, however, you find this chapter confusing and the decisions overwhelming, the following opinions may be just what you are looking for. Remember that these are the opinions of the authors, and there is no shortage of opinions in the computer field.

Server software

The authors believe that the best server software for IBM-compatible computers today is Novell Netware. That decision is easy, so other choices are not even discussed in this document. We also believe that most schools should also use either IBM ICLAS or a menu program on the server to make managing the network easier.

Networking hardware

This chapter will discuss five types of networking hardware: Ethernet, Arcnet, Token Ring, IBM Baseband and FDDI. We believe that most schools should install either Ethernet or Token Ring, just as most businesses do. The limited connectivity options of Arcnet and Baseband often outweigh their lower cost, and the high cost of FDDI takes it out of the range of most school budgets at this time. If a school has standardized on computers from IBM Corporation, has a strong relationship with their vendor, and/or wishes to connect to an IBM minicomputer at the district level, Token Ring is the best choice. For others, lower cost and wide connectivity options make Ethernet the best choice.

Cabling

The authors believe that the best cabling system for connecting computers in classrooms is a structured cabling system, with a separate cable leading from each computer to a hub in a wiring rack. This leads to our recommendation of either shielded or unshielded twisted pair cabling for most situations. Shielded twisted pair cable allows more flexibility for the future but its higher cost may exclude it for many schools. Coaxial cable or fiber optic cable may be required for a backbone in large schools.

Workstations

The majority of programs being used in education on IBM-compatible computers require the MS-DOS operating system. However, the popularity of new programs being written for the Microsoft Windows environment is certain to continue growing. Thus, while in the short term you may have to base your network on MS-DOS, in the long term your plans should include Windows or some similar graphic user interface.

Other considerations

You must also consider your own situation. If the school or district already has one type of network in place, you should consider staying with that standard; it is difficult and expensive to support more than one standard. A good relationship with a vendor experienced in one type of hardware or cabling may also influence your decision, since marginally better or faster hardware will be useless without good technical support. If you must have the least expensive solution, you will probably choose Arcnet, and if you must have the fastest, you will choose FDDI.

That's the simple advice. You can make your decision on the basis of this information if you need to; however, there are two important points to consider. First, if you are going to decide on your network design on the basis of the advice we just gave you, you must get expert help with the actual wiring and configuration of your network. Second, if you want to become more involved in the planning so you can better understand and customize your network, you need to dive into the information in the rest of this chapter.

More complete technical details

This section is for those who want more than just a few brief guidelines for network design and setup. It is technical and we assume that you understand the basic concepts of networking.

The network server software of choice: Novell Netware

A majority of all the networked computers in education and the business world are connected to servers which run Netware software, made by Novell, Inc. While it is complex to install, there are many sources of support. Netware has been expanded and stretched in just about every way possible; if anything has been done with a network it has probably been done with Netware. As of 1993, there are three versions of Netware: 2.2, 3.11, and 4.0. While version 2.2 is still more poplar in schools, added power is making version 3.11 more attractive. Version 4.0 is overkill for most schools and is still very expensive.

Our choice of Netware is based on what is available in 1993. The server software market is growing rapidly and is attracting more competition, so there may be equal or better alternatives within a year or two. For today, however, Netware is our choice for schools.

Netware's remote boot feature is of great importance to schools. By putting a special chip on the network adapter in each workstation, the workstations can start without floppy disks. Remote boot makes maintenance and management much easier and should be included in all plans for student networks.

Netware provides excellent flexibility for the future. It supports Macintosh computers as well as computers using the UNIX or OS/2 operating systems. It also has the greatest selection of connectivity options of any network server software. At the time of writing, the company which makes Netware also owns the rights to the UNIX operating system and to DR-DOS. We can expect to see these three support each other more fully in the near future.

Installing Netware is not a task for novices. Unless a school or district has experienced technicians, installation is best left up to the vendor. A district deciding to standardize on Netware is wise to employ a trained technician. After a few installations the task becomes easier and the understanding gained through the installations makes it possible to provide support.

IBM ICLAS or raw Netware?

Netware was designed for business, not education. To make it fit better into an educational setting, IBM developed the IBM Classroom LAN Administration System (ICLAS). Administering the network with ICLAS is easy, including such tasks as adding programs, making those programs available to users, and adding users (either teachers or students). ICLAS maintains a database of students and their classes and another database of the programs available to each class, giving teachers control over which students are in their class and what programs the students have access to.

Networks without ICLAS are referred to as "raw Netware," and are usually maintained by experienced and dedicated teachers. It is possible to do everything ICLAS does with raw Netware, but it is not as simple. Netware has built-in menus that can be set up for teachers and students, but the manuals for Netware fill a small bookcase and are not easy to follow. Most schools using raw Netware have a technical expert on staff. The danger is that a school with a complex network maintained only by one technical expert may be very dependent on that one person Thus districts using raw Netware are advised to have two people trained in its use, whenever possible.

The added expense of ICLAS is probably worth the cost. There are likely many ICLAS networks in your area and many educators available to provide advice. IBM also provides very good support for ICLAS networks. Although ICLAS is sold by IBM, in the authors' experience it seems to work well with most IBM-compatible computers as well. It should be noted, however, that IBM itself will provide much better support for those who are also using IBM-brand computers. Thus schools using IBM computers should definitely use ICLAS and those using compatibles should use either ICLAS or consider another menu program. While they are not designed for education, menu programs such as Sabre Menu are easier to use and maintain than Novell's menus. Such programs can make life easier for the network administrator.

Other options

Netware is not the only network server software available for IBM-compatible computers and there are some instances where other options should be considered. For small administration networks it is reasonable to consider a network without a dedicated file server, such as Lantastic or Tops. In making such a choice, you should consider whether the software you require will run on such a network and whether your technical staff can support it easily. In the longer term it is likely that server software based on OS/2 or UNIX or some other operating system will become a good choice for schools. The world of networking is constantly changing and good advice ages quickly.

Types of networking hardware

The choice of networking hardware is probably the most difficult decision to make, as there is no one right answer. The decision is made more difficult because each network system is evolving. Each was originally designed with one speed, one cable type and one topology, but most now allow a variety. The following table shows the types of cable which can now be used with each type of networking hardware:

Type of Cable	Ethernet	Token Ring	Arcnet	FDDI
Unshielded twisted pair	✔	✔	✔	Planned
Shielded twisted pair		✔		✔
Coaxial cable	✔		✔	
Fiber optic cable	✔	✔		✔

Ethernet

Ethernet is the most commonly used network in the IBM world, and is used often with Macintosh computers, computers using the UNIX operating system, and some minicomputers. It uses collision avoidance and runs at a speed of 10 megabits per second. Ethernet has a variety of standards which allow it to run on different cables:

Type of Cable	Ethernet Standard
Thick coaxial cable	10base-5 (thicknet)
Thin coaxial cable	10base-2 (thinnet)
Unshielded twisted pair cable	10base-T

The Ethernet standard is giving you three pieces of information. The first number indicates how fast the signal is travelling on the cabling – in all three cases here, the signal is travelling at 10 megabits per second (10 million bits per second). The word "base" indicates the type of signalling. "Base" indicates baseband signalling, a digital signal; "broad" indicates broadband signalling, an analog signal. The final piece of information is the maximum length of a cabling run before a signal is boosted. 10base-5 indicates a maximum length of 500 meters, 10base-2 indicates a maximum length of 200 meters, and 10base-T has a maximum length of 100 meters.

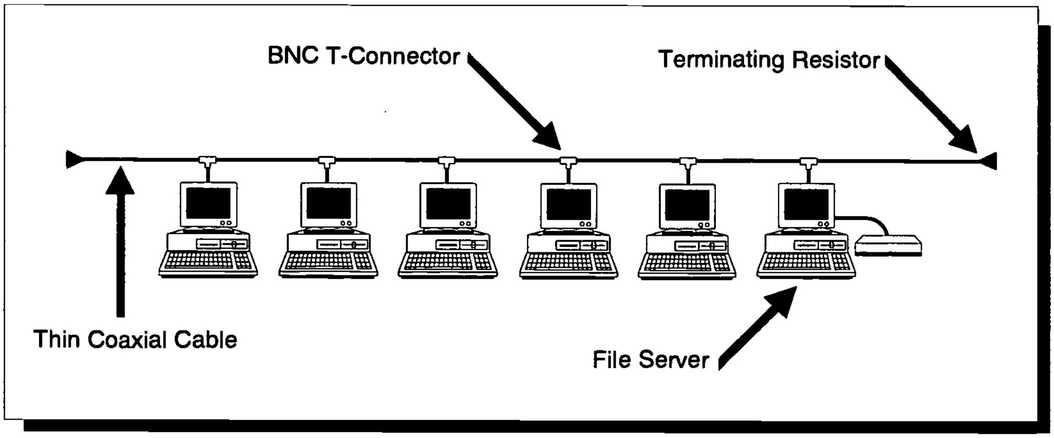

An Ethernet/Thinnet network using thin coaxial cable (RG 58). This network is configured as a bus and requires no hubs. Until 1992, this was the most popular network for IBM compatibles.

The original thick coaxial cable is seldom used in schools. If either type of coaxial cable is used, Ethernet is configured as a bus. With unshielded twisted pair cable, a star topology is used. Components are available to connect stars of unshielded twisted pair cable with a backbone bus of coaxial cable. Ethernet networks with this combination of topologies often have hundreds of workstations. We recommend 10base-T with a structured wiring system for Ethernet wiring.

Since there are such a variety of Ethernet connectors, it is important to use cards with the right connections. A 10base-T card designed for unshielded twisted pair will have a RJ45 connector, one designed for thin coaxial cable will have a BNC connector, and one designed for thick coaxial cable will have an AUI connector. Installing an Ethernet network will require outside expert help – don't try to do it all by yourself.

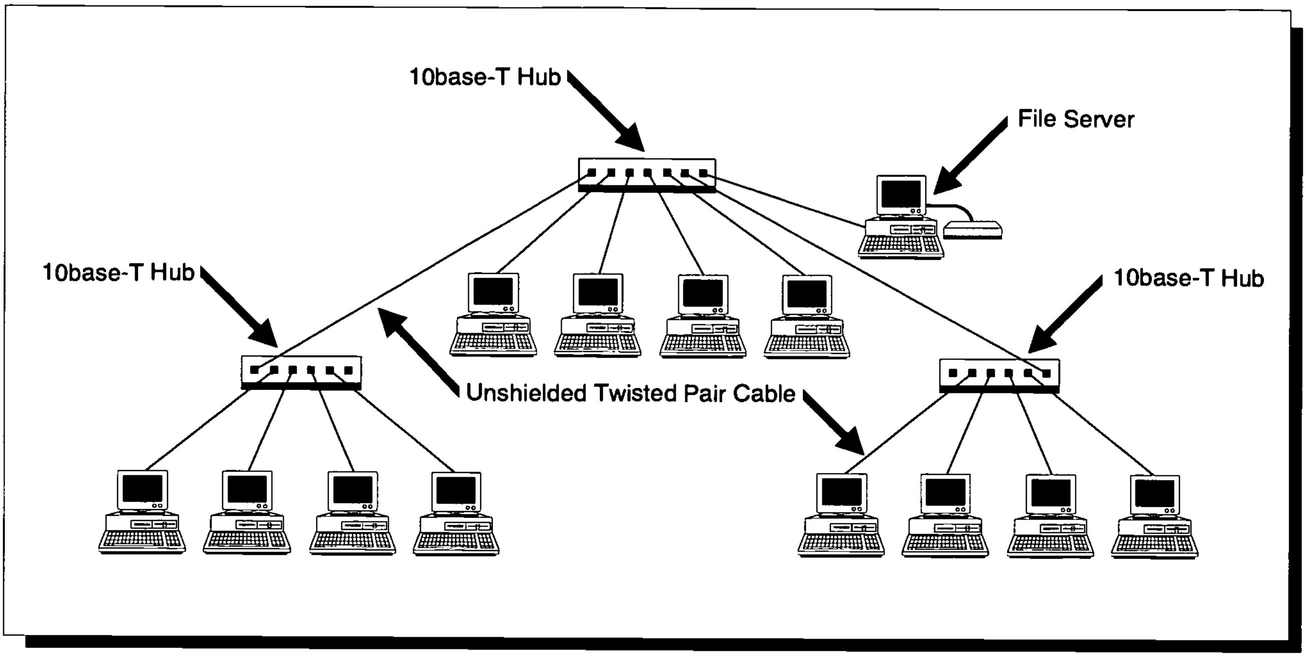

A 10base-T network using Unshielded Twisted Pair cable is the most common Ethernet topology. The network is configured as a distributed star using a structured wiring system. The hubs/concentrators are available with 4, 8, 12, 16, or 24 connections. This is the most popular type of network today.

In the near future Ethernet components will commonly come as part of computers. A few IBM compatibles and the Macintosh Quadra models already have built-in Ethernet, as do almost all UNIX-based computers.. For examples of Ethernet networks in schools, see Chapters Ten, Thirteen and Fourteen.

Token ring

Designed and supported by IBM, Token Ring is popular in business. It is a token passing network which runs at either 4 or 16 megabits per second. This makes it especially suitable for large high performance networks. While adapter cards rated for 4 megabit speed are less expensive, the future lies with the 16 megabit speed, or higher. This means you must use either shielded twisted pair wire or a minimum of Category 4 unshielded pair cable. Since all adapters on a network must use the same speed, it is probably worth the extra cost to buy the faster cards now.

Normally running on shielded twisted pair, Token Ring has been adapted to run on both unshielded twisted pair cable and fiber optic cable. A small Token Ring network is configured as a star, but for larger installations it is typically a ring of stars, with the ring serving as the backbone.

Token Ring adapters are available for Macs, but not commonly used, and are also used for IBM minicomputers. Token Ring networks can support hundreds of workstations. It is likely that some computers will soon come with built-in Token Ring hardware. For an example of a school network using Token Ring, see Chapter Ten.

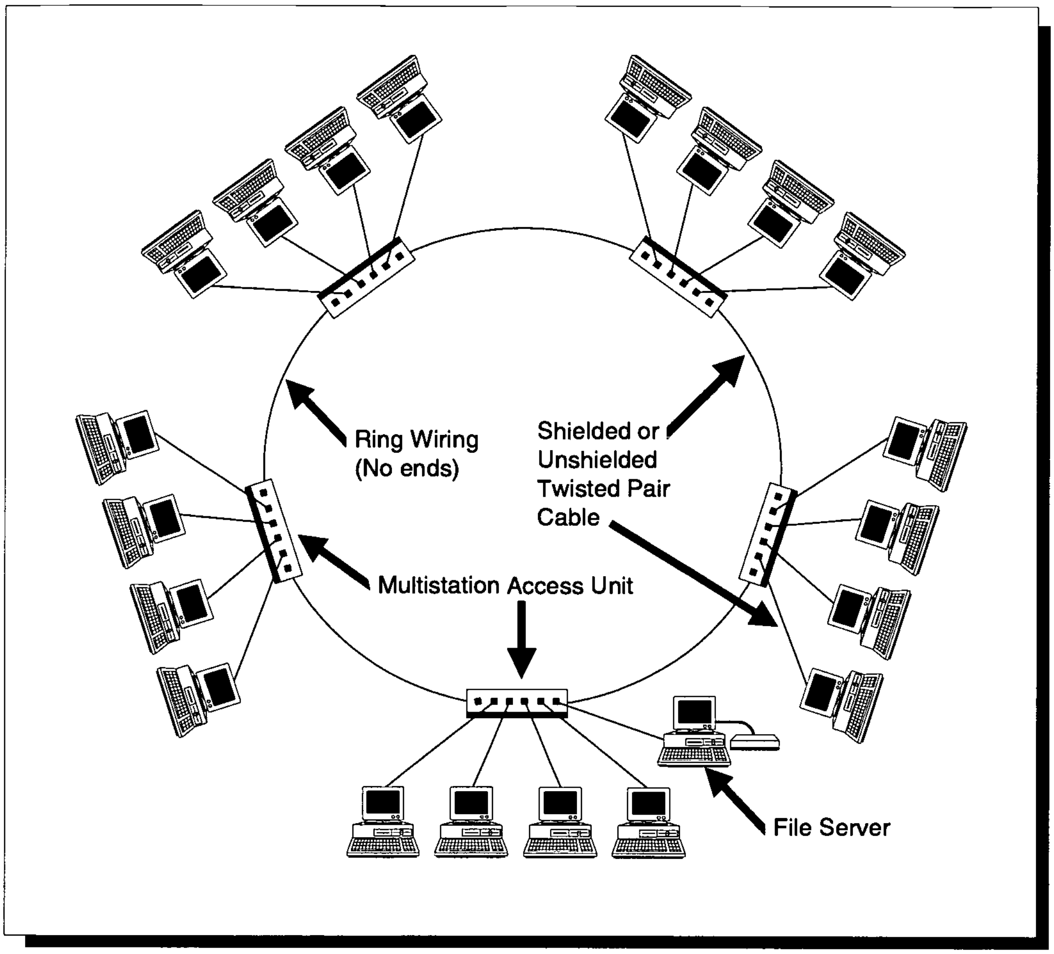

A Token Ring network. While each of the Multistation Access Units (MAUs) are wired as a star, the key to this network is the unending loop of cabling connecting all of the MAUs together. The network can run on either Shielded or Unshielded Twisted Pair cable. In most cases, the ring is not obvious because the MAUs are all together in the same wiring rack. Token Ring is a high-performance, high-quality, high-cost network.

Arcnet

Arcnet was the original network for connecting IBM-compatible computers, and is still used in many schools in Canada and the United States. It uses a token-passing communication strategy and has a speed of 2.5 megabits per second. Arcnet was originally designed as a distributed star, but has also been configured as a bus. Originally designed for coaxial cable, Arcnet has also been adapted to run on unshielded twisted pair cable.

It should be noted that Ethernet runs on two different types of coaxial cable and Arcnet runs on coaxial cable as well, but all three are incompatible. Arcnet is the least expensive of IBM networks, but is becoming less popular as Ethernet becomes competitively priced. There is no example using Arcnet in this book.

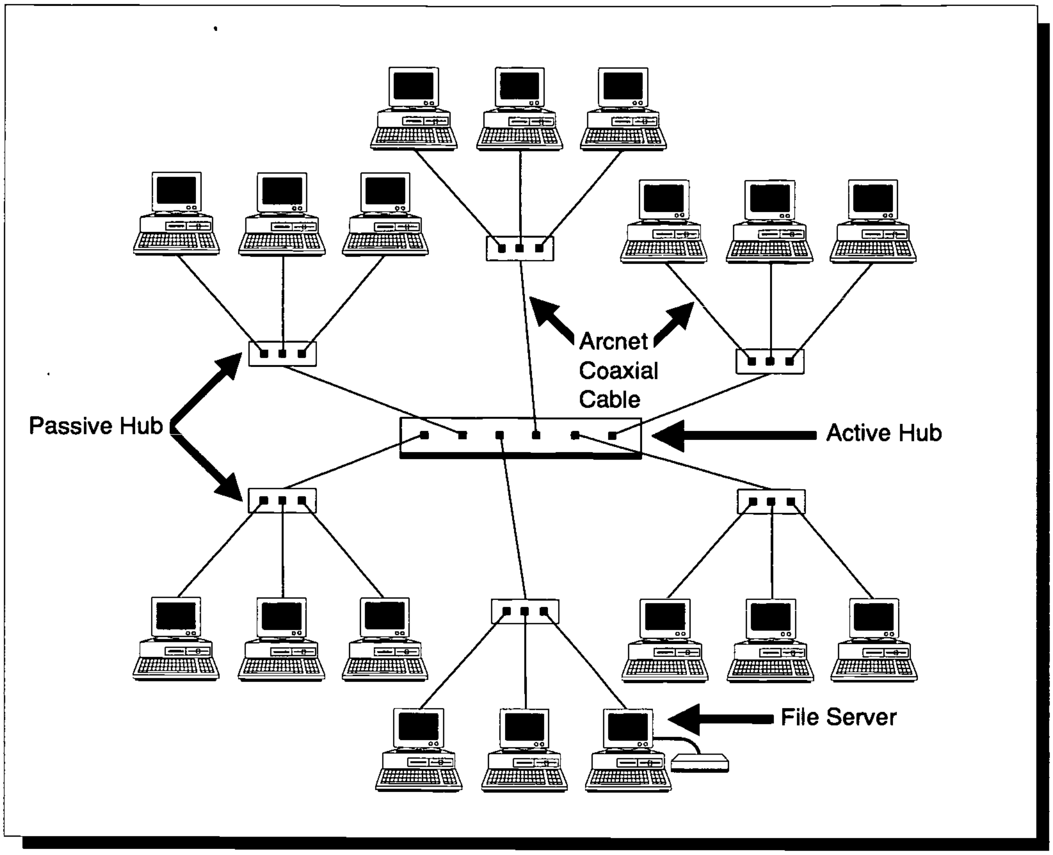

An Arcnet network using Arcnet coaxial cable (RG 62). The network is configured as a distributed star. Such Arcnet networks suffer from interference from malfunctioning workstation because the passive hubs do not isolate the workstations from each other. Thus, passive hubs should be avoided. A bus configuration is also possible. Ethernet is a better solution for new installations.

IBM Baseband

IBM Baseband is used in many schools, but is seldom seen in business. It uses unshielded twisted pair cable configured as a daisy-chain. Using a collision avoidance communications strategy, it runs at a speed of 2 megabits per second. Baseband is designed for use within a classroom, not throughout a building. It is relatively inexpensive, but does not seem to be a standard which allows for future growth. For an example of a school network using IBM Baseband, see the next page and also the network described in Chapter Ten.

FDDI

FDDI (Fiber Distributed Data Interface) is a 100-megabits-per-second network, by far the fastest available. It is a token passing network which has two separate rings for increased reliability. At present, FDDI is mainly used to inter-connect other networks, often over large distances, but is expected to become common for connecting individual workstations within the decade. Although designed for fiber optic cable, there is now a standard for using FDDI on shielded twisted pair cable, and the industry is working on FDDI for unshielded twisted pair cable. Its popularity is presently limited by its very high cost. There is no example using FDDI in this book.

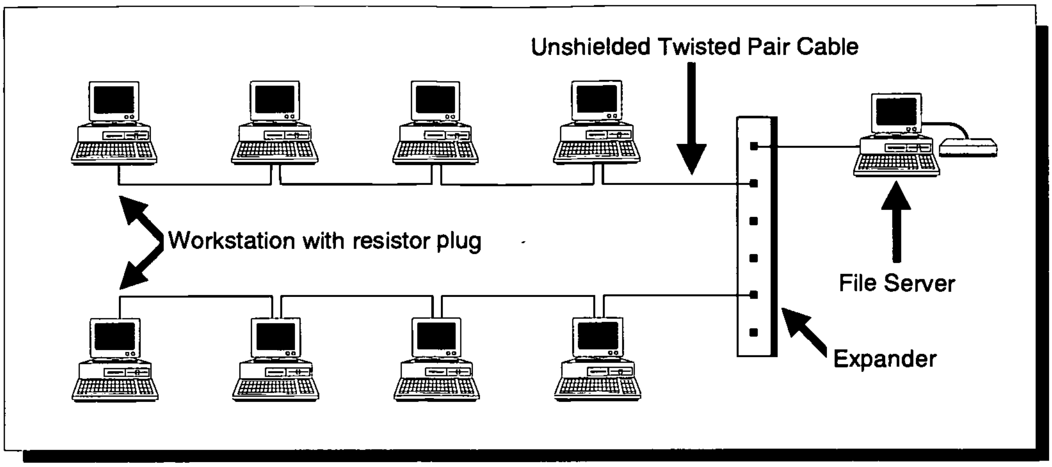

An IBM Baseband network using Unshielded Twisted Pair cable. This network is configured as a series of daisy-chains connected to a hub that IBM calls an expander. This system is designed for use in a single room and is not suitable for connecting classrooms together. Ethernet is generally a better solution.

Choosing a cable type and topology

Your choice of cable and the way it is laid out in the school will be limited by your choice of networking hardware, and vice versa. Your choice will depend on your goals. Three possible goals are maximum flexibility, maximum expandability, and minimum cost. Let's discuss these in detail.

Cabling for maximum flexibility

A star configuration of unshielded twisted pair cable is probably the most flexible cabling system. It consists of one or more wiring racks, with unshielded twisted pair cable going to each classroom. Since this layout requires a separate cable to each computer, it is best to have at least two cables going to each classroom (for two computers). Each cable should have four pairs of wires and use RJ45 modular connectors (with eight connections). If the entire school is being wired, it may also make sense to include an extra cable for telephone use.

This cabling system can support Arcnet, Ethernet or Token Ring, and in the future will probably support FDDI. If you are thinking of FDDI in the future, you should specify the 100 megabit cable that has recently arrived on the market.

Since unshielded twisted pair cable has a limit of 100 meters, large schools may require more than one wiring rack. One could be placed at each end of the school, linked by a "backbone" of coaxial cable (for Ethernet or Arcnet), shielded twisted pair cable (for Token Ring), or fiber optic cable. The example in Chapter Ten uses a Token Ring backbone of shielded twisted pair cable, and the examples in Chapters Thirteen and Fourteen use Ethernet backbones on coaxial cable.

Cabling for maximum expandability

A star configuration of fiber optic cable or shielded twisted pair cable will provide greater expansion capability. Either of these cables can carry video, telephone and computer networking. (With shielded twisted pair cable the telephone signal is carried on a separate telephone wire within the same covering.) Either can also be used for FDDI, if a higher speed network becomes necessary.

Since the cost of fiber optic cable is so high, it is likely that the network designer in a large school, looking for maximum expandability at a reasonable price, will choose shielded twisted pair cable to connect the workstations, with fiber optic cable used only to connect the wiring racks. Another possibility is to use unshielded twisted pair cable within labs and shielded twisted pair cable throughout the rest of the school. An example of this type of cabling system is provided in Chapter Ten.

Cabling for minimum cost using Arcnet

An Arcnet network is the least expensive, because its network adapters are the cheapest. Arcnet is wired as a distributed star, with a central active hub connected to passive hubs, which are in turn connected to three workstations each. The problem with this arrangement is that a broken workstation can affect other workstations on the network. Maintenance and troubleshooting become very difficult as the network grows. Therefore, even though this approach is cheap, it is not recommended. Arcnet can also be configured as a bus, but with the same drawbacks.

Arcnet can also be used without passive hubs, wired as a simple star with every workstation connected to a central hub. This is a great improvement, and the cost is still low if unshielded twisted pair cable is used instead of coaxial cable. However, we believe Ethernet is superior to Arcnet and should be used if possible.

While Arcnet uses only a single pair of wires in an unshielded twisted pair cable and RJ11 connectors (four connections), it will also work with four pair of wires in the cable and RJ45 connectors (eight connections). This works because the smaller RJ11 plug on a telephone wire will fit into a larger RJ45 wall jack. So, if unshielded twisted pair cabling is going to be put into the walls of a school, it makes sense to use category five, four-pair cable which will not only support Arcnet, but Ethernet and Token Ring as well.

Cabling for low cost using Ethernet

The type of network with the greatest number of installations is probably Ethernet coax. Actually two different types of coaxial cable are used with Ethernet; the one discussed here is thin coaxial cable, also called thinnet or cheapernet. The coaxial cable is installed as a bus. While this is a simple and obvious approach for a lab, it can also be used school-wide, with the cable dropping into each classroom as it winds its way around the school.

The problem with this cabling system is inherent in the bus topology: if the bus is disconnected at any point, all workstations are affected. This is quite a risk when the bus drops into every classroom. Troubleshooting is also difficult.

Choosing a network adapter

Once the type of network hardware and the cable type have been decided, the network adapters can be chosen. Some adapters can support more than one type of wire. For instance, Arcnet and Ethernet adapters are available which support both unshielded twisted pair and coax. Token ring adapters can support both unshielded and shielded twisted pair. Choosing a network adapter is largely a matter of testing for compatibility. Will the adapters work with current computers, as well as those planned for purchase? Will they work in the file server as well as the workstations? This may be a good time to call the adapter vendor's technical support line.

Choosing workstation software

The operating system

The computing world is constantly evolving, and this is certainly true for the users of IBM-compatible computers. As of 1993, the most commonly used operating system in the world is IBM or Microsoft DOS, which we will refer to as DOS. Two additions to the operating system, Microsoft Windows and GEOS, are becoming common in school networks, and will be part of your decisions. Finally, IBM's new operating system, OS/2, and Microsoft's latest (Windows NT) may or may not be in your future.

DOS

DOS comes from 3 companies (Microsoft, IBM and Novell) in a variety of versions, from 3.3 to 6.0, but all are rather primitive, unfriendly, character-based operating systems. The redeeming features of DOS are that it is reliable, it is less complex than the alternatives, and it has available the largest selection of software in the world. The goal of most network administrators is to isolate their users from DOS, and this can be done quite effectively. Network users typically choose applications from menus and never have to type in a DOS command. In fact, most student networks do not allow students to interact directly with the operating system. Thus, if you use application programs designed for DOS, you can usually install them and forget about the operating system.

Microsoft Windows

If you don't yet have Windows on your network, you probably will before too long. The advantages of programs which use Windows are too great to ignore. Windows is an addition to DOS which provides extra features to application programs and to the user. The extra user features are essentially a way of utilizing DOS features without typing difficult commands. Since network users seldom need to interact with DOS, these features are not critical and are often bypassed by network administrators. Thus, Windows programs can be chosen from a menu, similar to DOS programs.

Once within the application program, however, the difference is apparent. The interface is graphical, rather than character-based, and is similar for all programs. This is especially nice for desktop publishing, but is also useful for spreadsheets, databases, electronic mail, encyclopedias, multimedia, and just about anything else. Software for Windows is the fastest growing software category, and you can be sure your users will want access to it before too long.

Windows does have a down side. Since its graphical user interface requires much more data handling than any character-based system, it needs fast computers (with lots of memory), fast printers and fast networks. Most computers installed in schools prior to 1991 will not acceptably run Windows software, but virtually all computers being installed now will do so. The authors are presently buying only computers with a 486 processor and 4 - 8 megabytes of memory to take advantage of Windows.

The use of Windows software on educational networks is recent, but growing quickly. Many districts are now acquiring experience in this area, and are solving the problems you will need to overcome. Since educational uses are quite different from business uses, this is an area where you should seek advice from other schools. As an example, most business network users run Windows from a local hard disk to gain maximum performance. That approach poses problems for educational networks, and with reasonably quick workstations it is quite feasible to use Windows without having local hard disks. For network use, only purchase Windows version 3.1 or later.

GEOS

GEOS is a graphical addition to DOS, just like Windows. However, GEOS is much smaller and more efficient than Windows. This makes it very nice for use on educational networks. While the first version did not work on networks, the latest version is quite networkable. There are only a very small number of programs designed for GEOS, but it can be configured as a shell that will directly run any other program that uses the DOS operating system. GEOS will not run programs designed for Microsoft Windows. All in all, GEOS is a wonderful operating system and it runs well on virtually every model of IBM-compatible computer, even XTs and the original PC! Its major shortcoming is that it has a relatively small following and its future seems uncertain. If you like GEOS, definitely use it, but do not count on it to meet all your future needs.

IBM OS/2, Windows NT and others

OS/2 is the wild card in the operating system deck. It requires neither DOS nor Windows, but runs programs written for either. It has its own built-in graphic user interface. As of 1993, OS/2 is much less popular than Windows and few programs have been written for it. However, since IBM seems to be staking its future on the success of OS/2, the situation could change. Used largely by big businesses for in-house developments, OS/2 is still uncommon on educational networks. However, you can be assured that schools will be experimenting and finding the best ways of integrating it with present networks, if it becomes popular in the business world.

Windows NT is a newly released operating system from Microsoft which, like, OS/2, has a graphic user interface built in and does not require DOS. It is large and complex and will probably find its first use as server software rather than on workstations. It is too early to tell if and when it will affect networks in schools. Another new operating system appearing over the horizon will come from a collaboration between IBM and Apple. It will also be very powerful and may make your decision-making even more difficult in the future.

The application programs

When choosing programs for use on a network, it is important to check with the vendor to determine if it is compatible for network use. Often a specific network version is available, with instructions for installation on Netware and ICLAS. Most non-copy-protected programs which are not designed for network use can be made to run on a network with some extra effort. Once again, check with other schools and vendors for advice before investing a lot of your time. Read the usage license provided with the software to determine if this type of use is acceptable. IBM publishes a list of instructions for installing programs on ICLAS networks, and many schools have contributed additions to the list. Educational networks are becoming more common, and the choice of software available for network use is growing quickly – probably much faster than your software budget.

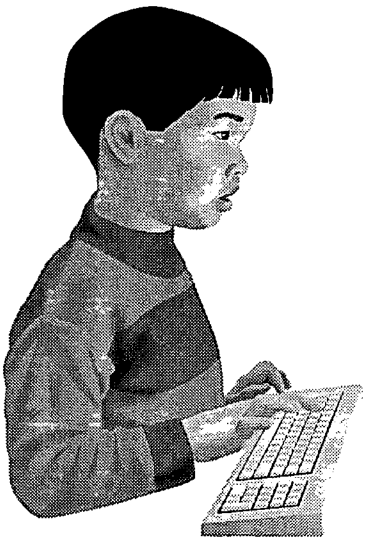

Chapter Nine

Simple networks for IBM-compatible computers

Chapter 9

Simple networks for IBM-compatible computers

Topics covered in this chapter...

- *Two simple network examples*
- *Where to use these networks*
- *A technical description of the networks*
- *Descriptions of what various users do on the networks*
- *A discussion of the potential for future growth*

Introduction

In the following six chapters we will describe examples of specific school-wide networks. The first two chapters use IBM-compatible computers, the next two use Macintosh computers, the next uses a combination of both, and one focuses on Apple II computers. The simple networks are suitable for small schools; the rest are more complex, suitable for larger schools.

This first example in this chapter is a simple Ethernet network running Novell Netware on IBM-compatible computers. The cabling system uses unshielded twisted pair cable installed in a star topology.

Where to use this network

The first example describes a network suitable for a small school (5 - 15 classrooms) that does not have a computer lab. The school is small enough that every computer in the school is within 100 meters of the wiring rack. Even though the plan is for a small school, the network has great expandability, and the most complex network in this book (see Chapter Fourteen) is really just an extension of this simple one.

A similar network could be used in a small school which has a computer lab but does not have computers in the classroom. The second example illustrates the configuration for such a network. Expansion to the classrooms would be straightforward when the need arises.

This network can also be used in a small administrative network. Small hubs can be purchased which connect 4, 8, or 12 computers and sit on a shelf, which eliminates the requirement for a wiring rack.

Description of the network

Diagrams

The first diagram shows the general layout of the cabling system in a small school. This is a conceptual diagram, not a wiring schematic. While the diagram shows wires going directly from the wiring rack to each computer, they would actually be routed in some neat path through the floor or ceiling. The task of routing these wires should be left to a district technician or a vendor.

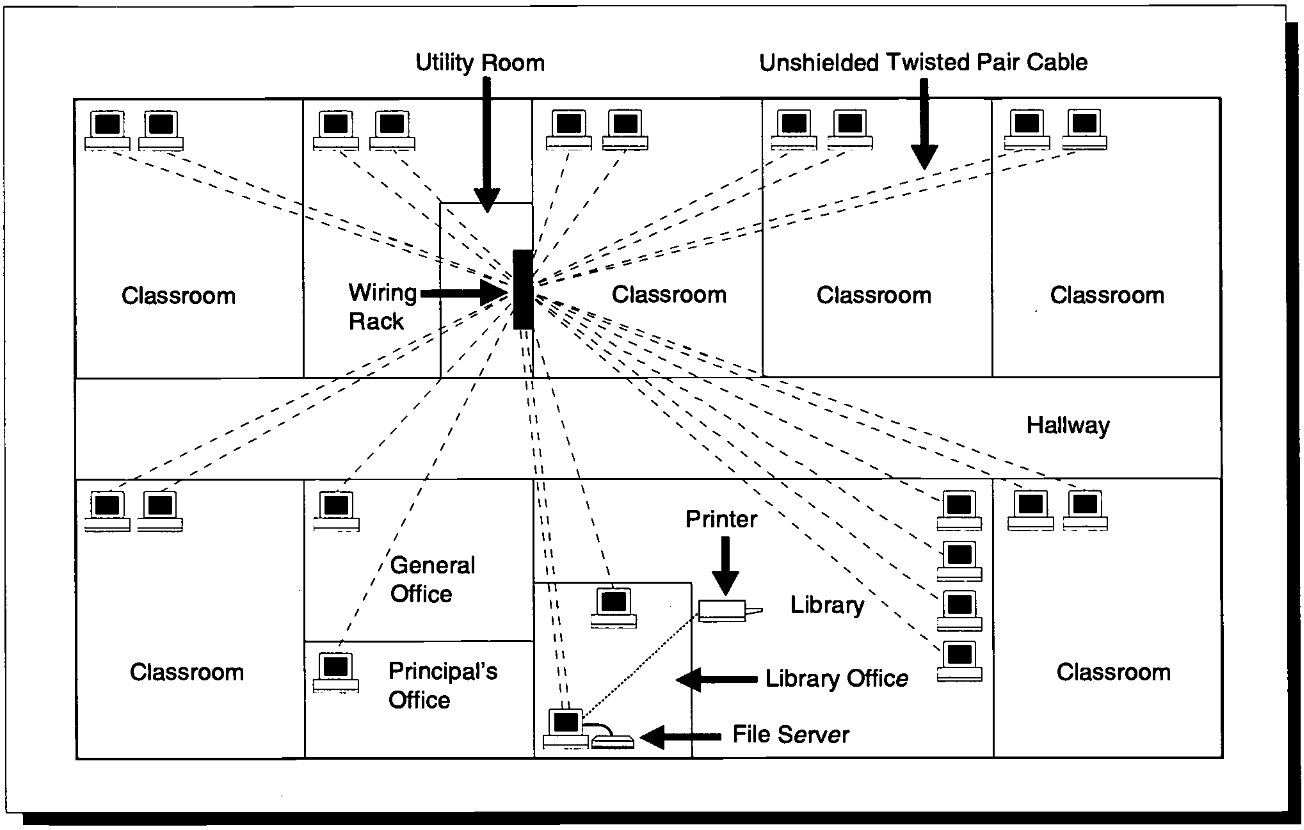

An Ethernet network in a small school. The network is based on a structured wiring system. The wiring rack contains two 10base-T hubs connected by Unshielded Twisted Pair cable to wall jacks in the classrooms. The printer is connected to the parallel port of the file server. Note the two lines running from the wiring rack to the file server. This allows the network to take advantage of the internal routing capabilities of a Novell Netware server. See the next diagram for a more complete description.

The cabling system has unshielded twisted pair cables running from two Ethernet hubs in the wiring rack to each computer (including the file server). There are two workstations in each classroom. There is also a local printer in each classroom, controlled by a switch connected to the parallel ports of both computers. Any student or teacher has a choice of printing on the local printer in the classroom or the laser printer in the library.

The wiring rack could be in any convenient location; in this case a utility room. Other locations could be a furnace room or a janitor's storage room. However, security of the room should be considered, since network problems can arise from accidental or intentional experimentation with the wiring rack. This experimentation could be at the hands of a well-meaning person, who may be unaware of its impact.

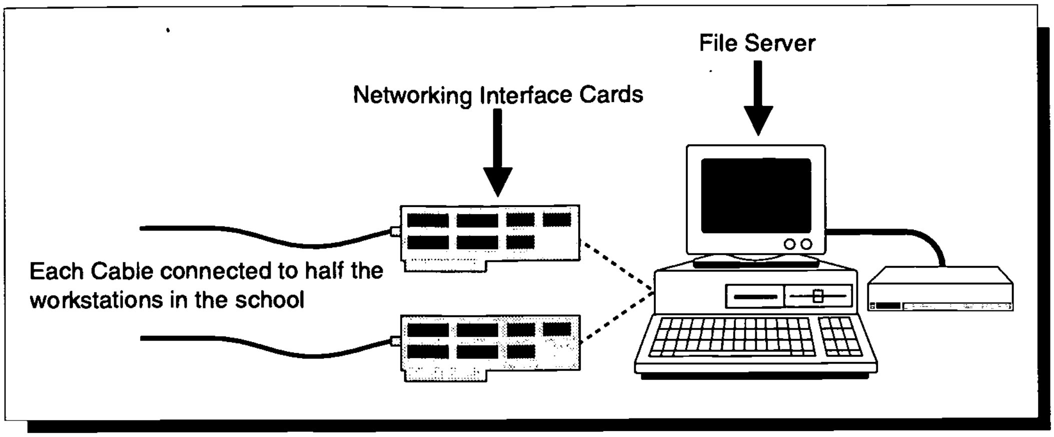

The Novell server for the small school network has two network interface cards inside. Each card is connected to half the workstations in the school. Novell Netware has the capability to isolate network signals so that communications between computers connected to one interface card do not cause delays for computers connected to the other. This is known as internal routing.

The third diagram shows a similar network set up in a single computer lab. Since there are more workstations in this diagram, it would require three Ethernet hubs rather than the two used in the first diagram.

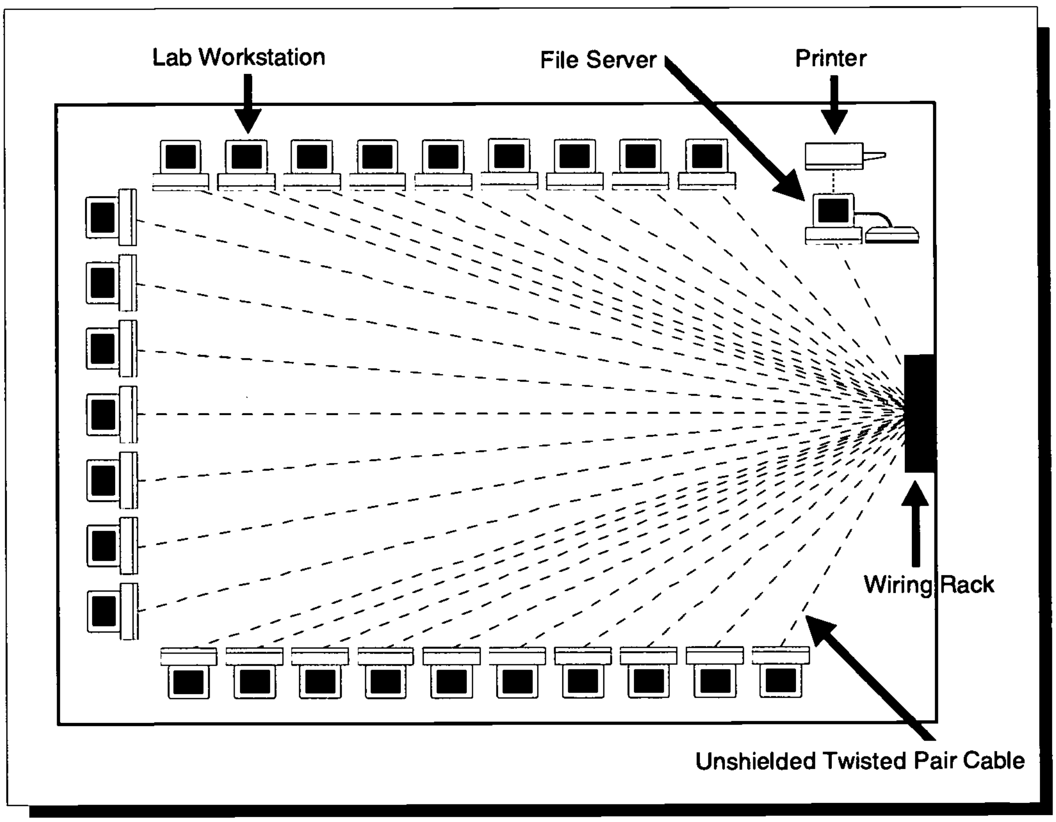

Here is an Ethemet network for a single computer lab also using a structured wiring system. The wiring rack contains three 10base-T hubs. The Unshielded Twisted Pair cabling to the workstations can be run under the tables or over the ceiling. Again, the printer is connected to parallel port on the file server.

The file server

The file server is running Novell Netware and IBM ICLAS. It is connected to a laser printer, through its parallel port. This computer may be left in operation for 10 months at a time, but its keyboard is seldom used. Leaving the file server turned on day and night will be necessary if users are to access it at all hours, and is also easier on the computer. It is in an out-of-the-way location in a utility room, but could be in virtually any location with reasonable security and some access by the network administrator. Since the easiest way to connect a printer to the network is via the server's parallel port, it simplifies matters if the server is within thirty feet of the main printer location.

The file server should be reasonably fast (at least a 33 MHz 486DX processor) and should have a large memory (at least 8 MB, preferably 16) and a large hard disk (at least 200-800 MB). The hard disk should be large enough to hold all the programs and data for the school, and should be as fast as you can afford.

The file server's Ethernet adapter is critical, so it is worthwhile to purchase faster adapters, as available. The performance of the network can also be improved by adding a second network adapter to the file server for internal routing, so at least one extra cable from the server to the wiring rack should be included in your plans.

The workstations

Each workstation has an Ethernet adapter with a remote boot chip, so none of the workstations require diskettes to start. The power needed for the workstation depends on the application software it will be running. An older computer with a 8086 or 80286 processor will run most DOS programs, but a 386 or 486 processor with 4 - 8 MB of memory should be used for Windows software.

If users throughout the school require access to different operating systems, such as DOS, GEOS, Windows and OS/2, ensure that the necessary software is on the network. The remote boot files on the server can be set up so that different workstations start up with different operating environments. In addition, multiple versions of application programs may be needed. For instance, a single network could have Microsoft Word for DOS and Microsoft Word for Windows, with different workstations accessing the appropriate program automatically by a menu choice.

A common problem faced by network administrators is a combination of old and new machines. If the machines have different features, such as different video adapters, extra work may be needed to get programs to take advantage of the capabilities of the new machines. In some cases programs may have to be installed twice, choosing different configurations for different machines. It is also possible to load different drivers for specific machines based upon their network addresses.

The cabling system

An unshielded twisted pair cable with four pairs of wires is routed to each of the computer locations, with at least one extra cable to the server location. Although the diagram shows the lines going straight to the wiring rack, they will be laid out in a neat, organized fashion. They may be in a crawl space under the floor or above the ceiling. Since there are two cables going to each classroom, they can share a single wall plate in the classroom, with two RJ45 modular jacks in the wall plate. A short cable with RJ45 modular plugs then connects each computer to the wall plate.

In the wiring rack, each of the cables is connected to a RJ45 modular jack in a "patch panel." Short "patch cords" with RJ45 modular plugs on each end are used to connect the patch panel to the Ethernet hubs. While there are other ways of connecting the wires in the walls to the hubs, using patch panels and patch cords provides maximum flexibility. If the school had only 12 computers when the wiring was installed, only one hub would be needed; a second hub could be added when more computers are purchased.

The two Ethernet hubs (also called concentrators) mounted in the rack have 12 RJ45 jacks each. Each hub also has a connection for a coaxial cable (thinnet). A short piece of coaxial cable is used to connect the hubs together. The instructions for making this connection come with the hubs. Alternatively, if the server has two network adapters, one adapter can be wired to each of the hubs. This would split the network into two, and thus improve its speed. It is wise to include an extra unshielded twisted pair cable between the rack and the server in the cabling plan to allow future splitting of the network.

The hubs conform to the 10base-T specification, so they will automatically disconnect a failing workstation before it interferes with other computers on the network. The hubs probably have two small lights for each connection; one which shows when a connection is made and another which shows when the computer is transmitting over the network.

What the network administrator does

Since this network uses ICLAS, the network administrator logs on with the user ID "sysop" (short for system operator) and accesses the "sysop menu." This menu includes options to directly access the programs normally used by a "raw Netware" administrator. However, most required tasks can be done from the sysop menu. These include adding new users, adding new programs, and changing the configuration of programs already installed.

Adding users

To add a new user, the administrator simply selects a menu choice, then enters the person's name, status (student or teacher), grade and a password (if desired). If many new users are being added at once, their information can be read in from a file. Since this information is commonly stored on an office administration system, all the users in the school are usually added each September, using a file exported from the office system.

Each fall, the list of users will have to be updated. ICLAS includes a feature to change the grade level of the previous year's students, but the authors have found it easier to erase all the students from the previous year and add all the current students. Since this is done with a file from the office administration system, the procedure is quick and easy.

If desired, other options can be set for each student, such as whether they have a password, how long the password must be, how often it should be changed, a limit on the amount of available student disk space, and whether a user can log-on to two computers at once. It is also possible to set up generic user IDs, which have no passwords and can be used by an entire class at once. If generic IDs are used, students will store their work on floppy diskettes.

Adding programs

Adding new programs to a file server is the most complex task for the administrator. The complexity depends on the program being installed. Programs designed to work on a Novell/ICLAS network will come with simple installation instructions. Many other programs are also easy to install, but there are some which are quite tricky, and others which will not run at all on the network. For complicated installations, the school network administrator may be able to call district support staff for assistance. Other school districts with similar networks are also good sources of help, as is IBM, which has prepared instructions for installing most common school software.

What the teachers do

On an ICLAS network, the teacher logs on with a user ID and a password and sees a menu. One menu option brings the teacher to their personal menu, from which they can use any of the software on the network. Another allows the teacher to work with a class. This class menu has options to add students to the class and make programs available to those students. Thus, the teachers can work only with the students they teach, and can control what programs those students have access to.

Teachers also have access to two special subdirectories for each of their classes which allow them to exchange information with their students. They can give out "electronic handouts" by placing files in one subdirectory and collect "electronic assignments" by having students place their work in the other subdirectory. In the first, students only have rights to read a file; in the second they can place a file in the directory but cannot read or change a file which is already there.

What the students do

When students log-on to any computer on the network with their own ID and password, they will have access to a menu for each of their classes. If a student has only one teacher, only one menu will be available, otherwise there will be one menu for each teacher the student has. Within the class menu, a student can choose any of the software that the teacher has made available.

Each student may be given an area on the server's hard disk to store their own information. This area is private, and can only be seen by the student and the network administrator. The student's data is secure if the passwords are kept secret and the network is backed up frequently. Data is usually the responsibility of the user, unless a hardware failure occurs, so the student may also keep a copy of important data on a floppy diskette.

Potential for future growth

This network has ample room for growth. The wiring rack has room for at least six more Ethernet hubs, servicing up to 72 more workstations. If the network grew to that size, it would probably require more file servers, which could be added easily. Even more growth could be accommodated by adding more wiring racks, in which case the network would be very similar to the example used in Chapter Fourteen. Macintosh computers could also be added to the network as shown in Chapter Fourteen; the addition of computers using UNIX or other operating systems would be accomplished in a similar way.

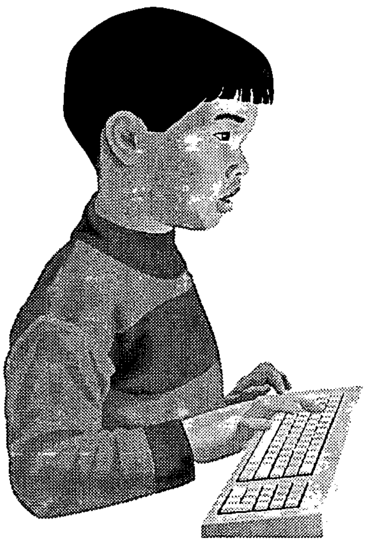

Chapter Ten

A complex network for IBM-compatible computers

Chapter 10

A complex network for IBM-compatible computers

Topics covered in this chapter...

- *Where to use this network*
- *A technical description of this network*
- *Descriptions of what various users do on the network*
- *A discussion of the potential for future growth*

Introduction

This example is more complex than the previous one, and is appropriate for a larger school. Using the cable and networking components normally sold by IBM, it provides for future growth and many connectivity options. The network consists of three labs utilizing IBM Baseband, which are connected together with a Token Ring. There are also a number of computers in classrooms connected directly to the Token Ring.

Where to use this network

A large school

A small school, such as the example in the previous chapter, can be served from a single wiring rack near the center of the school. In a larger school, the distance from the wiring rack to each workstation may be longer than the length recommended for the type of cable used, or it may just waste too much cable. One solution is to install two wiring racks, one at each end of the school. This example has a rack at each end of the school, with Token Ring hubs (called Multistation Access Units or MAUs by IBM) in each of the racks. These racks are then connected together to form one Token Ring network. Token Ring is an excellent choice for use throughout the school because its high speed allows for future growth and the shielded twisted pair cable is not easily affected by electrical interference.

A school with multiple computer labs

This school has three computer labs. Due to the cost of Token Ring components, the computers in the labs are networked with IBM Baseband. Baseband uses unshielded twisted pair cable, and is not recommended for use throughout the classrooms, but is fine for use within each lab.

In a smaller school with multiple computer labs it would still make sense to use a combination of Baseband in the labs and Token Ring throughout the school, but a single wiring rack near the center of the school would suffice.

Description of the network

Diagrams

The diagram on the next page shows the cabling system for the Token Ring network throughout the school. Again, this is a conceptual diagram, not a wiring schematic, and the cables would be laid out in a neat pattern in the floors or ceilings. The wiring rack at the top of the diagram is connected with shielded twisted pair cable to the file servers in two of the labs, the library, and the utility room. It is also connected to 24 workstations. The wiring rack in the bottom of the diagram is connected to the file server in the third computer lab, the file server in the office, and 22 workstations in classrooms and offices. On the following page, the second diagram shows how the Token Ring hubs are connected together.

The third diagram shows the cabling for one of the computer labs. The file server has two adapters: one for Token Ring and one for Baseband. The Token Ring adapter is wired with shielded twisted pair cable to one of the hubs in the wiring rack. The Baseband adapter in the file server is wired with unshielded twisted pair cable to a Baseband hub (extender) in the lab. A Baseband hub has eight connections, but each can connect to a daisy chain of up to eight workstations. Thus, one connection is used for the server, and other connections are used for short chains of workstations.

The file servers

There are six file servers in this school: three for computer labs, one in the library, one in the utility room, and one in the office. Each will have a large hard drive and one or more printers. The file servers in the computer labs have a Token Ring adapter and a Baseband adapter, while the file servers in the office, utility room, and library only have Token Ring adapters.

Each file server will have a different set of software and a different set of users. The server in the office will have no student users. The server in the library may have a user called “Catalog” for those who only want to use the library catalogue. The Business Education lab server will have Business Education software and Business Education users; the Humanities lab server will have English and Social Studies software and users from those courses; the server in the drafting lab will have CAD software and users taking technology education courses. Thus, each server will provide access to different resources, but will be available to users throughout the school. The server in the utility room will be a general purpose server, with software used in classrooms throughout the school.

The workstations

The only requirements for the workstations is that they be compatible with Novell Netware and that each have a network adapter. The computers in the labs will have IBM Baseband adapters; all the others will have Token Ring adapters. All the adapters should be capable of remote booting, so diskettes are not required.

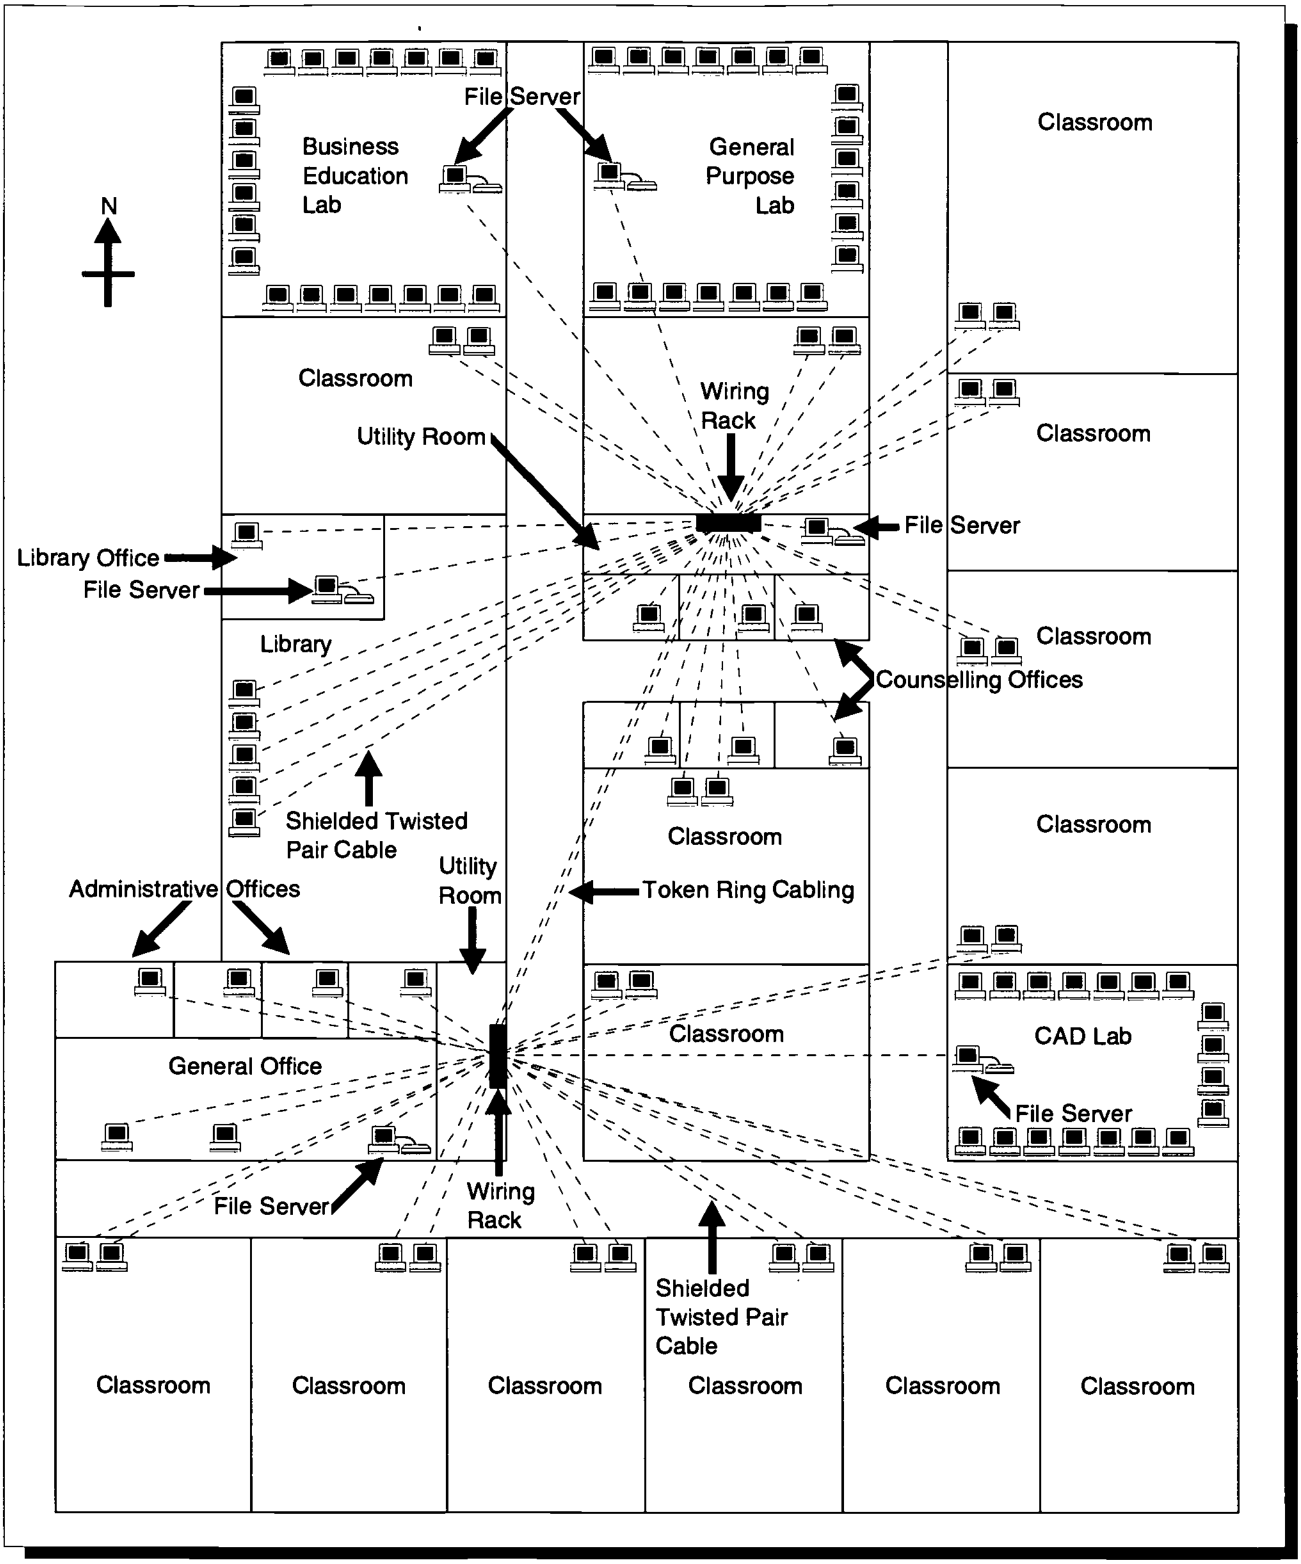

This large-school network is based on a Token Ring system that connects the labs and classrooms using Shielded Twisted Pair cable. The Token Ring network can handle a large volume of communication without seriously degrading the response time to individual workstations. See the next page for more detail on connecting the Token Ring cabling. The labs in this school use IBM Baseband networks, running on Unshielded Twisted Pair cable. See the next page for detail on the configuration of the networks in the labs. Each of the six file servers run Novell Netware. The servers in the labs contain two network interface cards so communications in the lab targeted for the local server do not affect any of the computers elsewhere in the school.

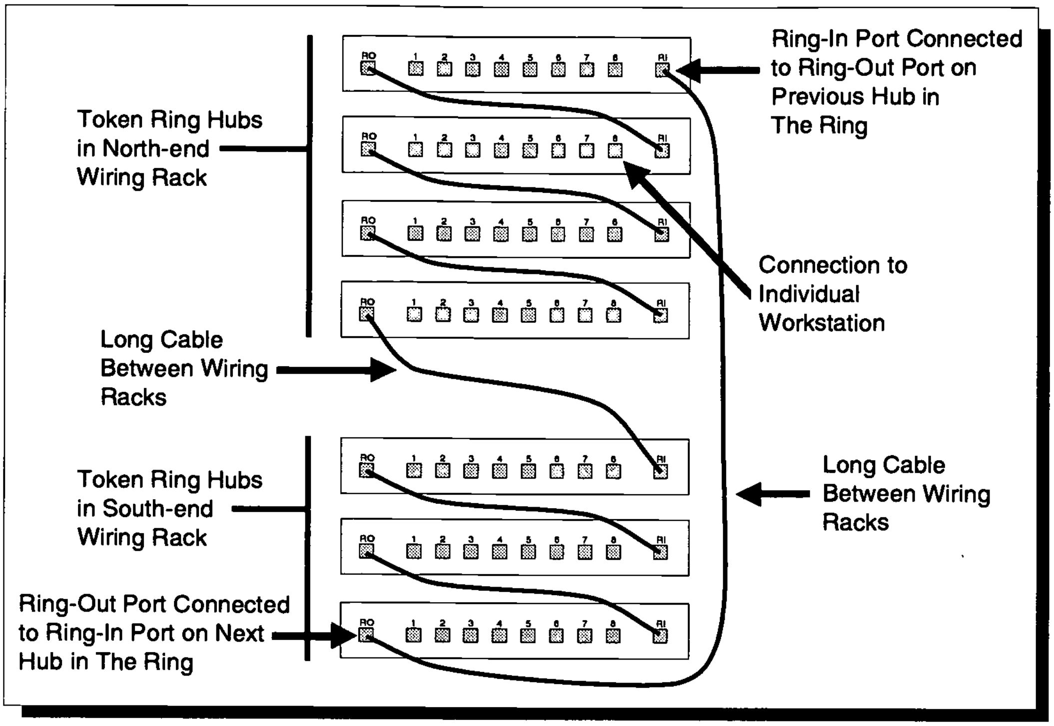

The circular nature of Token Ring wiring is not always obvious when all of the hubs are sitting in the wiring rack, especially when there are two or more wiring racks in different locations. Here, all of the hubs are shown together without the racks to help you see the endless loop of cable that creates the Token Ring.

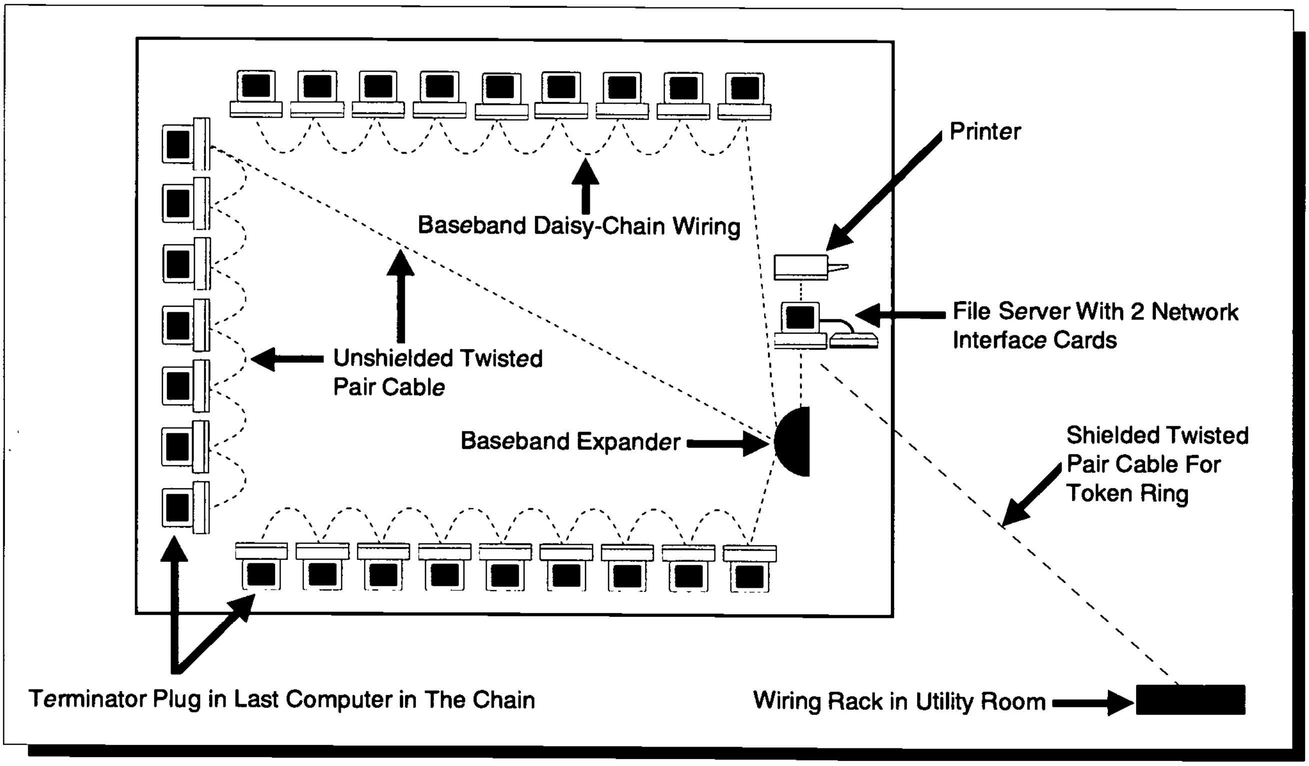

IBM Baseband configuration for labs. Baseband is not suitable for networks larger than a single room. The file server acts as a router to isolate local lab communications from the rest of the school-wide network. The printer is connected to the parallel port on the file server.

If there are a variety of machines, grouping similar machines together makes sense. For instance, the computers in the CAD lab may have 80486 processors and math coprocessors; the computers in the other labs may have 80386 processors; and the computers in classrooms may be older machines with 8086 or 80286 processors. While any recent purchases would have 486 processors and 4 - 8 MB of memory, the older machines can still be used productively for many tasks.

The cabling system

There are actually four cabling systems in this school: one in each of the computer labs and one throughout the rest of the school.

The cabling in the labs is unshielded twisted pair cable linking up to eight computers in a daisy chain and then connecting them to the IBM Baseband hub (called an extender by IBM). The cables have RJ11 modular plugs on each end. Each Baseband adapter has two modular jacks; one to take the signal from the hub and another to lead to the next computer in the chain. The last computer in the chain must have a special plug with a terminating resistor in the unused jack. The Baseband network is easy and inexpensive to set up and is usually trouble-free, but is not suitable for connecting all the classrooms together.

The Token Ring network throughout the rest of the school uses shielded twisted pair cable, also called IBM Type 1 cable. This cable is much bulkier and more expensive than unshielded twisted pair cable, and will require a professional installer. It may be hidden in the walls, under the floors, and in the ceilings. The only part of this hidden cabling the network administrator should have to think about is the special IBM wall jacks in the classrooms and on the wiring racks. Short cables called patch cords are used to connect the computers to the wall jacks and to connect the jacks in the wiring racks to the hubs (MAUs). The connections and cables are similar to the Ethernet system used in the previous chapter, except the cable and all connections are larger and more robust.

It should be noted that you could also use a very similar cabling system with unshielded twisted pair cable. However, one of the advantages of using the shielded twisted pair cable is that it can serve multiple purposes. If you go to the added expense of using shielded cable, you should definitely use cable that includes an extra pair of unshielded wires for a telephone. You should also plan on adding the special connectors which allow you to carry video signals over the same cable, as discussed below.

Each of the Token Ring hubs has eight connections for the wires leading to the computers. The following table shows how to calculate the number of hubs needed in each of the wiring racks:

	File server connections	Workstation connections	Total connections	Hubs required
Utility Room North	4	23	27	27/8 = 4
Utility Room South	2	22	24	24/8 = 3
Business Ed Lab	1	24	25	25/8 = 4
GP Lab	1	30	31	31/8 = 4
CAD Lab	1	20	21	21/8 = 3
Total	9	119	128	18

Each hub has two extra connections labelled ring in (RI) and ring out (RO). These are for wires leading to other hubs. The second diagram shows how these connections are used to connect all the hubs together. If you stretched out all the cables in the diagram you would see that the hubs are actually connected in a ring, hence the name Token Ring.

The two cables which connect the one wiring rack to the other are quite long. In the diagram of the whole school, you can see that they extend over half the length of the school. If the school were very long it would be necessary to use fiber optic cables for these two long runs. This would require special adapters to connect the fiber optic cables to the hubs.

A complex cabling system like this must be installed by experienced professionals. The vendor should not only install the cable, but also set up and connect the file servers and workstations. Technical support staff employed by the district may be able to do part of the installation, but the vendor must have the necessary expertise and be prepared to step in wherever required.

What the network administrator does

The tasks of the network administrator are similar to those described in the previous chapter, with the added complication of multiple servers. In almost every case the network administrator, like the users, works with only one file server at a time, setting up the programs and users on that server. It is quite simple for a user to log-on to more than one server at a time, but that is seldom of value. The one instance in which it is useful to log-on to two servers is when the user wishes to use a program that is stored on one server and print on a printer that is connected to a different server. To make this possible, the network administrator can create a menu choice which allows users to choose printers from a different server.

In a school of this size, it is possible that the computer labs will have different network administrators than the person responsible for the school-wide system. This can work well, but these administrators must have clear communication and work on some tasks as a team. It is best to have one person assigned the role of "super administrator" to deal with issues affecting the whole school.

What the teachers do

To a teacher, each of the file servers looks identical to the simpler example in the previous chapter, with the exception that before logging on, the teacher can choose which server to use. Each workstation has a "default" file server, which the user will log-on to if a different server is not requested. If a different server is wanted the user "passes through" to it by choosing from a list. Some teachers will only use one file server; others will use programs stored on a variety of servers and have classes set up on a variety of servers. Happily, once a server is selected, the operation of this network is no different from a simpler network.

What the students do

In most instances the students will not even know that the various servers are connected. The connection is only needed when a student wants access to resources on a different server, such as when a student in a classroom wants to use a program on the Business Education server, or when a student in the CAD Lab wants access to a file saved on the General Purpose server. All they need to do is pass through to the appropriate server before logging on.

Potential for future growth

This network has potential for growth beyond the previous example. The network could be expanded to support hundreds of workstations without causing long delays for users. If the 16 megabits-per-second speed of Token Ring is too slow for some multimedia applications, the shielded twisted pair cable in this school can support 100 megabits per second using a version of FDDI adapted for this cable.

The shielded twisted pair cable can also transmit video signals at the same time as the computer data. With a splitting device called an F-Coupler, available from IBM, the same cabling system can be used to send video signals such as cable TV, satellite TV, or multimedia programs, to workstations. The use of such a cabling system along with a centralized media system may do for the VCR what computer networks are doing for the stand-alone computer. The educational applications for this technology are still in their infancy, but the capability is there.

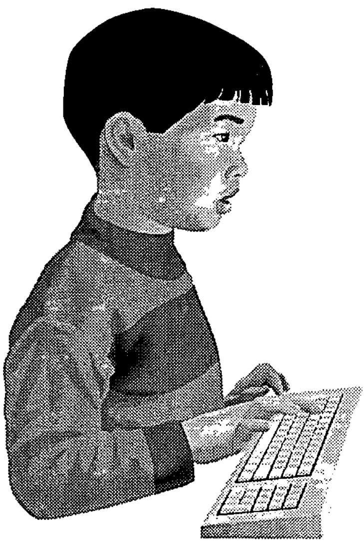

Chapter Eleven

Technical issues in Macintosh networking

Chapter 11

Technical issues in Macintosh networking

Topics covered in this chapter...

- *The power of AppleTalk*
- *Decisions to be made on Macintosh networks including:*
  - *Choosing the right network cabling*
  - *Choosing the network software*
  - *Choosing and installing the workstation software*

Introduction

The little computer that thought it could...

The Macintosh computer was introduced in January 1984. It was a radical departure from what computing had been up to that point. Computing had always been extremely technical, involving complex commands entered exactly according to a defined sequence. Consequently, computer use was limited to a relatively small group of technically and mathematically oriented people. The Mac, however, was based on a new philosophy, the notion that the computer was a wonderful productivity tool for virtually everyone. As the computer was applied to an ever-widening circle of tasks, users would be much more interested in the tasks than in the computer itself. To facilitate this focus on productivity for non-technical people, the computer had to be transparent and its use intuitive.

The Macintosh introduced a graphical interface that simulated the top of a person's desk. For example, files were stored in folders that looked like those you put in an actual filing cabinet. Commands were communicated to the computer by manipulating or moving graphical items in a way that mimicked the actions that people normally took on their desks, instead of typing complicated text commands. To delete a file, for example, the user dragged the picture of the file into the trash can, using a mouse. In fact, almost all commands can be communicated to a Mac without ever using a keyboard. This combination of easily recognized pictures and the ability to enter commands without using a keyboard proved very effective in making computer use intuitive and accessible to non-technical people.

For a number of years, only Mac fanatics seemed to be excited about this new paradigm. However, the power of this approach has now been widely recognized and the idea is driving interface design throughout the entire computer industry. OS/2, Windows, and GEOS are examples for IBM-compatible computers. A/UX for Macintosh computers and the operating system for the NeXT computer are examples for the UNIX operating system. From a very modest beginning, the Mac idea has emerged as one of the most significant advances in computing.

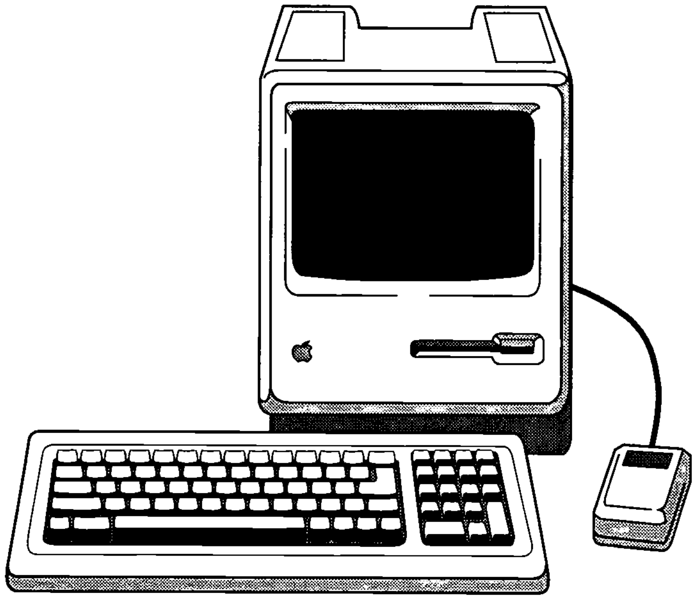

You've heard that big things can come in little packages – this little computer certainly fits that description! Based on a graphics-only, easy-to-use operating system, the Mac dramatically changed the direction of computing forever.

The power of AppleTalk

AppleTalk – powerful? The answer is definitely "Yes!" The AppleTalk networking system was a clever idea right from the outset. Instead of starting with the hardware, Apple Computer Inc. started with the design of how Mac computers would talk to each other. This design established a set of rules or protocols for network communication. This set of protocols became known as AppleTalk. It is hardware-independent, meaning it can be implemented on whatever cabling system you want.

There are four main reasons for the success of the AppleTalk system. First, every Mac computer comes equipped with the ability to operate on a basic low-speed AppleTalk network known as LocalTalk. This makes an AppleTalk network substantially less expensive to install than any other network for any other make of computer. There are no network interfaces to purchase because the interface is built-in and connected to the printer port of each Mac. All you have to do is buy network cabling connectors and the cabling and you can set up a fully functional network! This can be done for as little as $20 per workstation.

Second, LocalTalk networks are easy to install. They are designed to be "plug and play" systems. That means all you have to do is plug the network connectors into the Macs, plug the cables into the connectors, and the network is ready to begin carrying messages. It is a remarkably simple and easy to use system. Traditionally, networks of computers were very complex, so much so that highly trained specialists were required to set them up and keep them running. Then the Mac came along with built-in LocalTalk. LocalTalk did to networking what the Mac did to personal computing – made it much more accessible to non-experts. This was a significant advance in networking.

Third, following the philosophy behind the design of the Mac operating system, network application software for the Mac such as AppleShare, MacServe, TOPS, and MacJanet were designed to be significantly easier to use than anything else people had ever seen.

Fourth, AppleTalk networking is upwardly compatible with higher speed networking hardware. You can install AppleTalk networks that run on Ethernet hardware, token rings, and ultra-high speed fiber optic cabling systems.

The introduction of LocalTalk resulted in a substantial increase in the number of microcomputer networks in business and education. More importantly, these new networks were being installed and maintained by relative newcomers. The ease of both the hardware and the software components of LocalTalk networks set a new standard for administrator and user friendliness which has since been driving the development of new network systems on virtually every other make of microcomputer. AppleTalk powerful? You bet!

Technical issues on a non-technical computer

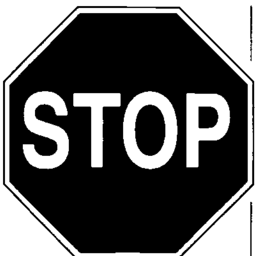

The Mac is not only easier for the average user, but also for network administration as well. Preparing boot disks for workstations, installing network server software, and most other administrative tasks are easier, because complex typed-in commands are not required. However, underneath the easy-to-use graphical interface, the Mac is just as technically demanding as any other computer. Thus, while accomplishing a specific task may be easier on the Mac, the network administrator for a Macintosh network must have the same conceptual understanding as any other network administrator. If you are new to networking and you haven't read the first two chapters of this book because you didn't think Mac users had to know all that information, you were wrong. Go back now and make sure that you understand the conceptual foundation we laid regarding networks, operating systems, network administration, and network planning – you will need it.

Using the computer is easy – networking it can be difficult

While we have said that the design of the Mac makes the specific jobs of network administration easier, the planning and implementation of Macintosh networks is often more difficult than doing the same with IBM-compatible computers. There are three main reason for this. First, the very fact that the Mac is easier to use, combined with its built-in networking capability, often leads to the expectation by school and district administrators that anyone can put a Mac network together. Thus, requests from classroom teachers for outside help can be met with disbelief. If you are installing an IBM-compatible network, the need for outside expert help is assumed (if you don't believe us, just read Chapter Eight!). But if you're putting in a Mac network, you would be surprised at how often the response is, "just read the instructions and everything will be just fine – after all, it's a Mac."

Second, as mentioned above, a Macintosh network is just as complex as any other network and the same level of understanding of computing and networking concepts is required to plan, install, and maintain one, especially if you begin to make internetwork connections. It may actually be harder with Macs because MacJanet (one of the best networking programs for schools) does not automate many of the administrative tasks that are transparent when using Novell network software on IBM compatibles. Don't worry, we'll help make the use of this product relatively straightforward later in this chapter.

Third, the slowest IBM-compatible network (Arcnet) is eight times faster than the Mac's built-in networking. This forces Mac network administrators to make decisions and deal with configurations that are simply not encountered in IBM-compatible networking, especially when attempts are made to speed up Mac networks.

Decisions to be made for Mac networks

AppleTalk networks are made up of the same basic components as any other network: cables, hardware, and software. The cables connect the computers so messages can travel between them. The hardware is the Mac itself which transmits and receives the information. The software has two components. First, there are the instructions for how data is sent and received. These instructions determine the protocols for network communication. They are burned onto a chip which is placed inside every Mac computer. The second component of the software is the network application program that provides the features of a network (file sharing, print spooling, electronic mail, etc.). This software is chosen by the user; we will help you with that decision later in this chapter.

The key to success is planning, both in the short term and in the long term. Long-term planning is especially important because the speed factor of Apple's built-in networking quickly becomes a major problem with larger networks. You must carefully consider all options for selecting a cabling system when developing your long-term networking plan.

There are three main decisions to make when planning a Mac network – what cabling system you will use, what network application software you will run on that cabling, and how you will configure the software on the workstation boot disks. Let's discuss each of these.

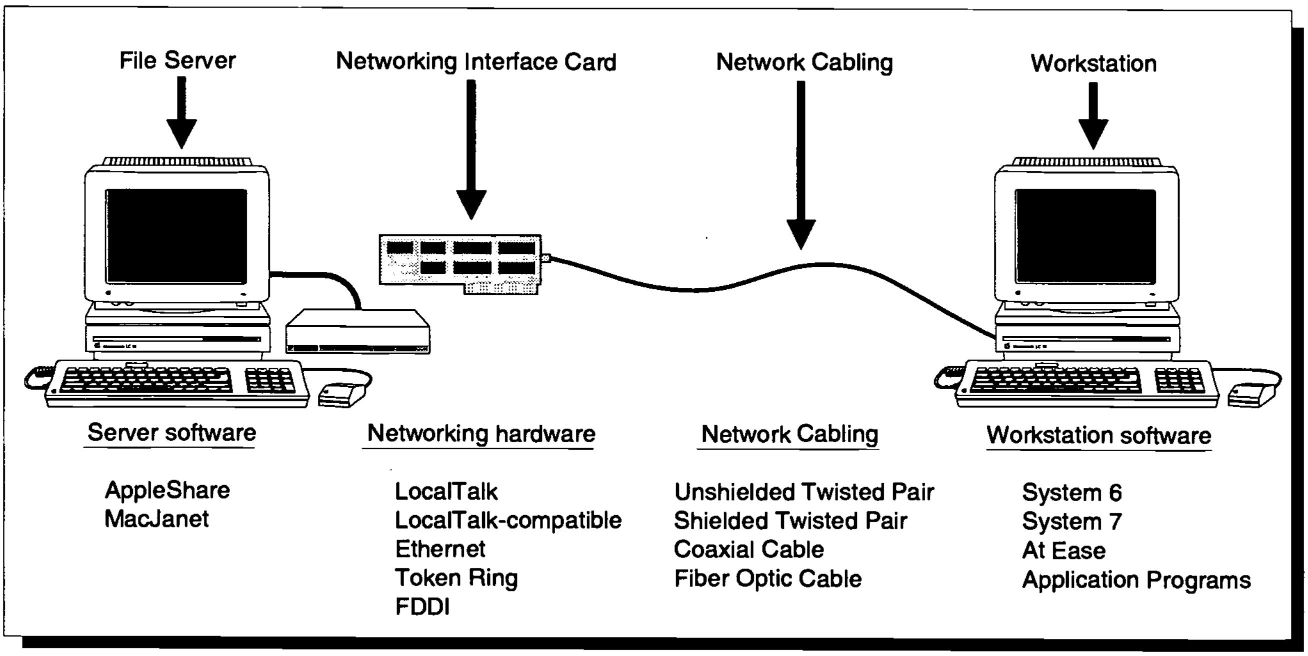

Here are the four major decisions to be made when designing any network. The various options for networking Macintosh computers are listed below each heading.

Choosing the right network cabling

Selecting a cabling configuration is a critical decision for Mac networks. Don't get a false sense of security from advertising that promotes the Mac as having "plug and play" networking (plug and play means that you simply plug it in and it works). While it is true that the Mac's built-in networking is the simplest available for microcomputers, there are still many factors to consider. If you decide to use the Mac's built-in networking, you should only do so after pondering the long-term implications of that decision. Let's discuss the various options for network cabling. Note: we will discuss network hardware and cabling systems together in this section. Network cabling and hardware systems for Mac networks fall into five categories – LocalTalk, LocalTalk-compatible, Ethernet, Token Ring, and fiber optic cabling.

LocalTalk

Every Mac is equipped with the logic required to communicate on a LocalTalk network. LocalTalk was originally designed as a print sharing network system. Consequently, LocalTalk has a relatively slow communication speed of 234,000 bits per second (that doesn't seem slow, but IBM-compatible networks start at 2,000,000 bits per second). However, LocalTalk employs a particularly good implementation of the collision avoidance communication strategy that makes up for some of the slowness. Nevertheless, speed is a problem on LocalTalk networks and we will provide some solutions as our discussion of Mac networking unfolds. The popularity of LocalTalk networks is due to the minimal cost and relative ease of installation.

The relatively straightforward components of the LocalTalk system are:

1. The instructions for using a LocalTalk network which are contained on a chip inside each Mac. No additional interfaces are required to connect a Mac to a LocalTalk network.
2. The cabling between the computers on the network. There are several cabling systems and several topologies, all of which will be described more completely below.
3. The connectors which attach a computer to the cabling. Connectors are always connected via the printer port on the back of the computer. Again, there are several options for Mac networks.
4. The terminators at the ends of the network cabling. Terminators are simply resistors which absorb the signals being sent over the cabling, to prevent their reflection back down the wires from the end of the cabling.

The LocalTalk networking system was designed by the engineers at Apple Computer Inc. in the early 1980s. Shortly after the introduction of the Mac, other companies began to make LocalTalk-compatible network systems.

LocalTalk is the name Apple Computer Inc. gives to its own cabling system. To make its networking system as simple as possible, Apple designed LocalTalk connectors to be daisy-chained together. To add a new computer to the network all you have to do is insert a new cable into the connector on the last computer on the network, run that cable to the new computer, and insert it into a connector attached to the new computer's printer port. The connectors are designed with two cable connections (called ports) to facilitate this daisy-chain topology. The following diagram illustrates the daisy-chain topology of LocalTalk cabling.

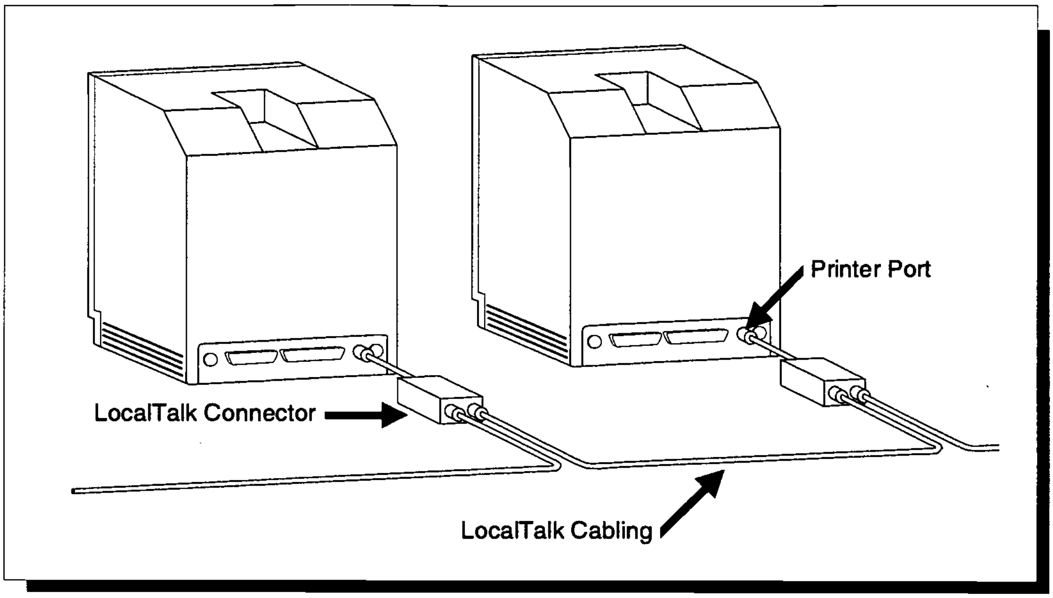

Apple Computer's LocalTalk cabling is configured as a daisy-chain. The last connector on the chain is left with an empty port because the connectors are self-terminating.

The connectors are also designed to be self-terminating. That means that each connector contains a terminating resistor that operates as long as one of the connector's cable ports is empty. The only two connectors that would have a port empty are the connectors for the computers at either end of the network. Thus, the terminators for the two end connectors would absorb the signal to prevent reflection back down the cabling. The terminator in a LocalTalk connector becomes inactive as soon as both ports on the connector are used. When a new computer is added to the network, the cable connecting the new computer automatically disengages the terminator in the second-to-last computer and the terminator in the connector for the new computer becomes active because, as the last computer on the network, its connector has one empty port. Make sure that you never connect all of the connectors together to form an unterminated circular network.

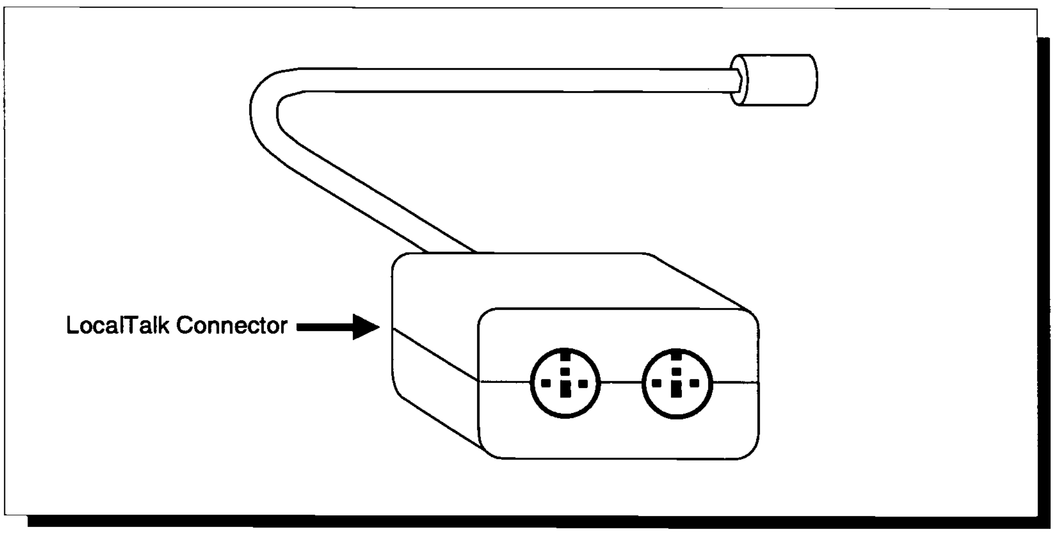

A LocalTalk connector. These connectors are easy to install and create the only true "plug and play" network available today.

LocalTalk cabling is not the most widely used system for Mac networks, for two main reasons. First, LocalTalk cabling is more expensive than the phone cabling systems described in the next section. Second, the daisy-chain topology of LocalTalk networks means that a bad connection to a connector's port or a faulty connector causes all computers further down the line to be cut off from the server. To prevent this problem you must go to a bus or trunkline topology, which is not possible with LocalTalk cabling. Third, LocalTalk is not practical in multi-room network installations.

LocalTalk-compatible network systems

PhoneNet, from Farallon Computing, was the first alternative to LocalTalk. PhoneNet is a substitute for LocalTalk cabling and connectors and, as such, runs at exactly the same speed as LocalTalk. The design is based on the fact that telephone cabling contains four wires. The phones only use two of these wires, so Farallon decided to make connectors that would insert into the existing telephone wall jacks (RJ11 jacks) in offices and send network communications over the two unused wires in the telephone cabling. If the location of the network did not have existing cabling, then simple unshielded twisted pair wire could be used, which was much cheaper than LocalTalk cabling. Farallon also made connectors that were substantially less expensive than Apple's. Most significantly, the PhoneNet system made it possible to use a trunkline or bus topology for the network, overcoming this limitation in LocalTalk. PhoneNet was so successful that it has become the most popular network system for the Mac. Several other companies have introduced similar network systems. One is the "Turbonet ST" system from Focus Enhancements Inc. The Turbonet ST system is completely interchangeable with PhoneNet equipment, but it is substantially less expensive. While Farallon still dominates the LocalTalk-compatible networking market, newer entries like the Turbonet ST system are having a significant impact on networking hardware sales. Long live competition!

LocalTalk, PhoneNet, and the other LocalTalk-compatible cabling systems operate identically. The only differences between a LocalTalk network and a LocalTalk-compatible network are the connectors and the cabling connecting the computers. Both networks rely on the LocalTalk chip inside the Mac. Thus, LocalTalk and phone wire-based cabling systems are both referred to as LocalTalk networks. This is in contrast to Ethernet, Token Ring, and fiber optic networks which do not rely completely on the LocalTalk chip inside the Mac.

Most schools do not have existing telephone cabling running to each classroom and they certainly don't have 24 to 30 telephone wall jacks in a single classroom to facilitate the connection of a lab of computers. This means that labs and schools must be wired with unshielded twisted pair wire to run PhoneNet or similar LocalTalk-compatible network systems. This is generally not a major problem because the cost of the wiring is so low. There are three topologies that can be used with phone wiring-based networks. The first is a simple daisy-chain. Each connector comes with a six-foot length of unshielded twisted pair wire that can be used to join two computers together. The connectors for the computers at either end of the network must have a terminator inserted into the empty port. Usually, each connector you buy comes with a terminator in the package, but check this in case terminators must be purchased separately for the particular cabling system you are using. Do not create a circular network by hooking all of the connectors together; always have two terminated ends.

Daisy-chain topologies are acceptable if your network will have 10 or fewer computers. Once you get above this number, a trunkline topology is more efficient. No matter how many computers you will have on the network, if you plan to have over 1000 feet of wiring, you will need a solid core trunkline to increase the distance your

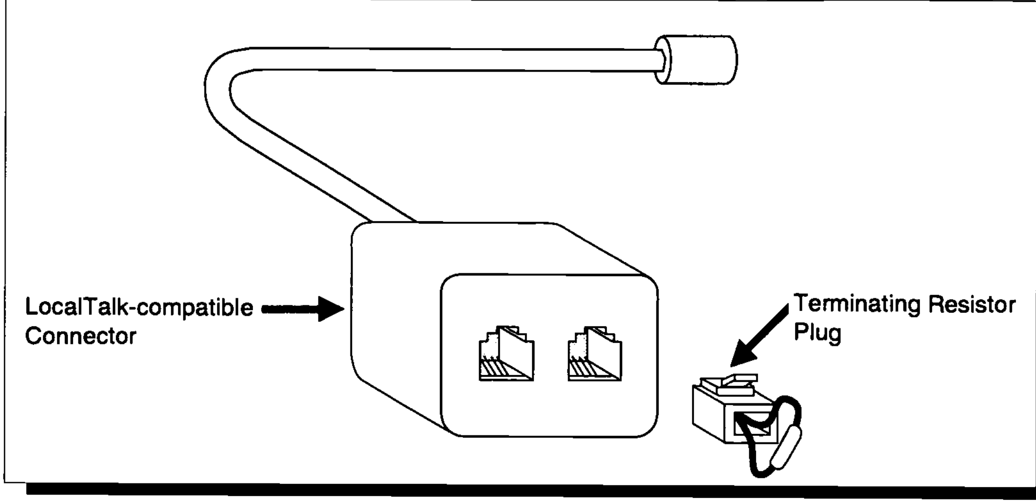

A LocalTalk-compatible connector with a terminating resistor. With the exception of placing the terminators at each end of the network, LocalTalk-compatible systems can be as easy to use as the original LocalTalk system.

signal can travel without becoming too weak. Solid core trunklines, which are much less expensive than signal-boosting equipment, increase the signal distance to 3000 – 4000 feet depending on the thickness of the cabling.

The second topology possible with phone wire-based networks is the trunkline or bus, which is recommended for two main reasons. First, unlike a daisy-chain, the trunkline allows network signals to continue past connectors that have malfunctioned. Second, as mentioned above, solid core copper wire increases the distance an unboosted network signal can travel. You connect RJ11 telephone wall jacks to the trunkline without cutting the solid core copper wire (which would decrease the distance signals can acceptably travel on the wire). Terminators are wired inside the wall jacks at both ends, thus eliminating the possibility that someone could mistakenly or mischievously remove a terminator from the network. Most network administrators will need to get some outside help to wire a trunkline. Just make sure that the help you get has had experience with installing computer networks. Not all electricians understand the requirements of computer networking.

Once the trunkline is installed, you connect the Macs to the network by inserting one end of the six-foot twisted pair wire that comes with each connector into the wall jack on the trunkline and inserting the other end into one of the two ports on the connector. The round DIN 8 jack on the short cable extending from the other end of the connector is inserted into the printer port on the back of the Mac. The other port on the connector is left empty. *DO NOT* place a terminator at each computer; they will absorb so much signal that your network will not operate at all.

The third topology possible with phone wire-based networks is the star. Many network administrators inadvertently create a star topology when they daisy-chain two or more computers off a single wall jack. This daisy-chained line off the trunkline is called a passive star, because the signal is unboosted. This should be avoided. Passive stars of unshielded twisted pair wire greatly decrease the efficiency of the network. Furthermore, the passive star is often terminated because that is what you are supposed to do with a daisy-chain topology. The problem is that the network already has the two terminators it needs and the extra one absorbs more signal than the network can tolerate. Unfortunately, these networks will operate, but problems with speed and

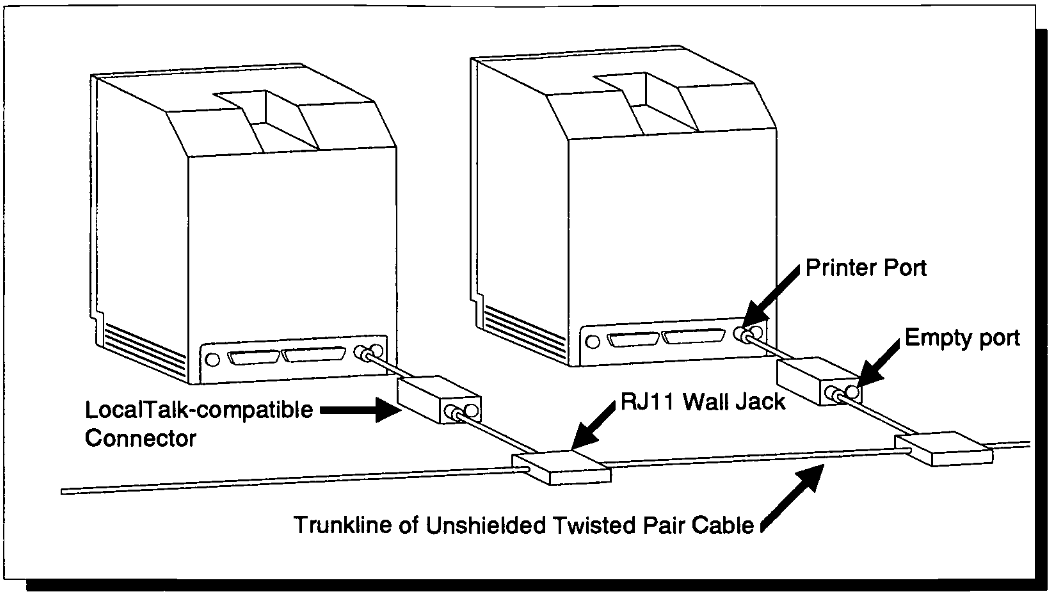

LocalTalk-compatible cabling systems are designed to run on a trunkline configuration. Notice that signals can still reach other workstations even if a connector fails.

transmission errors will result. If the network is ever expanded, the passive star can, suddenly, render a network inoperative A star topology can be used with phone wire-based cabling systems as long as a repeater is placed at the beginning of each arm that connects to the trunkline, or if a star controller is placed at the center of a true star topology. The repeater and/or the star controller boost the signals travelling on the network and ensure that the network will operate efficiently. They also greatly increase the permissible length of a network's cabling. If you plan to create a network that will have three or more arms, then an active star topology with a star controller is the recommended configuration.

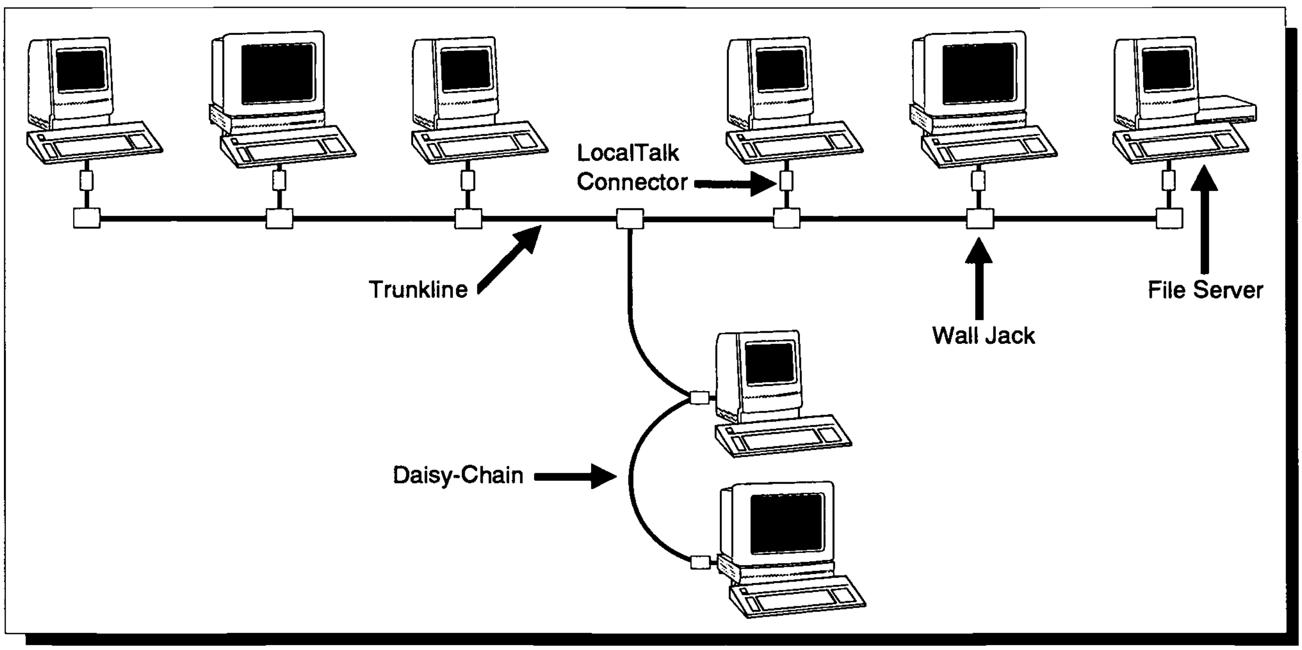

This trunkline has an unboosted daisy-chain arm branching off from a wall jack. This is called a passive star and should be avoided to ensure that network signals do not fade due to over- or under-termination, or excessive cable length.

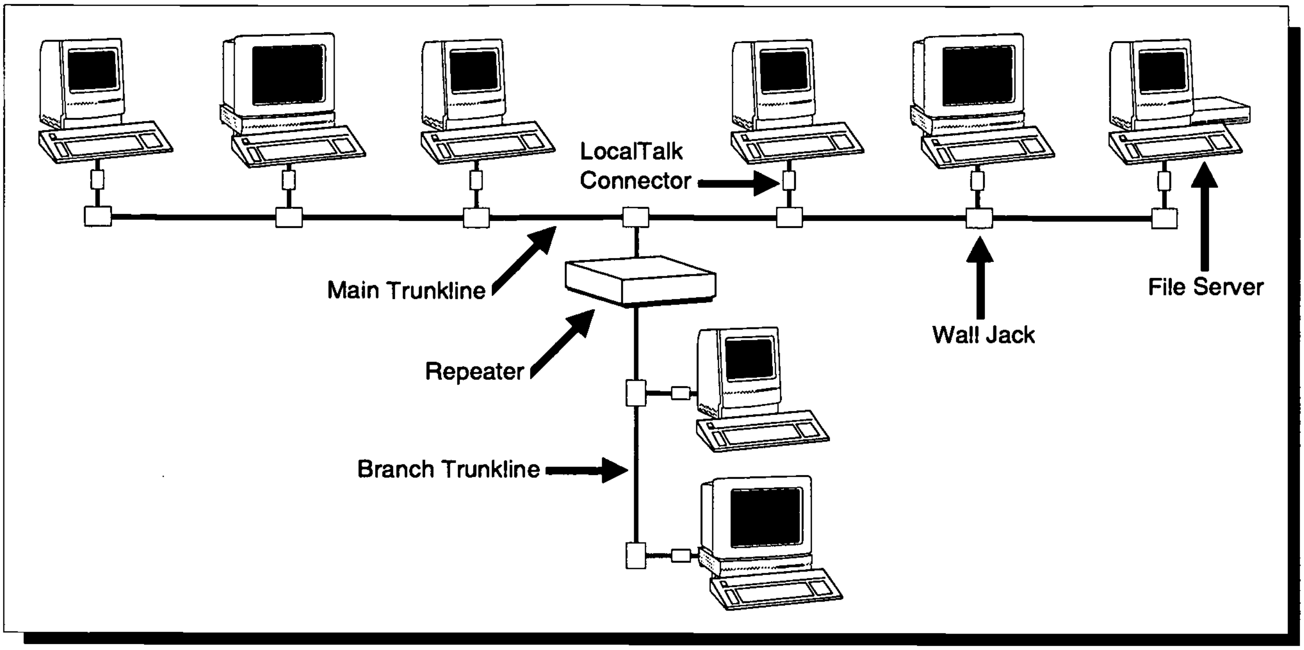

When you must create a new branch off the main trunkline, it is very easy to exceed the permissible length of network cabling. To avoid this, place a repeater at the beginning of branch to ensure that signals do not fade. The repeater does not create a new zone – it simply amplifies a single network's signals to make them travel greater distances.

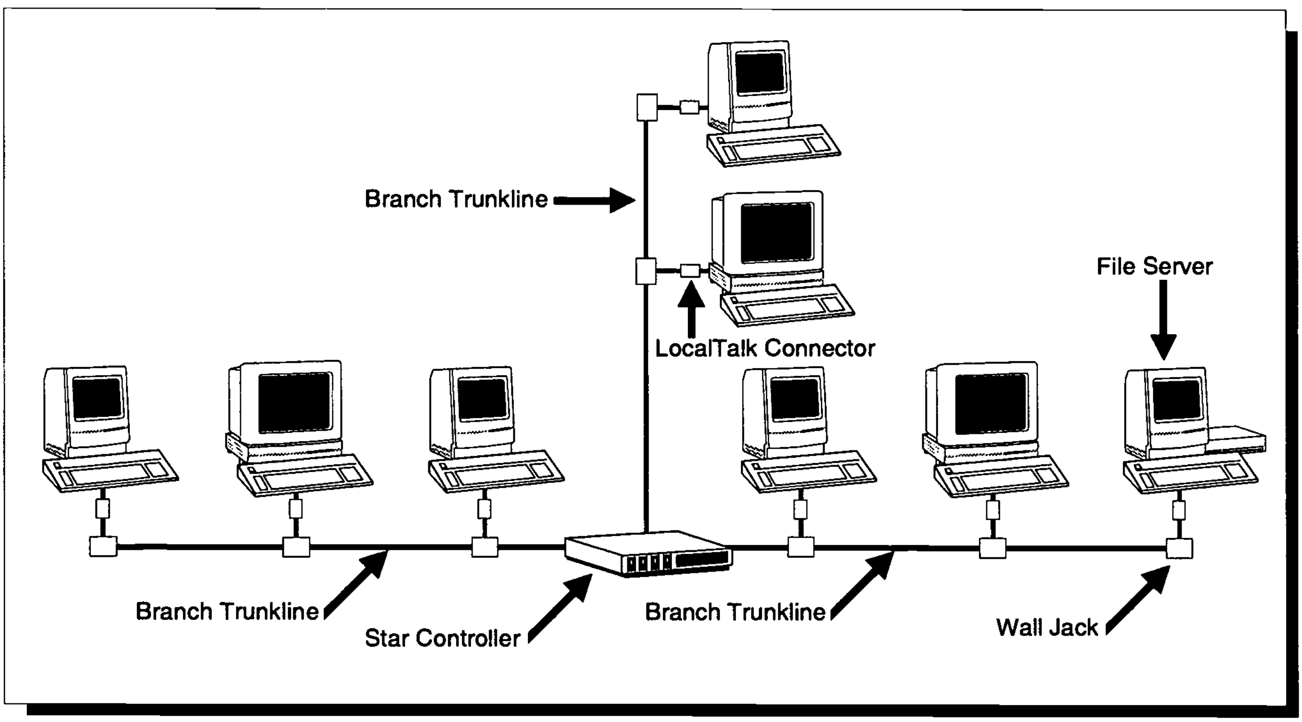

The star controller at the center of this star topology boosts signals to all branches or arms of the star. Notice that each branch of the star is a trunkline. The star controller not only boosts the network signals, it also contains instructions for shutting down workstations which begin sending garbaged signals. By so doing, the star controller can keep a network running smoothly that would otherwise have had major problems affecting many, if not all, workstations on the network. Note: star controllers do not create new zones for arms of the star. All arms are in the same network zone.

Ethernet

One of the great advantages of the AppleTalk networking system is its upward compatibility with faster networking hardware. Ethernet is a much faster networking system than either LocalTalk or LocalTalk-compatible systems. It requires its own interface to each computer and, in many cases, its own cabling. Ethernet does not rely on the LocalTalk chip inside the Mac for network communications, using its own set instead. These instructions are stored on chips inside the interface you purchase for each computer on an Ethernet network. Ethernet is able to achieve substantially faster network communication speed because the interfaces are designed for high speed cabling. Thus, Ethernet is an AppleTalk network using Ethernet hardware.

Ethernet is much more expensive than the other two solutions we have discussed. So why would anyone consider such a network for their Macs? The answer is speed. Ethernet on the Macintosh runs at a base speed of 10,000,000 bits per second, the same speed as on IBM-compatible computers. Most Mac network users would be amazed at how fast their network would respond if they upgraded to an Ethernet system. Ethernet can deliver information to a network workstation virtually as fast as a hard disk drive on a stand-alone computer. Currently, Ethernet can be purchased from third-party vendors, although Apple Computer Inc. has made statements that indicate that most new Mac computers will come with factory-installed Ethernet interfaces. High speed Ethernet is the direction that Macintosh networking is taking.

We are not saying that every Mac network must immediately be upgraded to Ethernet. Many LocalTalk and LocalTalk-compatible networks are a cost-efficient way to meet the needs of their users. But if speed is a problem, then Ethernet may be the solution you are looking for. Furthermore, as Macs are connected into school-wide configurations, some form of Ethernet installation will most likely be needed to make the network operate at an acceptable speed. We will discuss the ways in which to use Ethernet in school-wide networks in Chapter Eleven.

Now before you choke on the cost of installing an Ethernet network, just consider that this is a cost that must be faced by all computer owners who want to network their machines, except those who own Macs. Mac owners have become spoiled because their computers come with a built-in networking system. They don't have to pay any more to get a network up and running other than buying the cabling and connectors. It's actually a wonderful choice to have – you can install a LocalTalk or phone-wire based system and run with it until you have a problem with speed or until you can budget for the additional expense of upgrading to Ethernet. Part of a network may even run Ethernet and the other part run LocalTalk or a compatible system.

To install an Ethernet network, you must purchase an interface for each computer, printer, and any other network device. For Mac 512s, Pluses, Portables, Classics, Classic IIs, and Powerbooks, this will usually mean buying an external interface that connects to the computer through either the SCSI or printer port on the back of the computer. For all other Macs, interface cards can be purchased that fit in one of the internal slots in the machine.

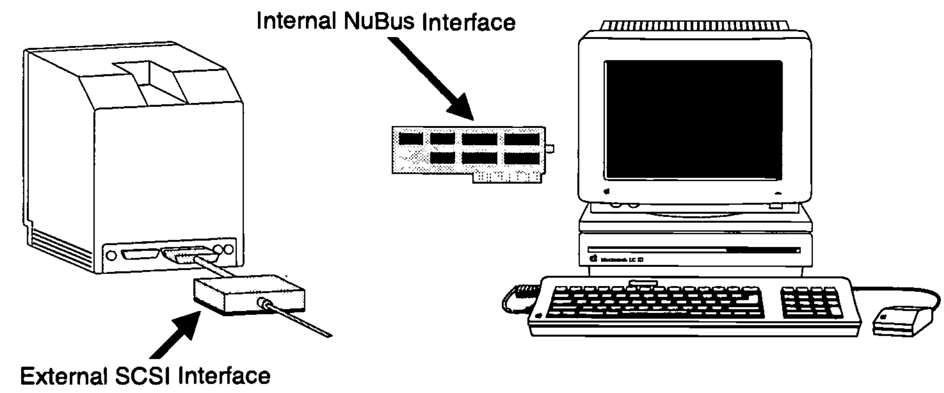

Most Macs do not come equipped to connect to an Ethernet network – additional interfaces are required. Older Macs and PowerBooks must use an external interface that plugs into the SCSI port on the back of the computer. Macs that have internal slots can use Ethernet interface cards.

Ethernet has a variety of standards which allow it to run on different cabling. The following table will tell you the type of cable needed for the three most common Ethernet standards.

Type of Cable	Ethernet Standard
Thick coaxial cable	10base-5
Thin coaxial cable	10base-2
Unshielded twisted pair cable	10base-T

The Ethernet standard is giving you three pieces of information. The first number indicates how fast the signal is travelling on the cabling – in all three cases here, the signal is travelling at 10 megabits per second (10 million bits per second). The word "base" indicates the type of signalling. "Base" indicates baseband signalling, a digital signal, and "broad" indicates broadband signalling, an analog signal. The final piece of information is the maximum length of a cabling run before a signal is boosted. 10base-5 indicates a maximum length of 500 meters, 10base-2 indicates a maximum length of 200 meters, and 10base-T has a maximum length of 100 meters.

The original thick coaxial cable is expensive and very difficult to work with, so it is seldom used in schools. If either type of coaxial cable is used, Ethernet is configured as a bus. With unshielded twisted pair cable, a star topology is used. Components are available to connect these types of cable, so a backbone bus of coaxial cable linking stars of unshielded cable is common. Ethernet networks with this combination of topologies often have hundreds of workstations. We recommend 10base-T with a structured wiring system for Ethernet wiring.

Since there is such a variety of Ethernet connectors, it is important to use cards that have the right connections. A 10base-T card designed for unshielded twisted pair will have a RJ45 connector, one designed for thin coaxial cable will have a BNC connector, and one designed for thick coaxial cable has an AUI connector. Installing an Ethernet network will require outside expert help – don't try to do it all by yourself.

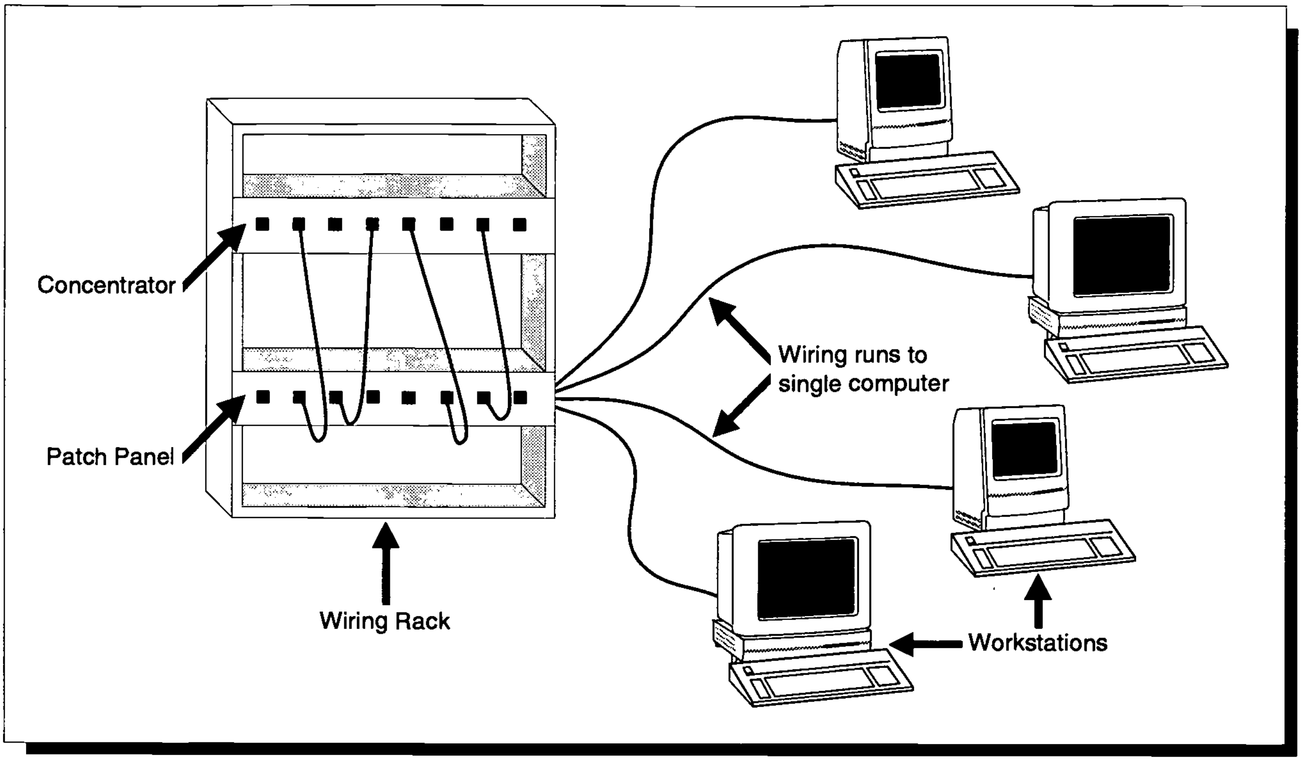

Ethernet networks are based on a structured wiring system where a single workstation is given its own cabling run. Ease of re-configuring and troubleshooting are the main advantages to structured wiring. The system utilizes a concentrator to boost and re-time signals on the network. Wiring racks for concentrators and patch panels are located in utility rooms or their own wiring closets.

Token Ring

The Token Ring network system was designed by IBM and is used extensively with IBM-compatible computers in business and, to a lesser extent, education. Token Ring networks run at high speed and suffer less degradation in response time when there is a high volume of network traffic. They are, however, expensive.

Token Ring is not common on Macintosh computers. The interface cards available from Apple and other manufacturers are intended for those Macs that must be connected to a Token Ring network created for IBM-compatible computers. The interfaces allow the Mac to function just like any other workstation.

Fiber optic cabling

Now we are talking real speed! Fiber optic cabling systems run at communication speeds that are at least ten times faster than Ethernet (depending on the system, up to thousands of times faster)! Fiber optic cable is capable of sending many more signals simultaneously at these terrific speeds than Ethernet. And it gets even better! Fiber optic networks are unaffected by electromagnetic interference generated by the other electrically operated devices in the building. Better yet! Unboosted signals on fiber optic cable can travel much greater distances, several kilometers if necessary!

Sound too good to be true? No, everything we just told you is true. So what's the catch? The catch is that this kind of network is relatively untested, very expensive for both cabling and workstation interfaces, and not widely used by businesses or schools. The cost will certainly decrease dramatically over the next few years as the sales volume of fiber optic systems increases. The power of fiber optic networking

will likely be needed sooner than you might think, because internetwork connections in many businesses and schools are becoming so complex that they are beginning to bog down with the sheer number of messages being sent.

Fiber optic networks require special interfaces for each workstation. They also require a device called a concentrator – an active device that is located at the center or hub of a network. Besides controlling the signals on the fiber optic network, the concentrator includes the logic required to make it function as a router and/or gateway, which converts the optic signal for use on LocalTalk or Ethernet networks. One of the most common configurations using fiber optic cabling will likely be a fiber optic backbone to which many LocalTalk and Ethernet networks are connected. See the Chapter Thirteen for more discussion of a fiber optic backbone.

Increasing the speed of Mac networks

Sooner or later access speed becomes an issue on Mac networks, unless you started right off with an Ethernet or fiber optic network. While these are the best solutions, we want to offer some other ideas for your consideration. First, however, if you are running a LocalTalk or LocalTalk-compatible network, do not buy a high powered Mac to act as your network file server. While you will get some improvement, our tests with high-powered Macs up to the Mac IIci and the Mac IIfx indicate that the most you can expect is a 25% to 30% increase in performance. The problem is not with the speed of the server as much as with the speed of the cabling. Using LocalTalk or a LocalTalk-compatible system is like trying to suck peanut butter through a straw – slow work! Using a high-powered Mac for the server on LocalTalk networks is a waste of machine. The speed of the hard drive connected to the server, however, can have an impact on network performance and at much lower cost than trying to upgrade the server. We recommend you get a hard drive with a seek time of 15 milliseconds or less.

The first step to take toward improving network speed is to add more RAM to the server. RAM (Random Access Memory) is the internal memory of a computer that requires electricity to remember what has been stored there. Network server software uses RAM to store the information it sends out to network users. Some networks use what is known as a RAM caching strategy where each user is given a part of the server's RAM for temporarily storing the information that the user requests from the server. Increasing the RAM in the server usually means that the server can respond to users more efficiently and, in some cases, handle more users altogether. Incidentally, do not use the Control Panel on the server to turn on the RAM cache for that machine. This removes memory from use by programs and "robs" the network server software of memory it desperately needs.

Another way to speed up a LocalTalk network is to use a file compression system. File compression is a technique that packs information into a smaller form than the original file without losing any of the data. It can be used on program and data files alike. Along with the compression program, you must also have a decompression program that expands the compressed file to its original form for use by the computer. On a network, the trick is to store compressed files on the server and send them in their compressed state over the cabling to a workstation, before expanding the files for use. This requires that each workstation be equipped with the decompression software. This software can be either loaded automatically from the boot disk or copied by the user from the server. There are several file compression programs on the market for the Mac, but you must be careful to get a system that will run as described on a network.

One of the easiest and cheapest ways to improve network speed is to physically cut your network cabling in half and put a server on both sides. A Mac Classic will do nicely for both servers. However, there are two factors to consider with this option. First, you must buy an extra computer and hard disk drive, or use one of the existing workstations as the second server. Second, you have doubled the amount of administration work for the person responsible for the network. Any file that is to be installed on the network must be installed twice because the networks are physically separated. It doesn't sound great and at times it's a pain in the "you-know-where," but this solution does dramatically improve network access speed at a fraction of the cost of other solutions.

Another solution to the speed problem is to install hard disk drives on every workstation on the network. This allows you to install the programs to be run by users on their own hard drive. This greatly speeds the system because requests to run programs from the server place great demands on the network, especially when an entire class wants to run a large program at the same time at the beginning of class. With the programs located on users' own hard disks, the network is only needed to save personal work, for public items such as assignment and test files, for print spooling, and for electronic mail. Only one network server is required because the demand has been greatly reduced. This solution, too, has its drawbacks. One of the major reasons for installing a network in the first place was to reduce the administrative work required to keep hard disk drives operating throughout a school. This solution also creates problems for ensuring the security of the programs. However, it is our experience that this solution can be a viable one for a lab where there is high demand for large programs and a competent network administrator. A desktop publishing lab using several large graphics and page layout programs, for example, may be a candidate for this configuration.

The optimum solution to the slow speed of the built-in networking that comes in every Mac is to upgrade to an all-Ethernet or all-fiber optic network. With Ethernet, at worst the cost will be only slightly higher than installing hard disk drives on every computer, and possibly even a little cheaper. With fiber optic cabling, the cost is substantially higher for both the cabling itself and the interfaces for each workstation. We will explore some of the possibilities for incorporating Ethernet and fiber optic cable into Mac networks in Chapter Thirteen.

Choosing the right networking software

There are two main types of networking strategies for the Mac – peer-to-peer systems and centralized file share networks. Let's discuss them both.

Peer-to-peer systems

Peer-to-peer systems represent a new networking strategy that is ideally suited to work groups. The idea is that members of a work group often want to share files but do not always have the time or the expertise to plan, install, and maintain a full traditional network. Thus, peer-to-peer networks are designed to run without an administrator overseeing who does what. Individual users decide whether they will connect to the network for any given work session and, on some systems, which files they want to share with other users. The main reason for these networks is to share files. Network users can "see" the files on other users' hard disks and copy each others' files. Using third party software, electronic mail can also be installed on peer-to-peer systems, but one of the computers on the network must act as the mail server for this to function. Peer-to-peer network systems are ideal for business work groups. They can be used effectively in some situations in education, but generally they are not the best solution for schools.

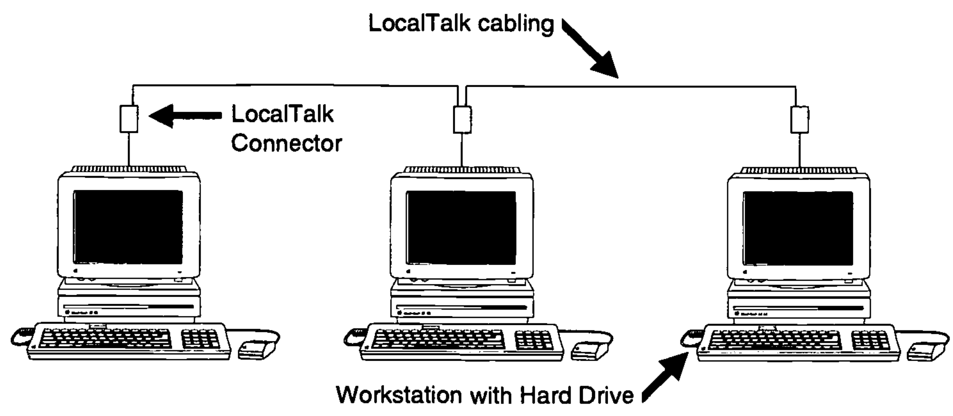

This is a peer-to-peer network system. Notice that there is no file server, just workstations connected by LocalTalk cabling. Each hard disk on the network can be accessed by any other user. The level of access to information and programs is determined by each individual user for his or her particular hard disk. Peer-to-peer networks are effective when you have groups of people who need to work together and the level of computer expertise is relatively high.

The two most widely used peer-to-peer network systems for the Mac are TOPS and System 7 file sharing. TOPS is the original peer-to-peer system for the Mac; it also runs on IBM-compatible computers. This means that files can be transferred not only between different computers on a TOPS network, but also between different makes of computers. When TOPS is used in conjunction with file translation software like MacLink, data transfer between Macs and IBM compatibles is simple and fast, and users don't even have to know what kind of computer others are using when they are working together.

In the spring of 1991, Apple Computer Inc. released System 7, a major upgrade of the operating system for Macintosh computers. Among other advances, System 7 allows users to connect two or more Macs together for exchanging files. The great advantage is that System 7 comes free with every new Mac you buy and it is much less expensive for older Macs than buying TOPS. However, if you want to run both IBM compatibles and Macs together on a peer-to-peer network, then TOPS is probably a better solution.

Centralized file share networks

We discussed the development of file sharing networks in Chapter One and you can refer back to that chapter for a more complete explanation than we will present here. Briefly, a centralized file share network consists of a dedicated file server providing access to private user spaces, public space for programs and shared data files, and additional features such as print spooling and electronic mail. Furthermore, file share networks usually offer more security for programs and data files than peer-to-peer systems. Clearly, centralized file share networks are better for meeting the majority of the needs of schools and we recommend them as the first choice for educators.

In some cases, centralized file sharing networks for the Mac are working with a considerable handicap because the communication speed of LocalTalk is so slow. IBM-compatible networks do not have to deal with this speed problem. Thus, you will generally find network software for the Mac to be more sophisticated in their communication strategies and sometimes more complex for network administrators than similar network software on the IBM compatibles.

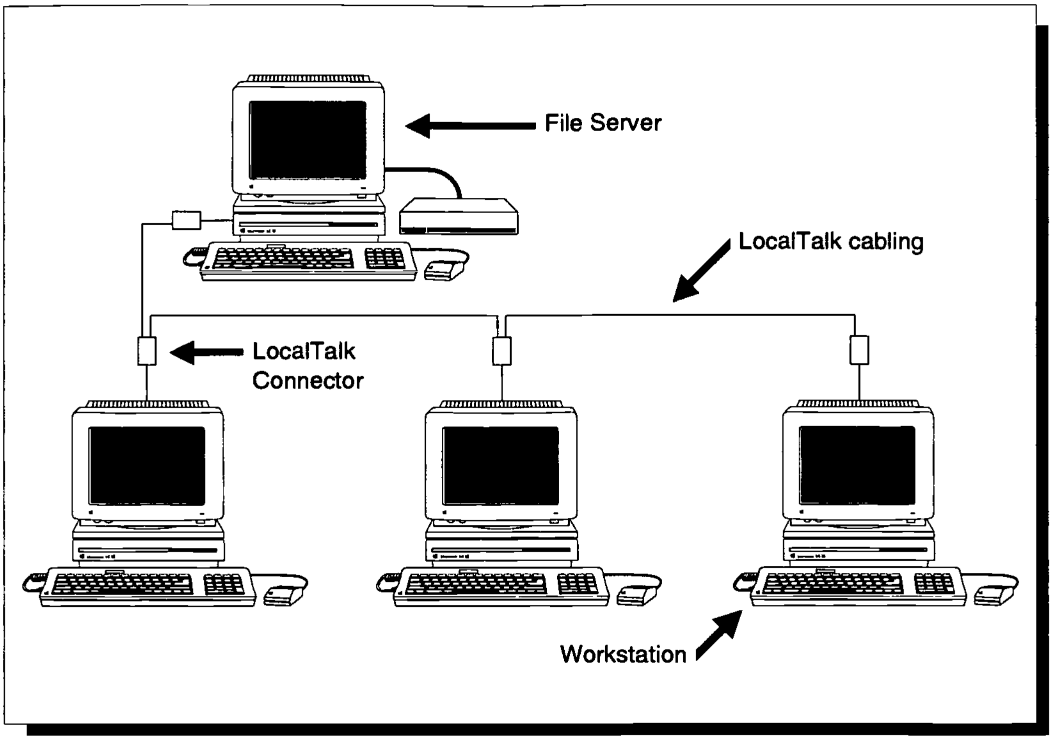

A centralized file share network requires a computer dedicated to the task of providing network services to all workstations. The level of access to network resources is determined by the network administrator. Centralized systems are excellent all-round networks and can be used by novices and experts alike.

There are two main file share network solutions for Macintosh computers – AppleShare and MacJanet. Many Mac users are not aware of the MacJanet networking software and many others have not encountered either AppleShare or MacJanet. Thus, an overview comparison of these two network solutions may assist you in making the right choice for your network.

Comparison of the AppleShare and MacJanet file share networks

MacJanet was the original file share network solution for the Macintosh. It was released in October of 1986. MacJanet was developed at the University of Waterloo in Waterloo, Ontario, Canada. Unlike AppleShare, MacJanet was intentionally targeted at meeting the unique needs of schools first, although MacJanet is also used in some businesses. AppleShare was released in late 1987. It is Apple's own networking software. It was originally targeted at business, although it is also used in education.

In this next section, we compare the MacJanet 3.2 with AppleShare 3.0. Recently, Apple introduced version 4.0 of AppleShare to run on its new workgroup servers. Although the internal workings of the software had to be adapted for the new hardware, the features remain virtually identical to version 3.0. Thus, while not accurate in every detail, the following comparison is largely valid for AppleShare 4.0 as well. Let's put the two networks up side-by-side to compare their features.

Comparison of AppleShare and MacJanet

Purchase options:

AppleShare 3.0
- Single copy
- No site licenses available

MacJanet 3.2
- Single copy
- Site license: Unlimited number of servers for an entire school district. The cost varies with the size of the school district. Site licenses are maintained by an annual fee.

System requirements for file server:

AppleShare 3.0
- 4 megabytes of RAM running System 7
- Cannot run server on anything but System 7

MacJanet 3.2
- 2 megabytes of RAM running System 5.1 or later
- 4 megabytes of RAM running System 7

Workstation requirements:

AppleShare 3.0
- Mac Plus or higher
- IBM-compatible
- Apple IIgs
- Apple IIe

MacJanet 3.2
- Mac 512KE or higher
- IBM-compatible
- Apple IIgs

Cabling:

AppleShare 3.0
- LocalTalk
- Phone wire cabling systems
- Ethernet
- Token Ring
- Fiber optic cabling

MacJanet 3.2
- LocalTalk
- Phone wire cabling systems
- Ethernet
- Token Ring
- Fiber optic cabling

User limits:

AppleShare 3.0
- Limit of 120 users (limit also imposed by cabling system, amount of space on server's hard disk, and available RAM in the server)

MacJanet 3.2
- No limit on number of users (limit solely imposed by cabling system, amount of space on server's hard disk, and available RAM in the server)

Network Administration features:

	AppleShare 3.0	MacJanet 3.2
• Groups/categories of users	yes	yes
• Private user spaces	yes	yes
• Public spaces	yes	yes
• Guest users	yes	yes
• User access without network space	no	yes
• Fixed size limit for user spaces	no	yes
• Auto-expansion of user space as needed	yes	no
• Create numbered users automatically	no	yes
• Create users automatically from a disk file	no	yes
• Global changes to an existing group of users	no	yes
• Delete all files from a group of student user spaces	no	yes

	AppleShare 3.0	MacJanet 3.2
• Remove log-on privileges for specific users	yes	yes
• Change user passwords	yes	yes
• View all network passwords	no	yes
• Multiple view options for looking at network volumes	no	yes
• Print-out or disk file of network configuration (group names, public area names, private user space names, access privileges, passwords)	yes	yes
• Network administration from any workstation	no	yes
• Network administration from server	yes	no
• Auto-save of user files when user deleted	yes	no
• Second-level administration for teachers to perform administrative tasks for just their students	no	yes
• Shut down server in specified number of minutes	yes	yes
• Shut down server at specified time	no	yes
• Password required to shut down server	no	yes
• Log-on to network denied a specified number of minutes before auto-shut down	no	yes
• Standard backup of server's hard disk allowed	no	yes
• Usage tracking of users, applications, printing, and mail	no	yes

Program security features:

	AppleShare 3.0	MacJanet 3.2
• Simple locking protection	yes	yes
• Encryption of programs so unauthorized copies will not run on a stand-alone computer	no	yes
• Quotas for the number of concurrent users of a program	yes	yes

Additional features:

	AppleShare 3.0	MacJanet 3.2
• Built-in print spooling	yes	yes
• Limit the number of pages users can print with the spooler	no	yes
• Hide the printer to force use of the spooler	yes	yes
• Built-in electronic mail	no	yes
• Built-in messaging to users	yes	yes
• Administrator control over who can send messages to whom	no	yes
• Send log-on messages to a group of users	yes	yes
• CD-ROM can be connected to server	yes	yes
• CD-ROM changes without shutting down the server	no	yes
• Built-in screen saver for the server	no	yes

User features:

	AppleShare 3.0	MacJanet 3.2
• Auto-log-on when workstation booted	yes	yes
• User can change his/her own password	no	yes
• User can set access permissions for folders	yes	no
• User control of own jobs in print spooler queue	no	yes
• User can send messages to other users (if allowed)	no	yes

AppleShare is an excellent solution for business work groups. The set-up and administration of AppleShare networks is relatively straightforward due primarily to the auto-expansion and contraction of private spaces as users add and delete files to and from the server. However, AppleShare was not designed for school use and has several shortcomings that cause problems in education. It lacks strong security features for protecting of programs. It relies on the "protect bit" to prevent unauthorized copying. Today, most of the data-recovery and file-copy utilities will easily by-pass this kind of protection. In the authors' experience, many students have copies of such programs (usually illegal ones!) and regularly take unauthorized copies of programs installed on AppleShare networks unless stronger steps are taken.

AppleShare also centralizes the administration of passwords for a network. While wielding all of this power may seem great at first, most administrators for school networks soon tire of the constant stream of users coming into their office because they have forgotten their passwords. The ability to decentralize as many administrative tasks as possible is key to maintaining the sanity of network administrators in schools. AppleShare does not allow a second level of administration for teachers to take care of daily tasks like viewing and changing passwords for their own groups of users.

There is no built-in electronic mail system in AppleShare and the messaging system is not complete. The administration program was not designed for schools, where it is often necessary to add a class of thirty users to the network. AppleShare requires you to add each student individually.

On the other hand, MacJanet was designed for schools and addresses the shortcomings of AppleShare. Program security via an encryption option in the network administrator program renders unauthorized copies of programs useless on stand-alone computers. Limited administration capabilities can be given to teachers for just their own students, thus relieving the network administrator from the tedium of dealing with a multitude of requests to perform various small administrative tasks. The ability to run the administrator program from any workstation is also a useful feature of MacJanet, especially if your network is distributed throughout the school. It may not seem like a problem, but always having to walk back to the server to do administrative tasks can waste a great deal of time. MacJanet's built-in electronic mail system is quite adequate for school use. The messaging feature is much more complete than that of AppleShare. The administrator program for MacJanet was clearly designed with education in mind, automating many features so administrators can quickly perform global changes for a whole group of users. The ability to create a group of users automatically from a disk file is ideal for schools where a file of student names can be generated from the main office computer.

All network spaces on a MacJanet network are given a size limit when they are created. The use of a fixed size for all spaces has both positive and negative aspects. For student spaces, a fixed size limit can be an excellent feature because it prevents a small number of students from saving large numbers of files and thus monopolizing the space on the server's hard disk. On the other hand, fixed volume size networks require a lot more planning and can result in problems that are simply not encountered with variable volume size networks. Nevertheless, because the other features of MacJanet are clearly better suited to school use than AppleShare, the authors recommend MacJanet as the best network solution for Macintosh computers.

Planning a fixed volume size network

Now that we have recommended MacJanet for most school networks, we are forced to deal with the issue of planning for networks with a fixed volume size structure. There are two major factors to consider when planning a MacJanet network: users need to save personal files on private user spaces, and the need for public spaces to hold the application software that users will access from the server.

The first step in planning private user space is to determine the underlying philosophy that will guide the use of your network. There are essentially three different approaches. Some network administrators feel that the network should store all student work and that students should not require any other storage of their files. Other administrators take exactly the opposite approach and feel that students should be completely responsible for saving their files on their own floppy disks. Yet other administrators want students to save copies of their work on both the network and on floppy disks.

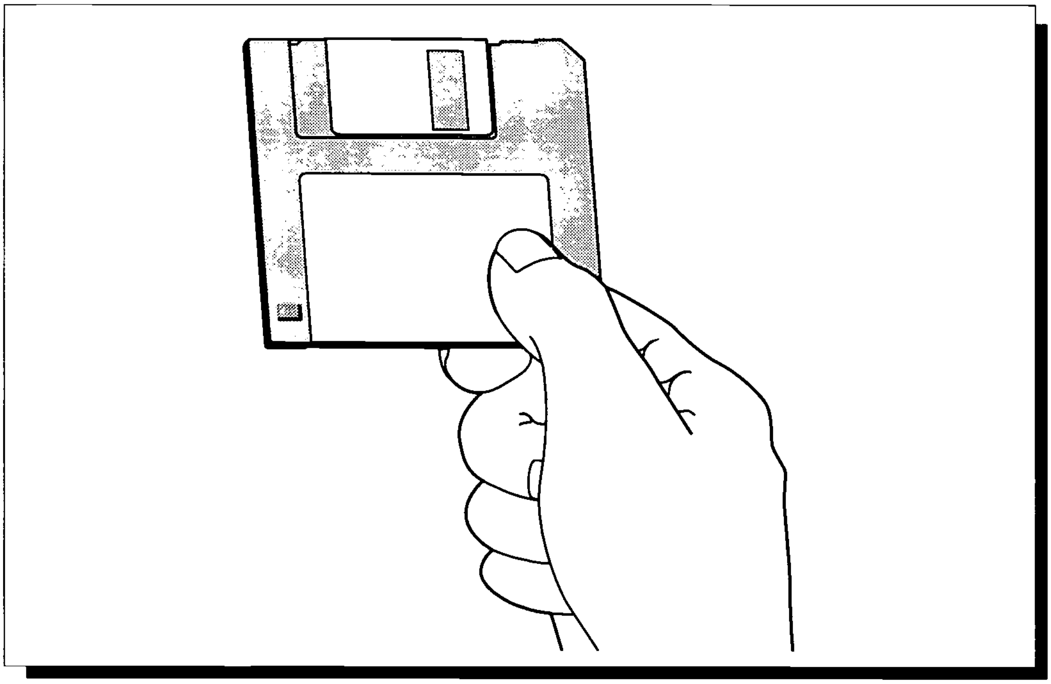

How floppy disks are to be used for storing files and backups is an important part of any network plan, especially when planning a fixed volume size network. We recommend you consider the use of floppy disks first, before any other planning for your network server.

Each approach has its pros and cons. If students rely completely on the network for saving their work, then computer use is very simple. On the other hand, if a file is damaged for any reason or there is a desire to retrieve some deleted work, those relying completely on the network will probably find they have no backup copies. The administrator must also plan for adequate space for students because once a private space is full, the student cannot save their work. Network administration is much easier if students save their work only on floppy disks. The administrator does not have to worry about creating private user spaces on the network. However, if students forget their disks at home or in their locker, then valuable class time may be wasted. In addition, students have been known on occasion to abuse items like floppy disks, resulting in unreadable files. With no backup copy saved on the network, the student is out of luck and the work must be redone. Saving copies on both

the network and on floppy disk would seem the best way to go, because two copies of files are kept. However, this approach, too, has its problems. Students must take the time to make two copies of a file whenever work is done. Students must also keep track of where the latest copy of a file is located or else they face the possibility of loading an old version and redoing work.

We recommend that you expect users to keep copies of files both in private user spaces on the network and on their own floppy disk. Making backups of computer work and keeping track of the most up-to-date work are essential life skills in an electronic world. Furthermore, the use of floppy disks to archive copies of files that are not currently needed relieves the network of providing huge private spaces for each user. We will assume that this is the approach you have taken with the majority of your users when we give you the planning guidelines below.

MacJanet provides unprecedented flexibility in handling network users. MacJanet networks often incorporate all three of the approaches just mentioned. One of the nice features of MacJanet is the ability to add new network users who have no private user space on the server. This means you can add users at any time since there will be no space demands put on the network. These users can be given all other network privileges including electronic mail, messaging, and print spooling. MacJanet also allows the administrator to create a set of user spaces that can be used by transient users, who use the space on the network for temporary storage of files. These features give MacJanet administrators a great deal of flexibility in meeting the needs of users and in solving problems that were unforeseen when the network was planned.

The second major consideration in planning fixed volume size networks is the space allocated to public areas for applications and data files. Programs for computers are always being upgraded. New versions, especially significant upgrades, provide users with new tools for increasing their productivity. However, if a network administrator has not planned for the additional space that an upgraded program requires, users will face the possibility of knowing that their favorite program has been upgraded, but it can't be loaded onto their network. The administrator could try to hide behind the fact that networks are very complicated, but the truth would be that the administrator did a poor job of planning. To avoid this scenario, follow this simple guideline – check the size of the current version of all the programs you plan to install on your network (including all related files required for a program to operate properly) and double that figure when you create the network space for the programs. Most of the time, this calculation will handle all updates to the programs before it is time to redo the network configuration due to the changing needs of your users.

It's time to give you a list of guidelines to follow that incorporate what we have just discussed (and a little more). Here are all of the steps that network administrators should take when planning a fixed volume size network like MacJanet for use in education. For the most part, administrators for any network for any computer should also take these steps to ensure that the configurations they create will be adequate.

1. Teachers who are potential network users must be interviewed to determine the space requirements for their personal private spaces, the private spaces for their students, and the public information space for a class, if necessary. Questions to ask should include when and how often the teacher and his/her students will use the network, what programs will be used, how extensively will the electronic mail system be used, and how much printing will students be doing.
2. These needs must be projected over the next two to three years if possible. Also, future users should be projected as accurately as possible. After years of

experience, one piece of advice we can give you is this – listen carefully to what the teachers say, write it down for future reference, and don't believe them! That is to say that most teachers will not be able to predict computer use correctly, especially if they are new to computing. It is often much more than they anticipate and, sometimes, much less. Our advice is to add a "fudge factor" based on what you think network users will require.

3. Group the users into categories with similar needs in terms of access to network resources.
4. A list of the programs to be installed on the network and their space requirements must be compiled. This usually involves the administrator projecting the needs and desires of network users to include programs that were not identified during the interviews. It may also involve doing some research into what is available on the market that will meet needs that were identified during the interviews. This is often a cyclical process, where you take the information from your research back to teachers before deciding on the best solution.
5. A list of public areas to be used for shared files and their space requirements must be compiled (shared files like network news, dictionaries, encyclopedias, databases, assignments, tests, etc. – often called info spaces).
6. The administrator must project the growth in space requirements for both public program areas and public file share areas, to ensure the network will remain viable for at least a year.
7. The amount of space needed for the print spooler to store its print queue on the hard disk of the server must be projected.
8. The amount of space needed to store electronic mail on the hard disk of the server must be projected.
9. The entire space for the network can now be calculated. If the total is less than 85% of the available space on the server's hard disk, then the network can be created. If the total is greater than 85% of the available space, then you have three choices. You can go ahead with the creation of the network if it will fit, or purchase another larger hard disk, or compromise on the space allocated to each user and public area. With the falling cost of hard disk drives today, we strongly recommend that you buy another for the server.

Following these guidelines won't guarantee that your network will be perfect, but it will ensure that your network plan is adequate to meet the current needs of your users and flexible enough to handle the new needs that will emerge over the next few years. Also, following these steps will make the process of creating a fixed volume size network like MacJanet relatively straightforward.

Choosing and installing workstation software

As you will recall from our discussion in Chapter Seven, the major concern for the workstation is loading the operating system and network extensions to the operating system when the computer is turned on. On the Mac, this usually means the preparation of floppy boot disks or the preparation of a hard disk drive for booting the computer. While it is possible to a limited degree, remote booting of a network workstation is not usually done on Mac networks. This has to do with the ROMs inside the Mac that provide the computer with all of its basic functions – the Mac's ROMs will not allow true remote booting even though the network software companies would like to provide that capability. Apple is rumored to be readying a ROM update that will allow true remote booting – we certainly hope this rumor is true! A company called Sonic has already developed an Ethernet boot ROM that enables Mac II class machines to boot remotely on Ethernet networks.

On the Mac, the operating system is not actually seen. Instead, the user interacts with the computer using a mouse to manipulate graphical objects. The program that makes this interaction possible is called the "Finder." Unseen to the user, the Finder generates all of the commands required by the operating system. The file that contains all of the instructions for the Mac's operating system is called "System." The System file, the Finder, and all of the other files needed by the operating system are stored in a folder called the "System Folder" on a boot disk. Creating this folder for new boot disks is extremely easy for Mac network administrators because Apple sends out installer programs with each update to the operating system. The installer for the operating system is located on the "System Tools" disk for System 6 and on the "Install 1" disk for System 7. All you have to do is run this installer and select the options you want to have on the boot disk and the installer takes care of the rest. When it is finished, the installer will have created a System Folder, complete with all files required to boot the computer. Mac users often are unaware of how difficult it can be on other computers to create boot disks and install operating system files. To put things in perspective, Mac users should watch someone trying to do the same process with UNIX or MS DOS operating systems (although even these operating systems are now beginning to use installer programs too).

Which version of the system for workstations?

The Macintosh operating system has gone through many minor updates and several major revisions since its introduction in 1984. Like all major software producers, Apple's numbering of its new versions of the operating system is consistent with the strategy we described in Chapter Seven. The latest major revision is numbered System 7. However, unlike most operating system developers, Apple is actively supporting two versions of its operating system – System 6 and System 7. The reason for this lies in the fact that System 7 will not run at acceptable speeds on older Macs, nor will all of its features run on those machines. Thus, Apple is continuing to both support and develop new versions of System 6 for its computers.

So which one do you use? Here are our unabashed suggestions for configuring workstation boot disks for your network. If your workstations are only equipped with floppy disk drives, then you must use a version of System 6 for your boot disks. System 7 is too large to fit on a floppy disk. The current versions of the system are 6.07 and 6.08 and either of them will do the job.

If your workstations are equipped with hard disk drives, then the decision becomes a little more complicated. The hard disk will certainly have enough room to store System 7 files, but the computer itself may not be adequate to the task. There are two factors to consider. First, you must consider the speed of the processor chip inside the Mac. When compared to the speed of System 6, we have found System 7 to be too slow on Mac Pluses, SE's, and Classics – any Mac using a 68000 processor chip running at 8 megahertz. All other Macs run System 7 at acceptable speeds.

Second, you must ensure that the workstation computers have sufficient RAM to run System 7. Apple says that System 7 will run on any Mac equipped with 2 megabytes of internal RAM. While technically this is true, we have found that it is not realistic. You must have 4 megabytes of RAM in the computer to run both System 7 and most applications on the market today.

So if you have a newer Mac equipped with a hard disk drive that is running at 16 megahertz and has 4 megabytes of RAM you can run System 7, right? Well, not exactly. Oh sure, System 7 will run just fine, but some of your favorite programs, Inits, and desk accessories might not. Before you even consider upgrading to System 7, do

some research into which of your programs and utilities will run under System 7, whether there is an update available for those items that will not run under System 7, and how much it's going to cost to update all the copies for all your workstations. In the short term you may find it best to stick with System 6.

There are two more important points to make regarding the operating system for Mac workstations. First, Apple has just released an add-on to System 7 called "At Ease," which presents the user with a Hypercard-like interface and locks them out of making any changes to the hard disk. It is a great solution for those networks where workstations are equipped with hard disk drives. It's pretty inexpensive, as well. This feature may make System 7 much more attractive for workstations because it greatly simplifies the job of network administration. (It's also great for stand-alone workstations in the school, too!)

Second, System 7 offers a number of new features that will make users more productive. Publish and Subscribe is one such feature that will allow a user to insert a picture that was created with a drawing program into a page layout program. There is nothing new with that, of course, but the System 7 Publish and Subscribe feature means that any changes made to the picture in the drawing program automatically show up in the page layout program without having to re-insert the picture file. As you can imagine, this will have far reaching implications for Mac users. The writing is on the wall that all Mac users will be using System 7 sooner or later. You should make sure that your long term plans include steps to upgrade or replace your network workstations so they are capable of running System 7.

Adding network files

To access a network, you must install additional files on the boot disk that contain the appropriate instructions for performing network tasks. These files are called network drivers; on the Mac, they are called "Inits" and are always placed inside the System Folder. This is the case for both AppleShare and MacJanet networks. AppleShare and MacJanet also require additional files to allow users to communicate with the network. AppleShare requires a Chooser extension that users access via the Chooser item under the Apple menu. AppleShare also requires a desk accessory called "Access Privileges," which shows up directly in the Apple menu. MacJanet requires only a desk accessory, which also shows up directly in the Apple menu.

AppleShare files are easily added to boot disks because, as Apple's own networking software, AppleShare files for boot disks are installed using the same installer that installs operating system files in the System Folder. Remember to use the installer program if any updates come with one – the installer not only copies the files to the appropriate place on your disk, it often makes important changes to the operating system or program files. AppleShare is a case in point. Do not simply copy the AppleShare files into the System Folder – use the installer.

MacJanet does not currently use an installer for adding network files to boot disks. Instead, the administrator must copy the network drivers to the boot disk by dragging them with the mouse from a master disk into the System Folder of the boot disk. There are two different ways to install the MacJanet desk accessory file, depending on the version of the operating system you are using. For operating systems numbered 6.0.x, the desk accessory must be installed by using the Font D/A Mover program that comes with the operating system disks from Apple. For operating systems numbered 7 or greater, the desk accessory is installed by dragging the file into the System Folder (System 7 is smart enough to know where to put the file for you – in the Apple Menu Items Folder).

Adding printer drivers

Printer drivers are files that contain the instructions necessary to operate the various printers that can be connected to the computer. Like all other system files, printer drivers are stored inside the System Folder of the boot disk. The only printer drivers that are applicable to network environments are those which drive AppleTalk printers. The candidates are the AppleTalk ImageWriter, the AppleTalk ImageWriter LQ, and the LaserWriter printers. If you will be using a LaserWriter on the network, be sure to install both the LaserWriter and the LaserPrep files onto the boot disk. Also be very careful that both LaserWriter files have the same version number (this can be checked using the Get Info option in the File menu of the Finder). If you will be using a print spooler, remember that the print jobs are stored temporarily on the server and then printed from there. This means that you must place the same printer drivers that are on the workstation boot disks in the System Folder of the server as well.

Adding Inits to your system

The System Folder can also contain files that extend the operating system to handle tasks not envisioned when the system files were originally designed. On the Mac, these extensions usually take the form of Inits that are automatically installed in the computer's memory when it is turned on. There are literally hundreds of Inits that you can get for your Mac, but we want to mention just five that are particularly useful for network environments – they're also free!

The first is called "Eradicat'Em." This Init removes a virus called "WDEF" from the desktop file on the boot disk each time the computer is turned on. The WDEF virus can cause a number of problems with the saving and retrieving of files from an infected disk, including very slow loading/saving, an inability to save on a disk that has free space, and system errors. WDEF replicates itself very quickly and is difficult to remove once users' disks are infected. Eradicat'Em can be a life saver for a network administrator because it removes WDEF every time a computer is booted. Thus, the virus is being eliminated continually and re-infections of the boot disk are handled automatically. You can get Eradicat'Em from most Mac electronic bulletin boards.

Eradicat'Em does not remove the WDEF virus from users' data disks, however. Nor does it look for or remove any other viruses. This is why we recommend another Init called "Disinfectant Init." This Init looks for all known viruses and will alert the user if one is found on any disk inserted into the disk drive. Once a virus has been found, you run the Disinfectant application to clean up the infected disk. The Init and the application are being continually updated to handle new viruses. You can get the files from most Mac electronic bulletin boards or from Northwestern University.

Two other Inits perform exactly the same functions. They are "SFVol Init" and "Drive Select." These Inits allow you to easily switch between network disk volumes on MacJanet networks without having to continually click on the "Drive" button. This drive switching is done via a pop-up menu that appears in both the "open" and "save" dialog boxes. All you have to do is select the network volume you wish to use from the menu and the switch is immediate. Drive Select was developed specifically for MacJanet and it is recommended for all MacJanet users. It is provided free of charge when you purchase MacJanet. You can get SFVol Init from most Mac electronic bulletin boards.

The last Init worth mentioning is one that was developed by Apple itself. This Init solves a potentially costly problem on older Macs with 800K disk drives. The Mac was designed to make computing easier, so the disk drive was made to automatically eject disks without having to manually remove the disks with physical force. While the idea

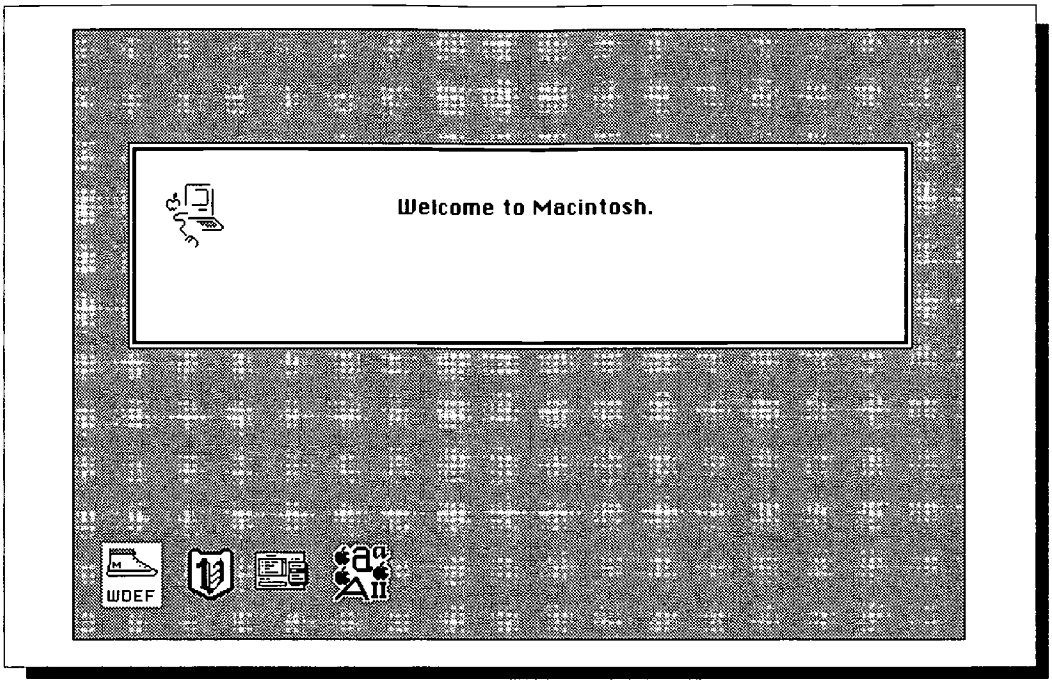

Inits are special programs that become active as soon as the computer is booted-up. Most Inits will display an icon on the screen as the computer boots. Here you can see the icons for Eradicat'Em, Disinfectant, SFVol Init, and Suitcase.

is a good one, the old 800K disk drives sometimes do not pull the read/write heads of the drive out of the way in time, so they are bumped by the disk as it is ejected. Once the heads are out of alignment the only solution is a complete drive replacement at a cost of a couple of hundred dollars. To remedy this situation, Apple wrote an Init called the "800K Eject Init." This Init pulls the heads of the disk drive back before ejecting a disk.

Keeping your System Folder small

You may have noticed that most of the files we have talked about in this section are stored in the System Folder on the boot disk. It also may have crossed your mind that the System Folder can get pretty large. This is, in fact, exactly the case – so much so that space on a boot disk quickly becomes a problem for Mac users, especially if you have 800K disk drives (although space can become a problem even with high density disk drives). There are many strategies for reducing the size of System Folders, but we would like to mention just two of the more useful ones.

The first method for reducing the size of the System Folder is provided by Apple. The operating system installer found on the System Tools disk contains a "Customize" option that allows you to select from a number of configurations for a boot disk. One of those configurations is called the "Minimal System." This option is available for each of the specific models of Macintosh computers. When selected, the Minimal System option only installs the files required to boot that specific model of Mac. The boot disk cannot be used to start up a different model. The normal installation places the operating system files for all of the Mac models on the boot disk. This is great because you can use the disk on any Mac, but it also makes the System Folder considerably larger than a minimal installation. If you have only one or two models of Macs for network workstations, then create boots disk using the Minimal System option to save space on the boot disks (be sure to label which model a boot disk is for if you have more than one model on your network).

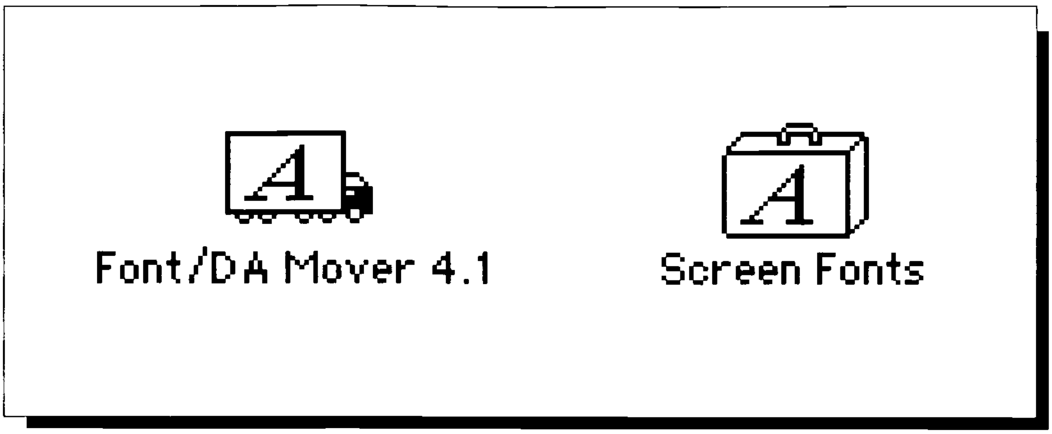

The Font/DA Mover is provided with the System disks that come with a Macintosh computer. This program allows you to make suitcase files containing fonts and/or desk accessories. It also allows you to install fonts and desk accessories directly into programs.

The second way to save space on a boot disk is to reduce the space taken up by font and desk accessory files that are stored in the System Folder. Font and desk accessory files are real space gobblers! There are two solutions to this problem. The first way to reduce the space taken up by fonts and desk accessories is to use the Font D/A Mover program to install the fonts and desk accessories directly into the application files that are saved on the network. Then all but the very basic fonts and desk accessories are removed from the boot disk. The fonts and D/A's will appear when a program is run. In this way, you can also customize which fonts and D/A's you want to have in different programs.

Another solution is to use a program that will access font and desk accessory files from the network instead of the boot disk. Either Suitcase from Fifth Generation Software or Master Juggler from AlSoft will handle this task. This method requires the purchase of a copy of a utility program for each workstation on the network. While the cost of this strategy may seem steep at first, you may find the money well spent when you consider the power of this solution. Using the Font D/A Mover program, place all of the bit-mapped fonts and desk accessories in one or two suitcase files. Then save these suitcase files on a publicly accessible network space. After purchasing the appropriate number of either program, install a copy inside the System Folder of each boot disk. When users log-on to the network, they use either program to open the suitcase files on the network to get their fonts and desk accessories. This means that the boot disk only needs to contain the absolute minimum number of fonts – Chicago 9-point, Geneva 9- and 12-point, and Monaco 9-point. The only desk accessories needed are those required to access the network. The result is a substantial saving of space on the boot disk, not to mention the ability to use the added capabilities of Suitcase or Master Juggler.

There are many other useful tricks and techniques that can streamline the creation and use of boot disks on Mac networks. The tips we have given you here will at least get you started. We strongly recommend, however, that you take advantage of any courses offered by experienced Mac network administrators or any more informal opportunities for discussing the various tricks and techniques of Mac networking with fellow network administrators. The information you will gain will be invaluable in making your network run efficiently.

Protecting software on hard disks

All new Macintosh computers are equipped with hard disk drives. The lack of remote boot capability for Macs means the operating system must be installed on these hard disks. In addition, many of the new versions of application programs are so large, they must be installed on local hard disks to ensure acceptable response times. However, the use of hard disks in workstations can result in a dramatic increase in work for administrators. Students can make illegal copies of programs, they can do inadvertent or malicious damage to software on the hard disk, and they can fill up the hard disk by saving their work there. Here are some strategies for making the management of workstation hard disks a little easier.

Apple Computer Inc. has created a program called At Ease to provide protection for hard disks. At Ease replaces the Finder and allows students to run programs on the hard disk without being able to access the disk's contents. Unfortunately, At Ease is amazingly simple to bypass and we have discovered that students find their way around this software in a few days.

To protect the operating system files, you can use the ResEdit program to make system files invisible. If you are using System 6, you can make the entire System Folder invisible. With System 7, the System Folder must be visible, but you can make all files and folders inside the System Folder invisible except the Apple Menu Items Folder. The desk accessories normally found in the Apple Menu Items Folder can be moved into another folder in the System Folder and replaced with aliases. The alias files can be re-installed easily if a student deletes them.

ResEdit can also be used to make application program files and folders invisible. Students run the program by double clicking on a data file that has been left visible. This technique can also be used to protect programs installed in public areas on AppleShare and MacJanet networks.

Hard disk partitioning software can also be used to protect programs. By creating a read-only partition, students are prevented from making changes. However, the System Folder must not be placed in a read-only area because the operating system is continually writing temporary files to the hard disk.

You can also purchase hard disk protection programs. DiskLock from Fifth Generation and FolderBolt from Kent Marsh both provide excellent protection for hard disks and folders. However, the cost of these programs is prohibitive for an entire lab or school-wide network. Another solution is FoolProof from SmartStuff Software (SmartStuff Software is a small company in Portland, Oregon (503) 775-2821). Not only will FoolProof protect the contents of the hard disk, site licences are available at very reasonable prices.

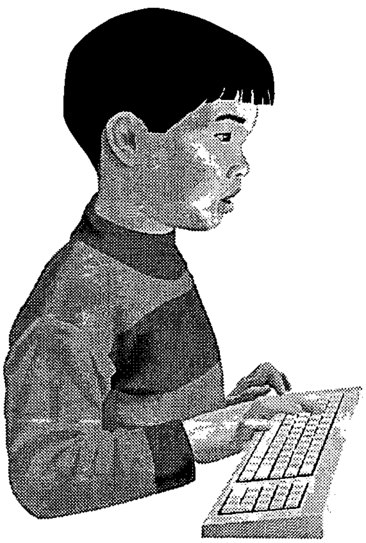

Chapter Twelve

Simple networks for Macintosh computers

Chapter 12

Simple networks for Macintosh computers

Topics covered in this chapter...

- *A discussion of small LocalTalk networks including:*
  - *Who can use this kind of network*
  - *A technical description of the networks*
  - *What the administrator does on these networks*
  - *What the teachers do on these networks*
  - *What the students do on these networks*
  - *A discussion of the potential for future growth*

Introduction

There are a multitude of ways to wire and configure networks of Macintosh computers. Obviously, we can't cover them all in this publication, but we can give you a few ideas. In this chapter we will deal with small networks that use the built-in networking capabilities of Macintosh computers and look at the most common places these networks are used.

Small LocalTalk networks

Overview

LocalTalk and LocalTalk-compatible networks are very popular for Macintosh computers because they work quite well for small networks of 30 users or less and because they are inexpensive to install when you compare them to other network solutions like Ethernet. Virtually all of the features of networking we discussed in this book are possible with a LocalTalk network, including file share access to programs, private user spaces for saving personal work in a secure place, centralized administration, security measures for programs and important data files, print spooling, electronic mail, log-on messages from teachers to students, and file share access to public areas for retrieving copies of assignments, tests, administrative data, library data, and any other information.

The examples in this chapter are designed for a centralized file sharing network run on a dedicated server. Each example will be based on a trunkline topology for the network and will use a phone-wire based cabling system. Either AppleShare or MacJanet can run on this cabling. We recommend that MacJanet be used for most networks intended for student use. We also recommend MacJanet for administrative networks, but we recognize that some administrative network software packages require AppleShare or TOPS to run properly.

Who can use this kind of network?

The most common place you will find LocalTalk networks is in classrooms housing full labs of 24 to 30 computers or mini-labs that generally consist of about 4 to 15 computers. These classroom networks are usually used for specific courses like English writing, Computer Science, Drafting, or Business Education courses. Mini-labs are often found in Learning Assistance Centers and Special Education classrooms. Libraries often use LocalTalk labs or mini-labs for general student use as well as for access to electronic card catalogues, encyclopedias, dictionaries, and other resources.

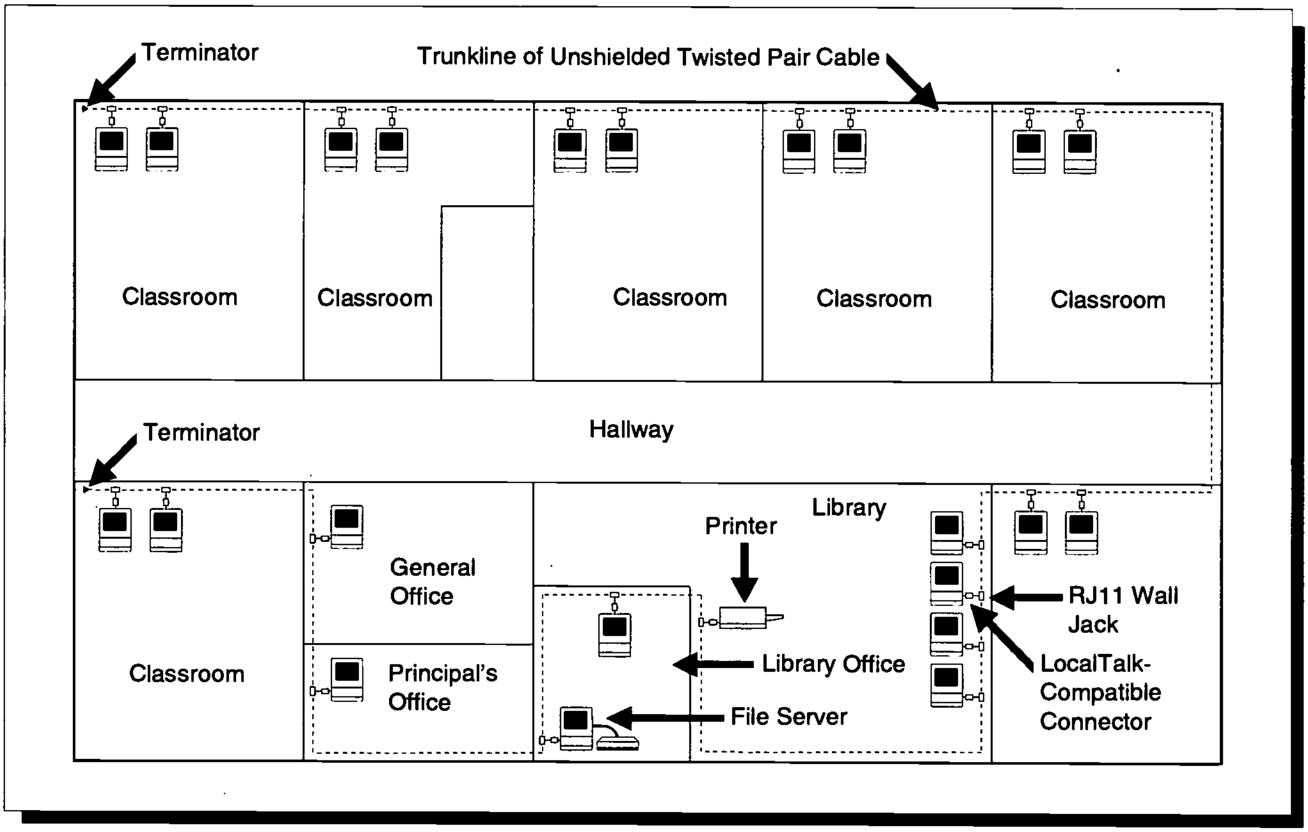

This small school is wired with a LocalTalk-compatible network. Here the single file server would have spaces for instructional programs, student and teacher private areas, and administrative programs and data. If the administrative programs required different network software than the instructional network, a second file server can be connected to the trunkline in the office. Both file servers can operate on the same wiring without interfering with each other. The trunkline of Unshielded Twisted Pair cable can be run in the ceiling, in the walls, or under the floor. Each workstation connects to the trunkline via an RJ11 wall jack. Notice that all devices on LocalTalk networks are nodes on the network. Thus, the printer connects directly to the trunkline, not to the file server.

Another place you will often find a small LocalTalk network is in the school's administrative office. The interconnection of computers here allows principals and counsellors to have simultaneous access to student achievement and attendance information. These networks can also provide an electronic mail feature which can be useful for administrators and counsellors who are always trying to get messages to each other. Print spooling and access to protected programs such as word processors and spreadsheets are also useful features of networks located in a school's administrative office.

LocalTalk can also be used to network an entire school provided the school is small (a school of 5 to 15 classrooms with a total cabling length not exceeding 3000 feet). Elementary schools fit this profile more often than high schools.

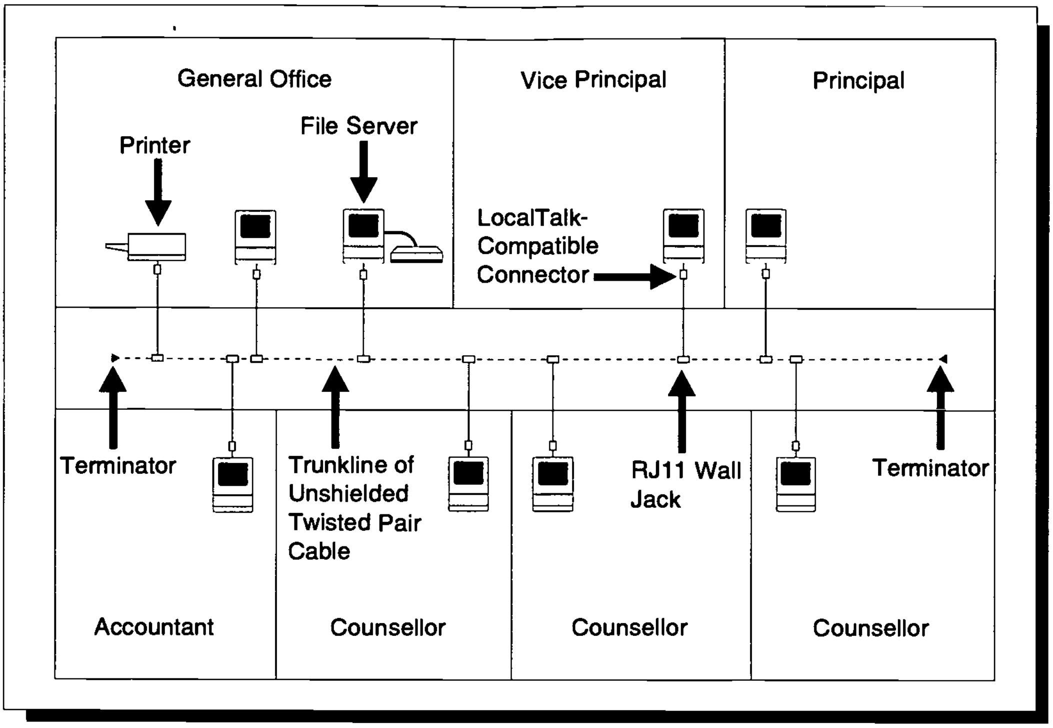

The network is strictly for the administrative area of the school. It will be used for storing student data and mark records, scheduling, and printing report cards as well as storing productivity programs for administrators, counsellors, and office staff.

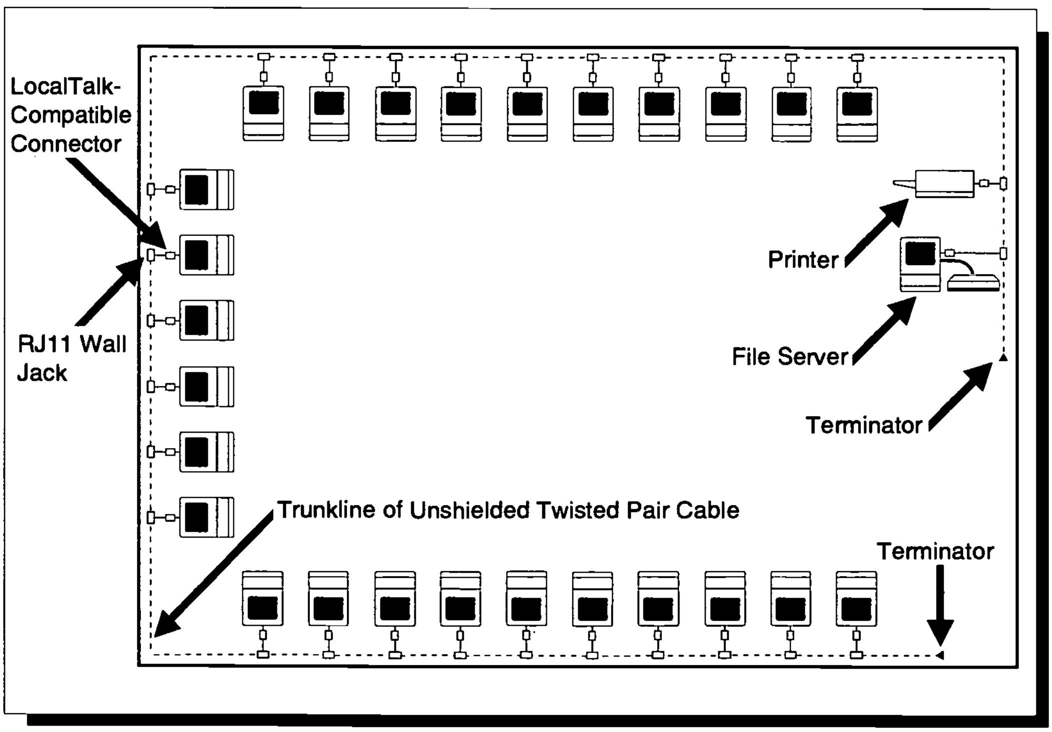

Here a computer lab uses a LocalTalk-compatible network configured as a trunkline. The file server would store instructional programs and private work spaces for students and teachers. The trunkline cabling can be run under counter-tops or along the wall.

Description of the network

The network examples you have just seen share a number of common characteristics. Let's discuss these shared attributes in more detail.

Network cabling

LocalTalk and LocalTalk-compatible networks use the built-in chip inside each Mac to transmit and receive network communications, so no additional interface is needed for each workstation. As we discussed in the previous chapter, the trunkline topology is preferred over a daisy-chain. This means that the cabling for these network examples must use a phone wire cabling system instead of LocalTalk cabling. It also means that the cabling system will be cheaper. The trunkline will consist of solid core 22-gauge copper wire that is run uncut from one end of the network to the other. Workstations are connected to the network via RJ11 wall jacks. The trunkline wire is stripped bare at each wall jack and screwed into the jack. By keeping the wire uncut you increase the distance an unboosted signal can travel.

There will be only two terminators on the trunkline. These will be wired inside the last wall jack at either end of the network. Workstations will be connected by the six-foot piece of unshielded twisted pair wire that comes with each network connector. One port of each connector will be left empty. Terminators are not required for these short cable runs out from each wall jack. Never will more than one workstation or peripheral device be connected to a single wall jack. If the network requires a line of wiring to branch off the trunkline, to run down a hallway, for example, then a repeater is used to make an active star. In this case, the arm that is leaving the main trunkline is constructed out of solid core wire and wall jacks, just like the original trunkline. If several branches will be required, then all branch lines are brought to a central location and connected to a star controller. Each arm of the star is terminated when you have a boosted signal on a star topology.

In computer labs and the library, the cabling is located under the tables by stapling it to the wall or in specially designed plastic cable and power trays that are usually screwed to the top of the tables. In the school's administrative office and in the small elementary school examples, the cabling is located above the drop ceiling. No special cable tray is required, although many schools are being equipped with trays to make the job of running network cabling easier.

Network server and server software

We recommend that the network server be a Mac Plus, Mac SE, Mac Classic, or Mac Classic II, although any model may be used. A minimum of 2 megabytes of internal RAM will be required, but we recommend that you go to 4 megabytes for the server. We recommend that you use the MacJanet network software (with the possible exception of the administrative office and/or the library, where the software being run may require AppleShare or TOPS). MacJanet will run with either System 6.0.x or System 7. We recommend that the server be running System 6.0.7 or 6.0.8. The server should be minimally equipped with an 80 megabyte hard disk.

You do not want students fiddling with the server. It should be placed in a lockable room out of student view. If this is not possible, remove the mouse and the keyboard from the server and turn the brightness completely down so that the screen is black. The server can be left running twenty four hours a day, 365 days a year.

Workstations

Any model of Mac can be used for network workstations. Having a mix of different models for workstations on the same network is not a problem as long as the boot disk for each computer contains the appropriate system files for that model of Mac and the computers are installed with the network driver (Init) and desk accessory for a MacJanet network. Individual workstations can run different versions of the Mac operating system, although we recommend that all workstations be running at least some version of System 6. However, all workstations must be running a network driver and a desk accessory that have the same version number as the software running on the server. Workstation boot disks must also contain copies of the appropriate printer drivers for the printers that will be used on the network.

Installing applications on Mac networks is usually very straightforward. Since the printer drivers are installed on the boot disk, applications need no customizing to communicate with the printers on the network. In addition, virtually all major Mac applications are programmed to self-adjust to whatever model of Mac is being used. There is no need for the administrator to modify the application in any way to make it run properly on various models of workstations. Thus, the administrator of a small AppleTalk network has only to copy applications to the public areas on the server to make the programs available to network users.

Some specialized applications, however, are machine dependent – color scan manipulation programs, for example. For these kinds of programs, the network administrator may have to modify the application for the specific color video cards installed in the workstations. Always consult the program's manual for the specific requirements of such specialized applications.

What the administrator does on the network

The network administrator's big job is planning the network and getting it up and running. Once this task has been accomplished, the work load should decrease. There are, however, certain on-going tasks that the administrator must perform to ensure that the network continues to run smoothly. The specifics for accomplishing these tasks vary depending on the network software you are running, although we will assume that you are running MacJanet. Let's discuss the three most common tasks an administrator must perform to maintain a network.

The first task is adding new users to the network. The first set of users will be set up according to the short-term plan that was developed before the network was created. New users, however, will inevitably have to be added. The long-term plan developed before the network was created should have resulted in a network configuration that has enough flexibility to handle these new users. See the first chapter for a complete discussion of how to go about developing these short- and long-term plans. When adding new users, keep an eye on the amount of free space on the server's hard disk. If the network expands to more than 80% of the hard disk, then its time to start looking for another hard disk for the server.

The second major administrative task is taking care of existing users on the network. The most common tasks are giving passwords to those who have forgotten them, recopying damaged boot disks, making new boot disks when new print drivers, network drivers, or any other useful files are released, and clearing student work files out of user spaces that are used on a temporary basis. Fortunately, MacJanet lets the network administrator pass the task of passwords over to each individual teacher who uses the network with their classes. Boot disks are the responsibility of the administrator. So, too, is clearing off files created by temporary users. MacJanet does make this

task quite easy, though, because the administrator can highlight a group of users and specify that all highlighted spaces are to be formatted. With this one command, MacJanet will remove all files saved in each space and return it to the same condition as when it was first created.

The third task is adding new programs or new copies of existing programs. Again, you will have to refer to your short- and long-term plans when new applications or updated versions must be installed. Usually, adding new programs is straightforward as long as there is sufficient free disk space. If you must create a new space on the network for the application, remember to change the access privileges for the appropriate groups of users so they can access the newly installed program.

Installing a new version of a program currently installed on the network can get a little more complicated. The fixed volume size structure of MacJanet can be a problem if the new version of the program exceeds the size the administrator gave the program's network volume when it was created. If the new version does not fit, then the current network volume must be deleted and a new and larger one created with the same name.

What the teachers do on the network

Teachers have three main tasks to perform on a network. They prepare instructional files, place those files in a public area accessible to their students, and take care of the day-to-day tasks required to keep their students productive. If the network runs throughout the entire school, as in the case of the small elementary school, teachers may also be responding to electronic mail from their colleagues and the school's administration.

MacJanet excels at handling the needs of teachers on a network because the administrator program can be set up to provide a limited subset of features directly to teachers. A teacher cannot create new groups of users, create or delete new public or private network spaces, or even view the network configuration for any users other than their own students. But MacJanet will allow any teacher to handle most, if not all, the normal administrative demands associated with having classes of students use a computer network. Teachers can view and change passwords for students, delete files from one or more student network spaces, control the print spooler, administer electronic mail sent by their students, change the network disks that are accessed automatically when a student logs on, send log-on messages to their students, and remove log-on privileges for specific students if necessary. Teachers are also the only users who are able to add or delete files to and from the public information space for their students.

What the students do on the network

Students will use the networks shown in these examples except the network for the school's administrative office. Students are primarily concerned with getting access to programs and saving their work in their private user spaces. With small AppleTalk networks, these tasks are straightforward. There is only one server on the network so the workstation automatically communicates directly with that server without needing to ask the user for a decision. All the student has to do is enter his or her name and password to gain access to programs and his or her own user space.

Potential for future growth

These small AppleTalk networks are an ideal solution to the networking needs of small numbers of users, either in a cluster such as a lab or distributed throughout a building such as a small elementary school or an administrative office. These networks are inexpensive and easy to install. However, since they are based on the built-in networking capability of the Mac, these networks are relatively slow and therefore cannot realistically handle any more than 30 concurrent users.

Once this number of users is met, there are only two options for expansion. The first is to split the network into two separate networks, each with its own server. This will allow you to handle up to 60 users. However, users on different networks will not be able to communicate with each other and the administrator must now take care of two servers. The second option is to create AppleTalk zones using routers and/or backbones. There are several different configurations possible with these network structures; they will be covered in the next chapter.

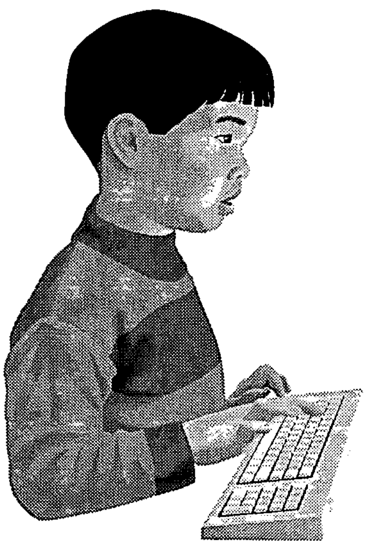

Chapter Thirteen

Complex networks for Macintosh computers

Chapter 13

Complex networks for Macintosh computers

Topics covered in this chapter...

- *Seven solutions for larger Mac networks including:*
  - *Repeaters and star controllers*
  - *Programs on workstation hard disk drives*
  - *Routers*
  - *LocalTalk backbones*
  - *Ethernet backbones*
  - *All Ethernet networks*
  - *Fiber optic backbones*
- *A complex Macintosh network example*

Introduction

We have already discussed the pros and cons of the built-in Macintosh networking system in Chapter Eleven. If you skipped over that discussion and you are new to Mac networking, go back and read the section covering the LocalTalk, LocalTalk-compatible, and Ethernet cabling systems. In this chapter we will assume you know that material.

Creating larger networks for Macs

The problem of speed

LocalTalk networks are the most inexpensive of any network available for any make of computer. While this makes them extremely popular for the majority of schools that use Macintosh computers, the slow speed of LocalTalk networks quickly causes problems when there are more than 25 to 30 concurrent users. All networks operate like a party telephone line in that all other machines must wait when one computer is sending a message on the network. With LocalTalk's slow communication speed, trying to load programs for more users simultaneously or trying to handle continuously high demand for information from the server simply overloads a LocalTalk network and results in long waits. Most schools will easily exceed the capabilities of simple LocalTalk networks when they begin to share large information resources or implement larger school-wide networking.

Another problem encountered with LocalTalk and LocalTalk-compatible networks is the limit in cable length that network signals can travel. LocalTalk cabling has a limit of 1000 feet. Phone wire-based cabling systems generally have limits of 3000-4000 feet. Cable runs that exceed these limits result in networks that run very slowly or just don't run at all because the signals fade too much.

So what do you do when you want to expand a LocalTalk network? We offer seven possible solutions. The sample network in this chapter is by no means the only way to configure a network for a large school. It will provide a conceptual base from which you can plan an expanded network. Even if you are only considering a simple expansion of your network, it is a good idea to look at how all of these examples work, so that you are better able to plan for the future.

Seven solutions for larger Mac networks

Repeaters and star controllers

Repeaters and star controllers boost the signals travelling on network cabling so they can be transmitted over much greater distances. They will not help you if you have reached the 25 to 30 user limit on your network, but they will help you run your cabling out to workstations located at the end of a large building. A repeater is ideal if your network cabling is linear and your overall length exceeds the limit for the type of cabling you are using. The repeater will boost the signal so that all computers can be "heard." The repeater is also ideal if you have to create a branch line off the main trunkline for the network – to run out to another building or to run down a hallway that branches off the main building, for example. If your building has several hallways or several different buildings, resulting in a number of arms radiating out from a central point, then the best solution is a star controller. The star controller will sit in the center or hub of the network cabling and boost the signal it receives from one arm so that the signal can travel down another arm. The star controller can handle a significantly longer cable length than the repeater. However, it costs more.

Programs on workstation hard drives

One way to address the problem of the 25- to 30-user limit on LocalTalk networks is to install hard disk drives in each of the workstations. This greatly increases the speed in which programs are loaded into each workstation and significantly reduces the amount of network traffic. The network becomes a place for accessing shared files in public areas, a place for storing personal work files in a secure private area, a print spooling service, and an electronic mail service. When providing this minimal level of service, an LocalTalk-type network can handle 50 concurrent users or more.

This sounds like this is an ideal solution for LocalTalk networks, but there are some mitigating factors to consider. First, there is the cost associated with buying hard disk drives for each workstation, if your computers are not already equipped with internal hard disks. Second, the network administrator must now take care of all of the hard disks on the workstations. This cannot easily be done from a central location, so the administrator must physically move from one computer to the next to keep the workstations running (getting out of this job was probably one of the reasons you got a network in the first place). Third, security for the programs installed on the hard disks is a problem. Setting them up to prevent illegal copying is expensive and time consuming. However, if the use of hard disks is the only way for you to handle the loading of large programs in a desktop publishing lab, for example, consider using a product like At Ease, Folderbolt, Spot-On, or FoolProof. Refer to the end of Chapter Eleven for a more complete description of these products and how to configure them for use on a network.

Routers

Routers are sometimes called bridges. Their function is to isolate network signals into specific AppleTalk zones so that the number of computers waiting for a particular computer to finish transmitting a signal on the network is greatly reduced. Suppose you had two LocalTalk or LocalTalk-compatible networks in your school, one with 20 workstations and another with 16. Each network is running fine and is meeting the basic needs of its users. But now you decide to create a school-wide network so that all users will be able to access the resources installed on both servers as well as communicate with each other via electronic mail. When the two networks are joined together by a connecting cable the total number of users on the new combined network jumps to 46. Now the wait for a response from either server would be so slow that network users would get very frustrated. The reason is that 45 computers on the combined network would have to wait when one computer was communicating with one of the servers. This would be unacceptable.

The solution is to use a router and create two AppleTalk zones. You start by splitting your network into two sections and connecting one end of each section to a router. Each section is called a zone and is given a unique identification number. The router isolates network signals generated by a workstation in one zone that are directed at a server in the same zone. Only the workstations in the same zone have to wait while that communication takes place. The workstations in the other zone can be involved in a completely different communication. The router only lets signals pass that are intended for a computer in the other zone; for instance, when a user wants to access some clip-art from the other server in the combined network.

You must be very careful not to create a daisy-chain of AppleTalk zones. When several zones are connected in a line, a single signal from a workstation in a zone at one end intended for a server in the zone at the other end will cause every workstation to stop and wait until the communication is completed. As soon as you need to create three or more zones, you should create a backbone.

LocalTalk backbones

A backbone is like a highway between two or more cities. It lets signals travel quickly between networks without interfering with local network communications. No workstations are connected to the cable that forms a backbone. The only devices connected to the backbone are routers for each AppleTalk zone. The only signals that travel on the backbone are those which originate in one zone, but are intended for another. The beauty of the backbone topology is that inter-zone signals affect only the two zones involved in the communication, because the routers for all other zones do not let the signal pass onto their cabling.

You can create a backbone using phone wire-based cabling. All that is needed are the routers for each zone and a piece of phone wire to make the backbone. Repeaters are often used in conjunction with a backbone topology to ensure that signals are strong enough to travel from a workstation to the backbone. Fading signals can be a problem when the backbone is located a considerable distance from the workstations. The repeater is usually placed in the same room as the backbone. The repeater is not connected to the backbone, but rather placed between the workstations and the router for that particular zone.

Ethernet backbones

Backbones are an effective way to handle more than 25 to 30 concurrent users on Mac networks. They are also effective for creating a school-wide network that allows users to communicate with each other and access the resources located on other servers in the building. However, a simple LocalTalk backbone is usually not adequate to meet the inter-zone communication needs of a large combined network. Any inter-zone signal must travel along the cabling for the zone where the communication originates, pass through the router onto the backbone, pass through the router for the destination zone, and finally along the cabling in the destination zone. Each step in this process takes time and the wait may be unacceptably long. To address this problem, backbones are commonly made with a cabling system that has a higher speed than LocalTalk. The most common cabling used for backbones is Ethernet. Ethernet passes signals along the backbone so quickly that users may not notice any significant increase in response time when communicating with servers in other zones. Ethernet backbones are an excellent solution for expanding networks or for creating a school-wide network because they provide a means for breaking the user limit of LocalTalk-type networks while also ensuring that inter-zone communications do not take so long as to become unworkable. In addition, Ethernet backbones are relatively easy to install and are less expensive than buying hard disk drives for each workstation or running Ethernet to every computer on the network.

With a backbone topology, network administrators may find it very useful to connect their own personal workstation directly to the Ethernet backbone. In this way, the administrator has to make a jump through only one router to get access to the server and workstations in any zone. Network administration is much faster. The administrator can also monitor traffic on the backbone very easily.

All-Ethernet networks

All-Ethernet networks offer one major advantage when creating school-wide Mac networks that you don't have with LocalTalk – speed. Ethernet is over 40 times faster. There is no problem with many concurrent users and several servers all running on the same cabling. Thus, it is very easy to create a school-wide network because you simply run the Ethernet cabling and connect your workstations. No other devices are required. If this sounds too good to be true, you must be aware of a major negative factor: the cost of the Ethernet interfaces required for every workstation. However, if you have enough money to consider the purchase of individual hard disk drives for every workstation on a network, then you should definitely look at the all-Ethernet option before proceeding.

All-Ethernet networks do not have to be installed throughout an entire school. If an Ethernet backbone has been installed in the school, then you can choose to upgrade to Ethernet in only one zone. You may be able to use existing wiring for the new Ethernet network. Besides purchasing the Ethernet interfaces for each workstation and peripheral device in the zone, you will also need a new router to connect the Ethernet zone with the Ethernet backbone.

Even when you have upgraded every workstation in the building to Ethernet, you will probably want to keep the backbone topology for your network. Despite the faster communication speed, Ethernet networks can bog down when too many devices are trying to "talk" on the same cabling. The backbone will isolate network signals in order to minimize the number of machines waiting while a message is being sent, so that the quick response from the network is maintained. An Ethernet backbone for an all-Ethernet network will probably meet the needs of a school-wide network for many

years. However, even Ethernet may not be fast enough for school-wide networking, especially if the network is being used to transmit disk-based live video out to the workstations. The next step is to incorporate fiber optic cabling into the network.

Fiber optic backbone

The easiest way to utilize the blinding speed of fiber optic cabling is to use it for the backbone of a network. The individual zones in the network can be running either LocalTalk or Ethernet. However, if the network will be transmitting disk-based live video out to the workstations, those workstations must be running at least Ethernet to be able to handle the volume of information. The server will likely be connected directly to the fiber optic backbone to ensure a high-speed response anywhere in the school. A concentrator will be required to convert the optic signal to either Ethernet or LocalTalk.

As far-fetched as it might seem, the information that will soon be travelling over network cabling may overwhelm even a fiber optic backbone network. When this happens, the only option is to consider an all-fiber optic network or at least an all-fiber optic zone. These networks are still very costly and are not common. But the need for this kind of network speed can be easily projected over the next few years.

A complex Macintosh network example

Here is a network for a large school that incorporates a number of the solutions we have just discussed. This network will provide school-wide access for students, teachers, and administrators. Ethernet is used throughout the school. While there are other ways to configure a network for a school this size, this example is a good solution with plenty of room for growth.

The file servers

In a school this size, there are file servers for various purposes. The three computer labs have local servers running MacJanet that provide students access to instructional programs and their own private work spaces. These servers also allow teachers to place instructional files in public areas for student access. The Library has a server running AppleShare. AppleShare is required by the MacSchool Library Administration software. The server in the general office is also running AppleShare for the MacSchool Administration software being used for handling administrative tasks for the school. The server in the north utility room is connected directly to the backbone. In this way, any workstation in the school can access this server by making a jump through a single router. This server is running MacJanet and provides school-wide access to programs and information resources.

Unlike LocalTalk networks, Ethernet can benefit greatly from a having a powerful computer as a server. Speed is required from both the computer's processor and the hard disk drive. It is not uncommon for a Macintosh Quadra or Centris to be used as a file server on Ethernet networks.

The workstations

Any model of Macintosh can be used on this network as long as it is equipped with an Ethernet interface. For older Macs like the 512, Plus, Classic, Classic II, Portable, as well as the new Powerbooks, an external interface that connects to the SCSI port is required. All other models will need an interface card that fits into one of the slots inside the computer.

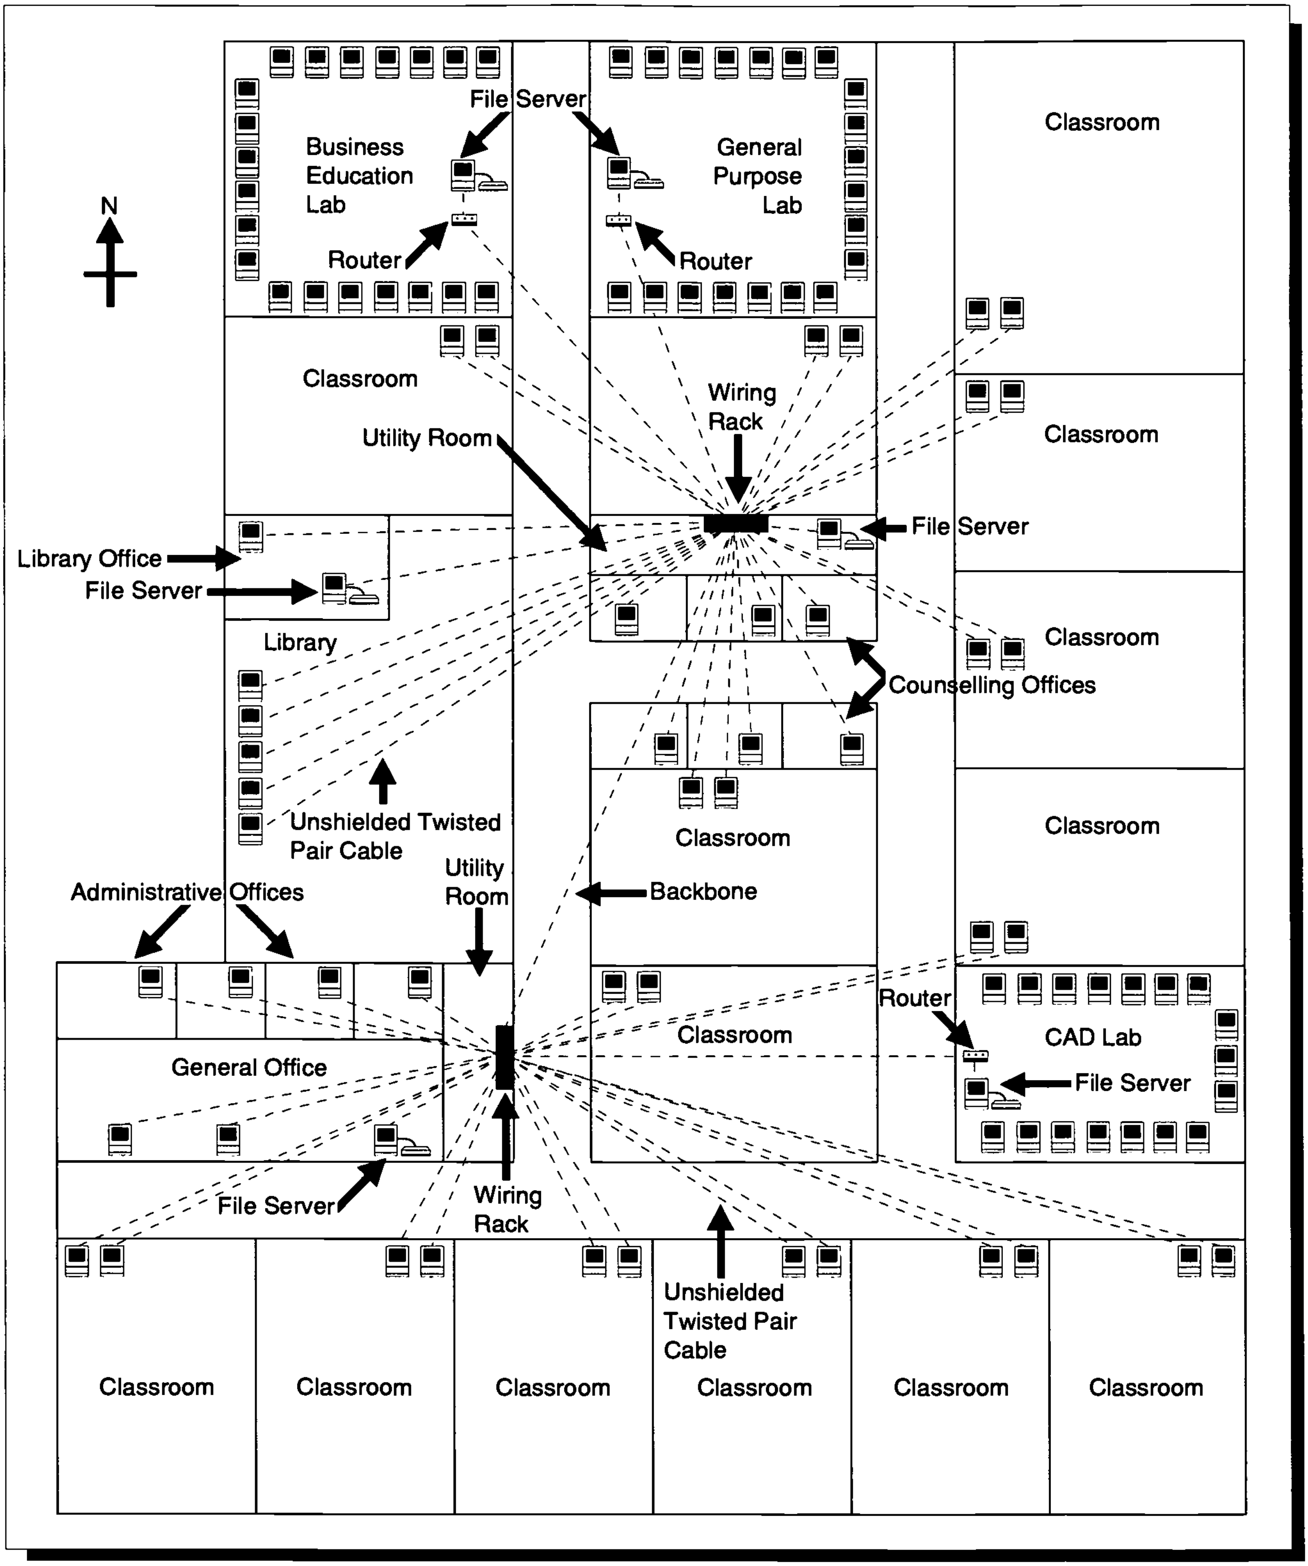

This large-school network uses Ethernet throughout to ensure quick a quick response when requests are made of file servers. The wiring racks in the utility rooms are connected by a backbone. See the next diagram for more detail on the backbone wiring. The labs in this school can be configured at least two different ways. See following diagrams for detail on both the Ethernet and LocalTalk options for the labs. The networks in the labs are separated by routers so communications in the lab targeted for the local server do not affect any of the computers elsewhere in the school. The six servers run MacJanet and AppleShare simultaneously over the same cabling.

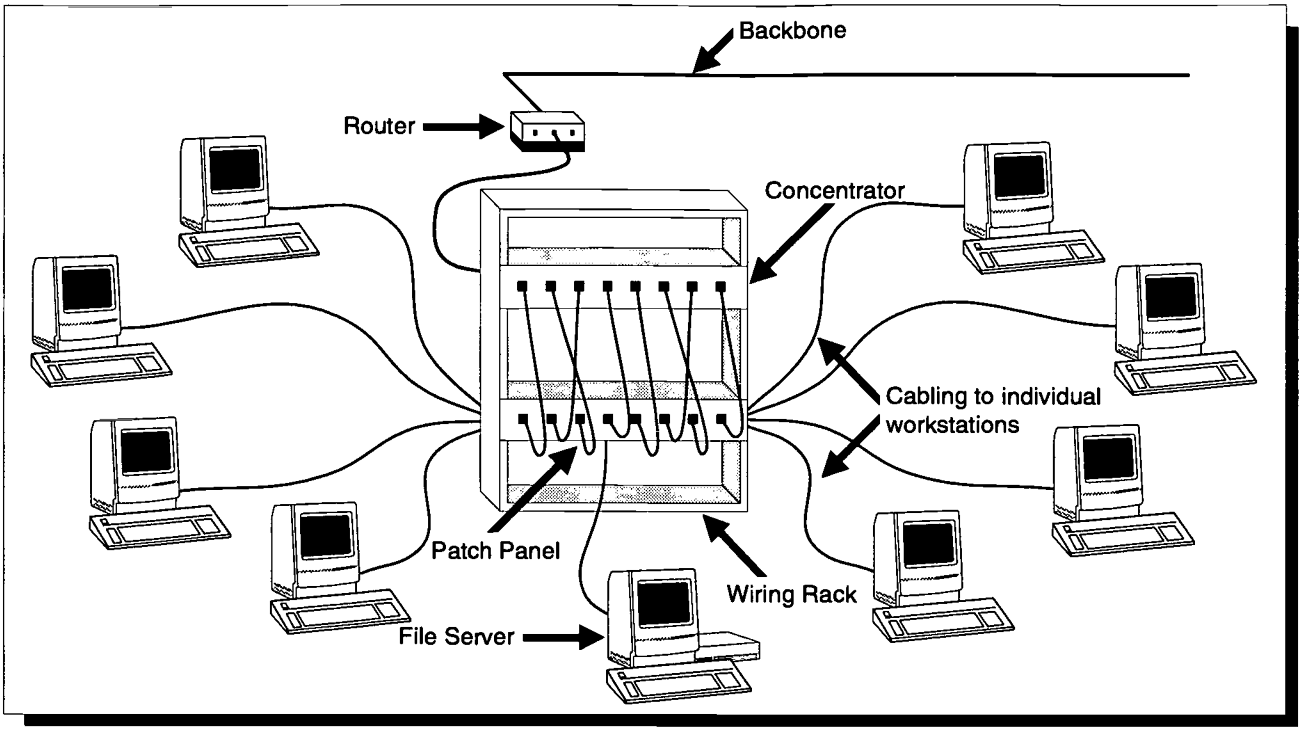

The backbone in the school makes faster access to the school-wide server possible. The server in the north-end utility room is connected directly to the backbone. Routers are used to keep communications targeted at local servers from affecting the school-wide server. Here is the configuration for the south-end wiring rack. The file server is located in the general office and is used only for administrative purposes. The router connected to the back of the concentrator keeps requests for information from the office server from passing onto the backbone and affecting other workstations from accessing other servers in the building.

The cabling system

This network is based on a structured wiring system with wiring racks in two utility rooms connected by a backbone. All wiring is done with category 5 Unshielded Twisted Pair cable for a 10base-T Ethernet system throughout the school. Cabling from wiring racks to individual workstations and the backbone cabling is run in cable trays located above the drop ceiling in the halls and classrooms. RJ45 wall jacks for workstations are located in each room. The wiring racks in the utility rooms contain patch panels and 10base-T Ethernet concentrators.

There are two configurations possible for wiring the computer labs. The first is a structured wiring system radiating out from a wiring rack in each of the labs. The lab network is separated from the rest of the school-wide network by a router so that requests from lab workstations for information from the local server do not cause delays for other computers throughout the school. The second configuration for the labs is a LocalTalk network. While this seems like an unacceptably slow solution, if each of the workstations is equipped with a hard disk drive containing the programs to be used in the lab, the network traffic will be limited to saving and retrieving work from private work spaces, and accessing instructional files from public areas. This LocalTalk configuration is much less expensive than the Ethernet network.

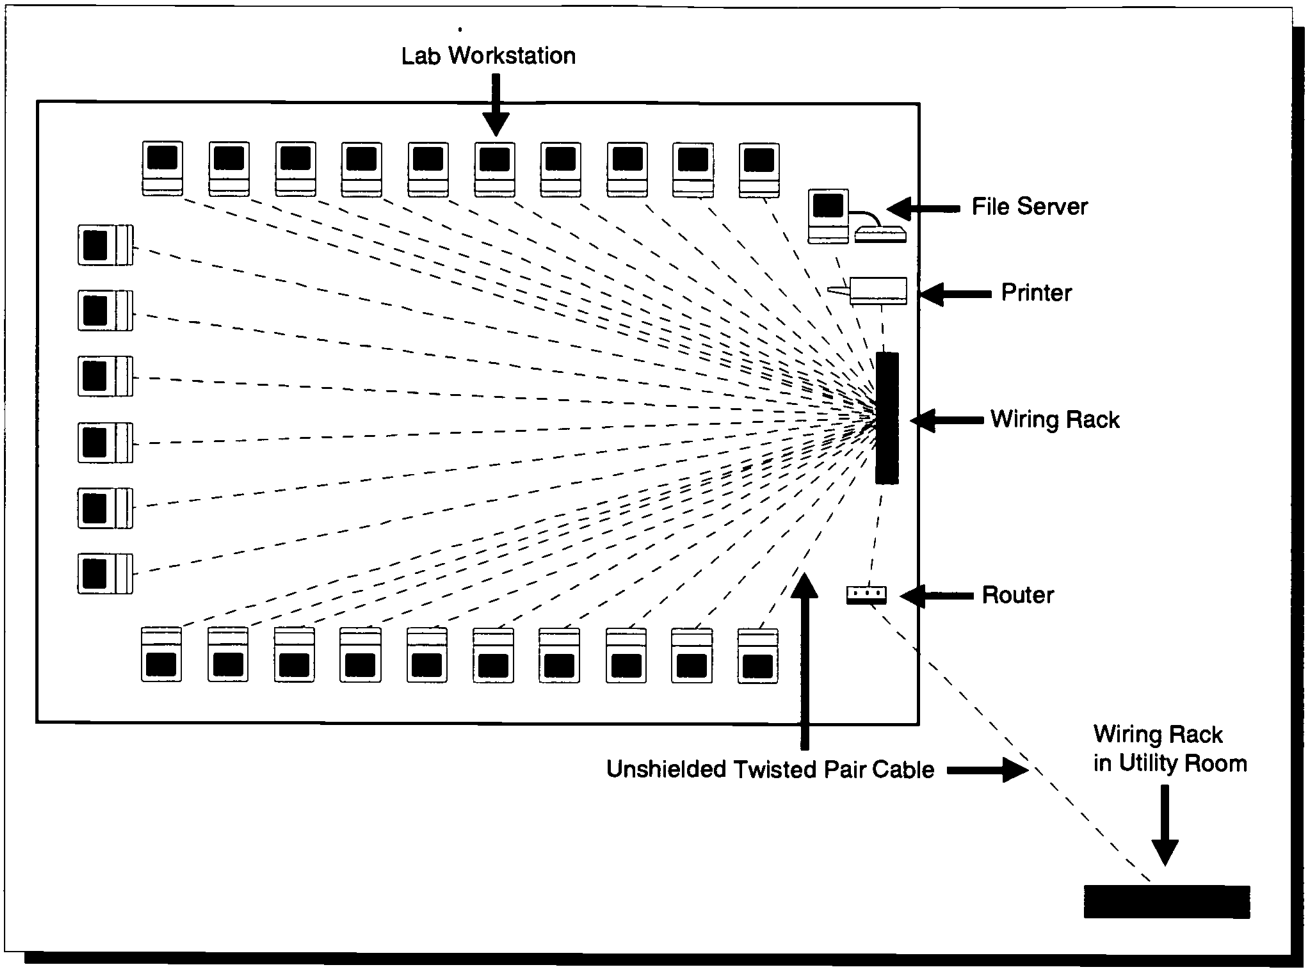

Here is an Ethernet network configuration for a single computer lab using a structured wiring system. Ethernet is much faster than LocalTalk and is recommended. The wiring rack contains three 10base-T hubs. The Unshielded Twisted Pair cabling to the workstations can be run under the tables or over the ceiling. The router can be placed either in the lab or in the utility room.

What the administrator does

Administrative tasks for this network are the same as for smaller networks with the added complication of multiple file servers and routers. In addition to setting up the usual access privileges, now the administrator must determine the level of access users from other network zones will have when accessing a particular server.

Moving around a large building can become quite a problem when a network requires some re-configuring. A better solution is to access servers via software. Some administrative tasks can be done by simply logging on from any workstation. However, other tasks require more control of the file server computer. Those tasks can be accomplished by using products like Timbuktu from Farallon or ScreenLink from Data Watch. These programs allow you to take control of another computer over the network cabling, just as though you were sitting there using the keyboard and mouse. In this way, the administrator can perform the majority of his or her tasks without having to leave his or her office.

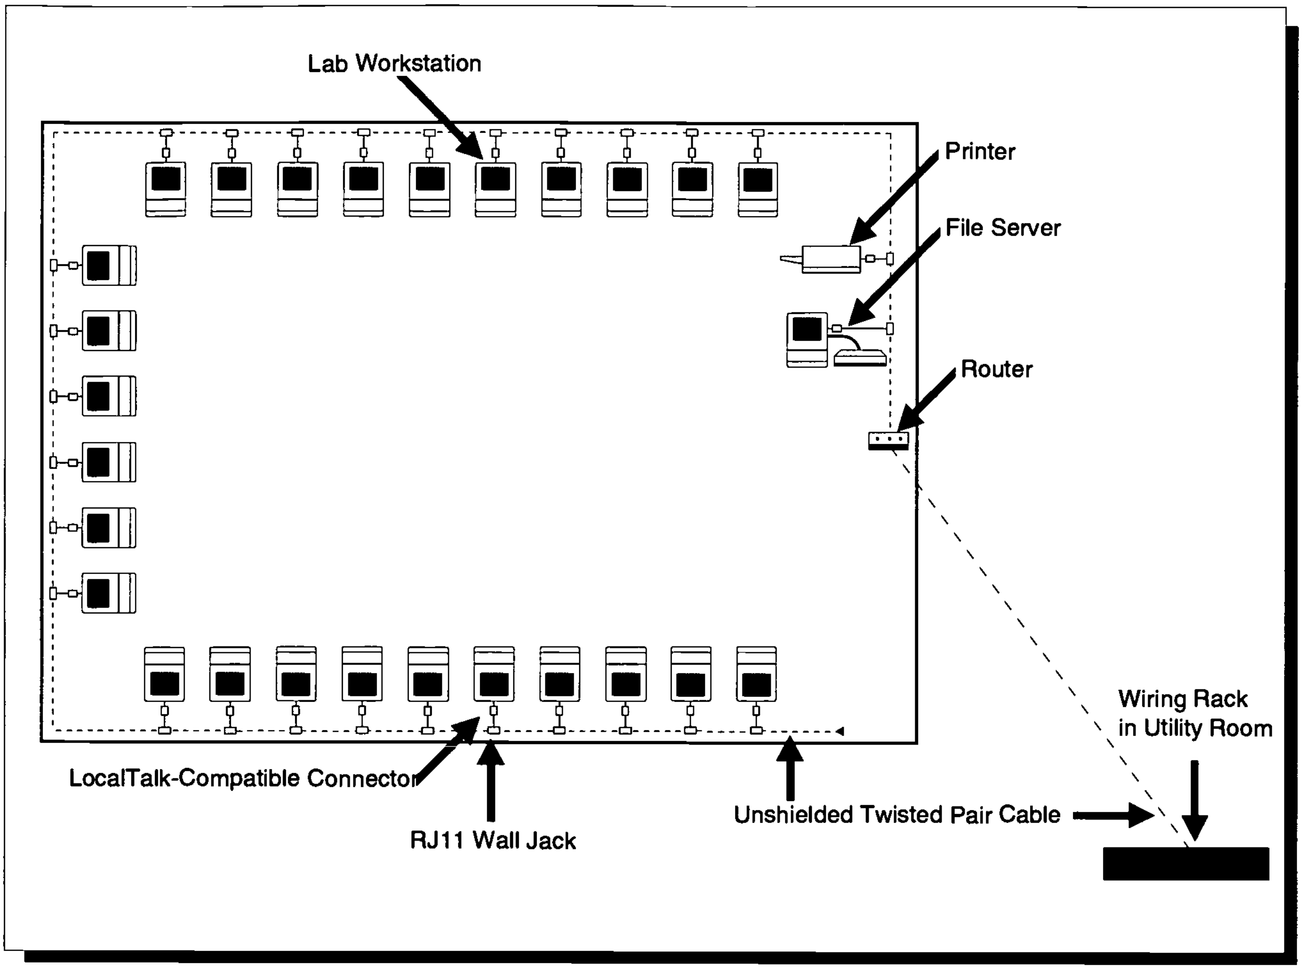

Here is the same computer lab configured as a LocalTalk trunkline network. This network is much slower than Ethernet, but it is also much less expensive. It can run at acceptable speeds if the programs to be used in the lab are installed on the hard disks in each workstation.

Another program that is very useful for an administrator is GraceLAN Update Manager from Techworks. This product allows you to add, delete, or replace files on all workstation hard disks from a single computer. It is extremely useful when updated versions of programs are received. It can also be used to remove unwanted files from all computers on the network – games or unwanted Inits, for example.

What the teachers do

Teachers will usually use only one of the file servers on the network. The only difference they will notice on a large network like this is the additional choices for the AppleTalk zone and the file server name. The routers create zones on the network, each with a unique name. Once a zone has been selected, all servers in that zone are displayed. If there is only one server in that zone, it is selected automatically. Boot disks can be configured to remember zone and server names, making the log-on process much easier. Once they have accessed the network, teachers can use programs, access information resources, and place files for students in public areas as they do on other networks.

What the students do

Students must make the same selections of zone and server when they log-on as well. The great advantage that a school-wide network has is that students can access information resources and personal work saved on other servers in the school without having to move to another location. All that is required is for the student to access another server over the network cabling using either a Desk Accessory or the Chooser. While the network can be used to access information and personal work, students should be taught to run programs only from the local server or their own hard disk. Running programs from servers in other zones is not only slow for the student, it also causes unnecessary delays for other users.

Potential for future growth

The beauty of a structured wiring system is its expandability. This network can be expanded to hundreds of workstations by adding more wiring racks. The only requirement is the use of routers to create zones of users to isolate local communications. If the network were to become sluggish due to an increase in usage, the Ethernet backbone could be replaced with a fiber optic backbone. Most concentrators can handle connections to fiber optic cabling.

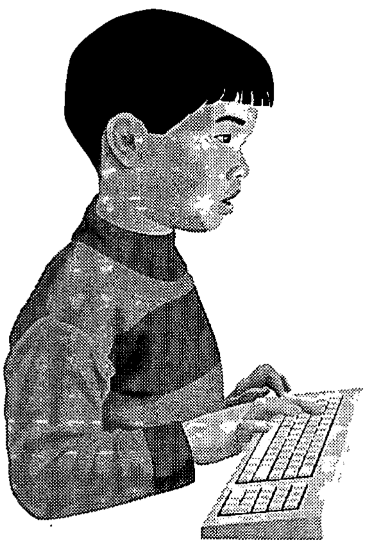

Chapter Fourteen

A complex network for IBM-compatible and Macintosh computers

Chapter 14

A complex network for IBM-compatible and Macintosh computers

Topics covered in this chapter...

- *A discussion of complex networks for Macs and IBM compatibles*
- *A technical description of a network example including:*
  - *The file servers*
  - *The workstations*
  - *The cabling system*
  - *What the administrator does*
  - *What the teachers do*
  - *What the students do*
  - *A discussion of the potential for future growth*

Introduction

This network brings together features found in many of the previous examples. The layout of the computers in the school is almost identical to the examples in Chapters Ten and Thirteen, while the networking hardware and type of wiring is the same as the example in Chapter Nine. If you have not read those chapters you may want to do so now, as this example builds on those simpler ones. The wiring configuration outlined in this chapter will support both Macintosh and IBM-compatible computers. This Ethernet wiring system will allow you to use any of the three major microcomputer network software packages: Novell Netware, AppleShare, or MacJanet software. You can even use more than one at the same time.

The designer of a complex network such as this has many decisions to make. One of the most important is the choice of server software. Networking a school that uses both IBM-compatible and Macintosh computers is a real challenge, especially if you want some school-wide services like electronic mail and CD-ROM access. But it is possible to provide these services to both hardware platforms simultaneously on the same wiring. Here is what we suggest.

First, when the computer system in a school reaches the complexity of the example shown in this chapter, there is an acute need to isolate network traffic in order to keep response times for the entire school-wide network at acceptable levels. This will require routers. You can, of course, buy individual routers for each zone of your network, but you may find it much more manageable to install a Novell Netware server instead. Novell has programmed its software to act as a router between the interface cards installed in the server. You can run different zones of the network into

the server and Netware will only let inter-zone network signals travel between the interface cards in the server. Netware can even translate the signals if the zones use different wiring and different communication protocols. By using the server as the routing device, network traffic will travel more quickly and the cost of creating new network zones is greatly reduced.

Novell Netware is an obvious solution if you are using all IBM-compatible computers or mostly IBM compatibles with only a few Macintoshes. What is surprising is that a Novell server is a great solution for all-Macintosh or mostly-Macintosh networks as well. On Mac networks, Netware is used for its routing capabilities and MacJanet or AppleShare can be installed on another server on the network. Users will not be aware that Netware is running. They will simply log-on to the MacJanet or AppleShare server as normal. Furthermore, many Mac applications will run directly on Netware because Novell has programmed it to support the Apple File Protocol.

Let's summarize your options. If the network will have mostly IBM-compatible computers with only a few Macintoshes, Novell Netware is all you will need. If there are more Macintoshes, add either a MacJanet or AppleShare server for your Macs. As discussed in earlier chapters, the authors believe that MacJanet will probably fit better into most instructional settings than AppleShare.

Where to use this network

The network shown in this chapter is appropriate for a school which has a combination of Macintosh and IBM-compatible computers. The Macintoshes may be in the humanities classrooms and a general purpose lab and the IBM-compatible computers in science classrooms and in Business Education and Drafting labs, but any other combination is possible. This network could be developed by adding Macintoshes to a previous IBM-compatible network, or the other way around. Studying this example may be instructive even if you have only one type of machine in your school, as you can design a network flexible enough to adapt to other types of computers later.

Description of the network

Diagram

The layout of the computers in the school is identical to the previous two examples, except here is a mixture of IBM-compatible and Macintosh computers. All computers shown in the following diagram have Ethernet adapters. There are five wiring racks: one for the north end of the school, one for the south end of the school, and one in each of the computer labs. The Ethernet hubs in these wiring racks are connected to the computers with unshielded twisted pair cable.

The file servers

Five of the seven file servers in this example are IBM-compatible computers running Novell Netware. The other two file servers are Macintosh computers running MacJanet. Each of the Novell servers has two network interface cards to take advantage of Netware's internal routing feature. Two Ethernet network adapters are installed on each of these servers, one with a twisted pair connection to a 10base-T hub, and a second with a connection to the coaxial cable backbone. Each file server has two printers, except the one in the CAD Lab, which has one printer and one plotter. Any file servers which have Macintosh programs run a special program to enable this access, while any file servers which have programs for IBM-compatible use have IBM ICLAS installed. Both can be used on the same server.

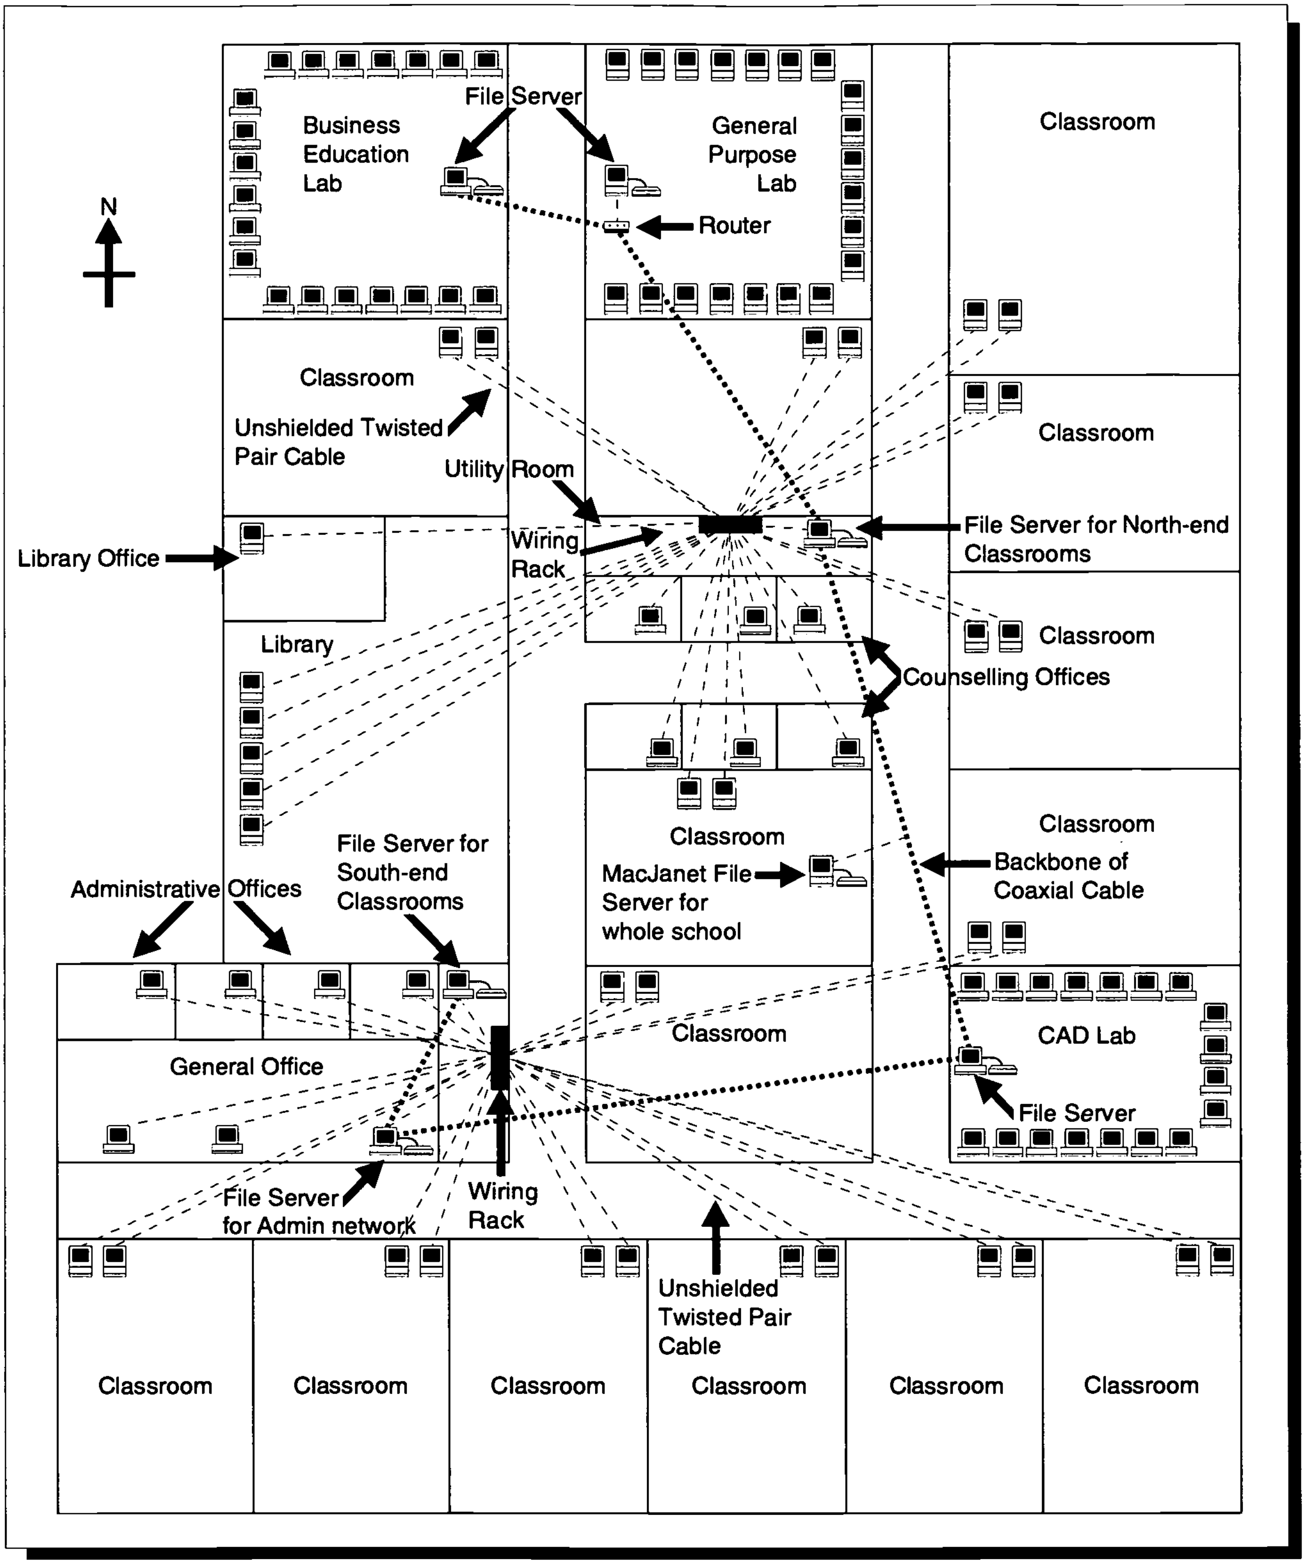

This large-school network is based on a backbone of coaxial cable that connects all the file servers in the building. The backbone, when used in conjunction with routers or servers acting as routers, can isolate signals so that several workstations can communicate with their local server simultaneously. See the next diagram for more detail on connecting the backbone cabling. Cabling from the wiring racks to workstations is Unshielded Twisted Pair cable. Five of the file servers run Novell Netware; the other two are Macintoshes running MacJanet,. The servers in the labs contain two network interface cards so communications in the lab targeted for the local server do not affect any of the computers elsewhere in the school. The server connected directly to the backbone would actually be located in one of the utility rooms.

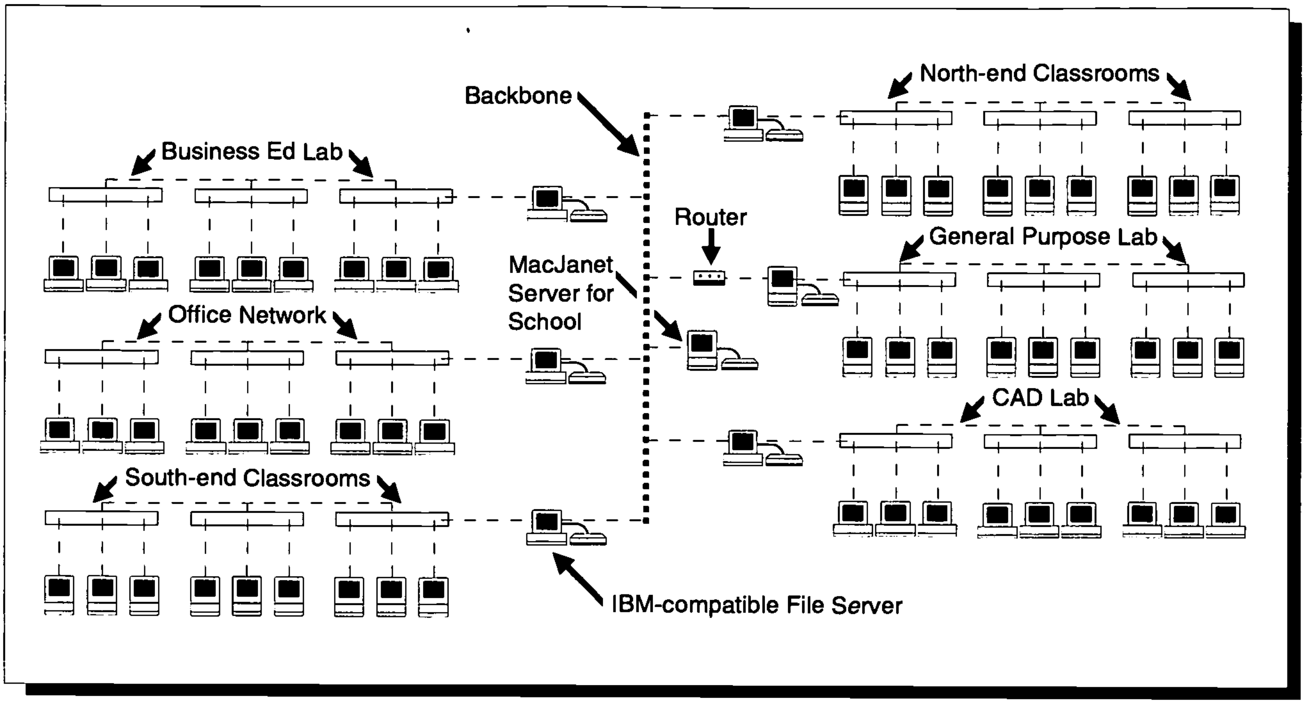

This diagram shows the configuration of the backbone. Notice that the Novell servers do not require routers. Also notice that any workstation can access the MacJanet server connected directly to the backbone with only a single jump through a routing device. This configuration speeds the response of the MacJanet server throughout the school.

Setting up the software on the servers will be quite complex, as there will be programs for IBM-compatible users as well as for Macintosh users. While the programs are quite separate and cannot be used on the other type of computer, some of the data will be transferable. If a program has a version for the Macintosh and a version for the IBM-compatible it may be able to exchange files between the two. For instance a file created by Microsoft Word on an IBM-compatible computer could be opened by Microsoft Word on a Macintosh computer and then put into a Macintosh desktop publishing program. Since the two computers use different rules for naming files, an IBM-compatible user accessing Macintosh files with long names will see the file name shortened to the DOS limit of eight characters. As this can be quite confusing, Macintosh users should take this into account when thinking up names for files that may later be accessed by IBM-compatible computers.

A MacJanet server can be placed anywhere on the network, but we suggest connecting it directly to the backbone as shown in the diagram. Connected to the backbone, this server can be accessed by any workstation in the building with no more than a single jump through a router (or a server's network interface card acting as a router). The MacJanet server will provide CD-ROM access and private user spaces for Mac users. Mac applications can be stored on this server or on the Netware server. Since the wiring throughout the school is Ethernet, Mac users will not notice any problems with response time when accessing the MacJanet server from anywhere on the network. In fact, if they are used to the speed of a LocalTalk network, they will be pleasantly surprised at the faster response on Ethernet.

The workstations

The teachers and students using Macintosh computers on Novell do not have network server software designed for education - they are using "raw Netware." This fits in well with the Macintosh desktop: they see folders and applications on their desktop almost as if they were on a local hard disk. The network administrator has control over what programs users can access. If they do not have the rights to a particular program, they just don't see it on their desktop.

Since Macintosh computers cannot boot remotely, they will require boot diskettes with the appropriate drivers and Init files. If the workstations have no hard disk drive, then System 6 boot diskettes must be used. The drivers and Inits in the System 7 system folder are too large to fit on a diskette, so users of System 7 must have a local hard disk. System 6 can also be installed on hard disks. Refer to Chapter Eleven for a more complete discussion of when to use System 6 or 7. Local hard disks make network management more difficult. For example, with Novell Netware 2.2, a student can easily drag a program file from the network to a local hard disk. This not only puts unnecessary traffic on the network, but will require the network administrator to go around and clean up all the files left on the local hard disks. It is also illegal if the student then copies the program onto a diskette and takes it home. Netware version 2.2 cannot control Macintosh users in the way it can control IBM-compatible users, but Netware 3.11 can. For this reason, Macintosh users should purchase Netware version 3.11 or later.

The IBM-compatible workstations are set up exactly the same as those on other Novell networks (see Chapters Nine and Ten). They make use of the software on the servers through ICLAS or through some other menu program, and can choose which server to use by "passing through." Since these workstations can be set up to remote boot, they will not require boot diskettes.

The cabling system

The following table shows how many servers and workstations are connected to each wiring rack, and the number of hubs required. The number of hubs is calculated by dividing the total number of connections by 12 (assuming each hub has 12 ports):

	File server connections	Workstation connections	Total connections	Hubs required
Utility Room North	1	26	27	27/12 = 3
Utility Room South	2	22	24	24/12 = 2
Business Ed Lab	1	24	25	25/12 = 3
GP Lab	1	30	31	31/12 = 3
CAD Lab	1	20	21	21/12 = 2
Total	6	122	128	13

The complete network will require 13 Ethernet hubs. This is not the most economical solution, as it could be done with fewer hubs, but this solution separates the networks very clearly and provides some redundancy. In addition, the school may not require all these hubs when the network is first installed. It is not unreasonable to install all the cabling in the walls when the school has only half this number of computers, purchasing only half as many hubs. Then, as more computers are added, more hubs can be purchased. Note that the total of 13 Ethernet hubs allows considerable expansion, without adding more hubs.

The Ethernet Backbone is a coaxial cable which connects each of the servers in the school. Since this is coaxial cable, one of the Ethernet cards in each of the file servers must have a connection for coaxial cable. The advantage of using a backbone like this is that the network is broken into six smaller networks. This is shown more clearly in the previous diagram. The servers act as routers, and only pass information onto the backbone if it is going to a different network. With this topology, most of the traffic stays within the smaller networks, and the performance is not affected by the large number of workstations in the school.

In a very large school the length of the backbone may exceed the length limitations of thin coaxial cable. If so, there are three options. Repeaters could be added to allow a longer backbone, thick coaxial cable could be used, or fiber optic cable could be used. The third option is preferable, and since the cost of using fiber optic cable is quickly coming down, we can expect to see more backbones using fiber optic within a few years.

What the network administrator does

Administering a network as large and complex as this one is neither simple nor quick. When combined with the task of providing computer support and in-service for the teachers in the school, this role should comprise at least a half-time position. Some of the tasks are done using a Macintosh workstation, others are done using an IBM-compatible workstation. The most common practice is to manage each file server independently, with each user given a separate ID and password on each file server they need to use. There are now products available which automate this process, letting a user be added only once with access to all servers. Novell calls this the Netware Naming Service.

There are Macintosh versions of the administration programs used to add users, work with files, manage security and carry out most of the other tasks of the network administrator. Macintosh software is usually installed on a local hard disk and then copied onto the file server. Software not designed for a network will then require some “tweaking” by the administrator. Since the Macintosh computers cannot use remote boot, the network administrator must also create boot files for all the workstations, either on hard disks or floppy diskettes. In all of these operations, the administrator is well advised to seek out the advice of other schools using Macintosh computers on Novell networks.

The network administrator uses an IBM-compatible workstation and the ICLAS program (or other menu program) to manage the IBM-compatible programs and users. Administrative tasks, such as adding users, are exactly the same as with an IBM-compatible-only network with multiple file servers,. Since the IBM-compatible computers will remote boot, the network administrator will put the appropriate boot files on the file server.

What the teachers do

For classes using Macintosh computers, the teacher probably does the same as a student. The teacher does not determine what programs are available to their students - all the students see all the programs. While this could be changed, it would probably be changed by the network administrator, and the teacher would not have any additional duties.

For classes using IBM-compatible computers, the teacher does the same as with any other ICLAS network, as described in previous chapters. This includes choosing which students are in the class and which programs and information resources should be made available to those students.

What the students do

Students using Macintosh computers just see the programs in folders and choose them as they would from a local disk drive. If they are not given space on the file server to store their files, they are not much influenced by the fact that their computer is on a network. If they are given space on the file server, it will likely be quite limited, and since Macintosh users are used to working with very large files, they will have to be careful with their disk space. Any students using large graphics files will probably have to store them on diskettes.

Students using IBM-compatible computers choose programs from menus just as they do with the IBM-compatible networks in the previous examples. The only difference the IBM-compatible users may notice is the presence of files they created with Macintosh workstations. The names of these files will be shortened from their Macintosh versions, and may or may not be compatible with their IBM programs.

Potential for future growth

This network has already grown substantially from the first example in this book, but it is by no means at its limit. Adding more workstations, either Macintosh or IBM-compatible, is a simple process of connecting to an existing Ethernet hub or adding another hub if necessary. If more file servers are added it will be most convenient to put them beside present servers to make it easier to connect them to the backbone. The limitations of networks such as this will come when even higher speeds are required for multimedia applications. The computer networking industry is now working on ways of boosting the speed of these networks, to extend their life into the next century.

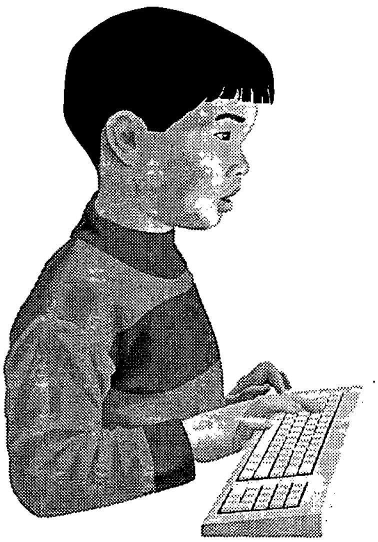

Chapter Fifteen

Networking for Apple II computers

Chapter 15

Networking for Apple II computers

Topics covered in this chapter...

- *The rags to riches story of Apple Computer Inc.*
- *Networking for the Apple IIe computer*
- *Networking for the Apple IIgs computer*

Truly a rags to riches story

The development of the first Apple computer and the emergence of Apple as a major producer of microcomputers is one of the most remarkable stories in the history of business. In the mid 1970s, two college students with an interest in electronics decided to build their own computer. By selling a car, Steven Jobs and Steve Wozniak raised $800 in capital and proceeded to create the very first Apple computer, in a garage. Within a few months, their computer design had generated enough interest for these young entrepreneurs to obtain a substantial amount of venture capital. By the late 1970s, they had developed the Apple II and Apple II+ computers as well as disk drives and a disk operating system. The design of the Apple II computers was innovative, with high resolution color graphics (at least compared to the other microcomputers available at that time), easy connection to TVs to act as large color monitors, and an open architecture with expansion slots that encouraged third-party electronic firms to develop devices that extended the usefulness of the computer. By the mid 1980s, Apple Computer Inc. had reached annual sales in the billions of dollars. In less than seven years, Steven Jobs and Steve Wozniak had taken their computer company from a garage to a major multi-national corporation. Today, Apple has become the world's largest single producer of microcomputers, even surpassing the perennial leader – IBM. The meteoric rise of Apple Computer Inc. is a story that has few, if any, equals in the history of business.

Apple II computers were used in both in business and education, but it was education that overwhelmingly embraced these computers and gave Apple its major market. Even today, no other computer has more educational software than the Apple II line. In fact, Apple IIe and Apple IIgs computers are still being sold, solely on the sales to schools. This is truly remarkable when you consider that the basic design of these computers is over 15 years old!

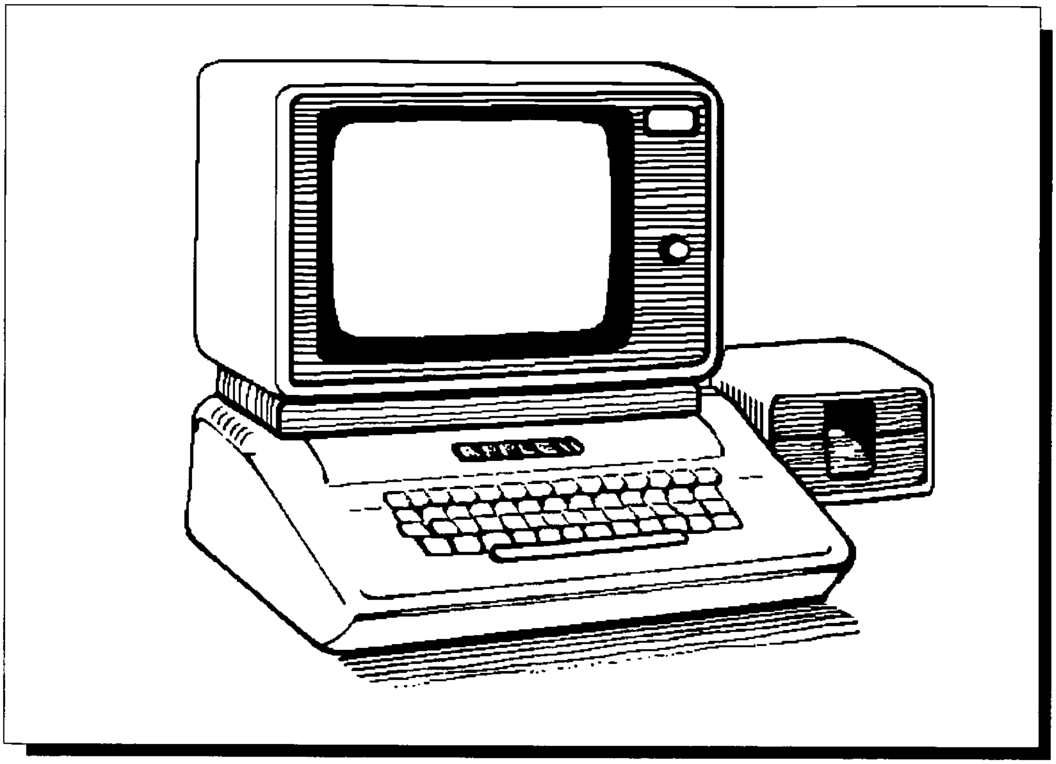

Starting in a garage, Apple Computer Inc. had a humble beginning. But Apple II computers caught the fancy of educators and the company rode the phenomenal success of these machines to become the world's largest single producer of microcomputers.

Networking the Apple IIe – the cold hard truth

Apple IIe computers can be networked using AppleShare. A Macintosh computer is required to act as the server. However, the Apple IIe does not come equipped with built-in LocalTalk capability like Macintosh computers, so you must install a workstation card in each machine. This card, usually installed in slot 7, makes connections to an AppleShare network possible using LocalTalk or LocalTalk-compatible cabling and connectors. Unfortunately, Apple no longer sells these cards so if you are not already running an Apple IIe network, you're out of luck.

In addition to the lack of networking, there is another important issue you must face if you are currently using Apple II computers. The Apple II line is a dead-end product line. When Apple brought out the Macintosh in 1984, it signalled a new direction for the company. No significant development has taken place on the Apple II computers for a number of years. Apple's present focus is clearly on the Macintosh. Furthermore, Apple will soon bring a whole range of new computers to the market based on the PowerPC processor chip, so even the long-term future of many Macintosh computers is questionable. It is now time to move on from Apple II computers. If you have Apple II equipment in your school, you should be making plans to replace those machines in the near future. It is cheaper to buy a Mac than it is to buy an Apple IIe. If your investment Apple II software is substantial, then consider the purchase of a Macintosh LC II or LC III. These computers can be equipped with a card that enables them to run Apple II programs.

Networking the Apple IIgs

Like the Macintosh, the Apple IIgs is equipped with AppleTalk ROMs that make it possible to connect the computers to a LocalTalk network without the need for an additional network interface. Both AppleShare and MacJanet support Apple IIgs networking, although neither system gives the IIgs all the features given to Macs. As with the Apple IIe, you must face the fact that the Apple IIgs is a dead-end product. Although newer than the IIe, the IIgs will suffer the same fate. Apple is putting its energies into the Macintosh line, PowerPC computers, and Newton hand-held personal electronic assistants. The Apple IIgs should not be the basis of any long term technology plans for schools. We recommend that you spend as little as possible on this equipment and begin planning for its replacement immediately.

The file server

A Macintosh computer must act as the file server for Apple IIgs networks. You may choose MacJanet or AppleShare network software since either will communicate with the IIgs. AppleShare offers a feature called the Aristotle Menu Software for managing Apple II computers on the network. Aristotle gives users basic networking features such as access to programs and information files.

Cabling

Apple IIgs computers are designed to work with LocalTalk or LocalTalk-compatible cabling systems. Either daisy-chain or trunkline topologies can be used. Apple IIgs computers can be connected to existing Macintosh networks. Ethernet, Token Ring, and fiber optic cabling systems are not available for Apple IIgs equipment.

Workstations

Networks offer Apple IIgs computers access to basic features – access to programs and data files stored on the server and print spooling. However, not all features of MacJanet or AppleShare are extended to the IIgs.

The IIgs is not capable of remote booting and each computer will require a boot disk. One of the strengths of Apple is its use of installer programs for creating and updating the operating system on disks. Once the operating system has been installed, you must also add the appropriate network drivers for the network software you have chosen. Accessing the network from an Apple IIgs is similar to a Mac – the Chooser is used with AppleShare and a desk accessory is used with MacJanet.

Potential for future growth

Very little. Networks for Apple IIgs computers can be used for access to productivity programs like AppleWorks and various instructional programs. Plans should be made to upgrade to Macintosh LC III computers and equip them with Apple II emulator cards if access to Apple II instructional software is critical.

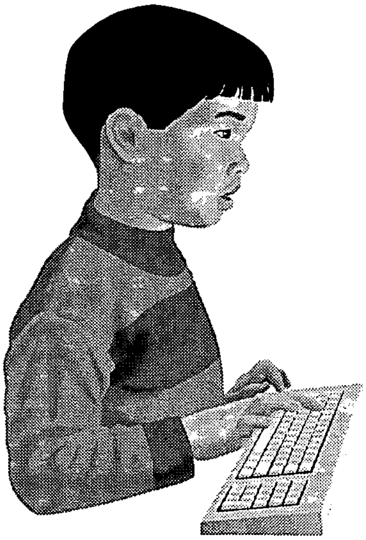

Chapter Sixteen

The never-ending chapter

Chapter 16

The never-ending chapter

Topics covered in this chapter...

- *Our predictions of future developments including:*
  - *Changes in network operating systems*
  - *Changes in client operating systems*
  - *Changes in networking technologies*
  - *Changes in telecommunications*
- *Our unabashed recommendations*
  - *For the small school*
  - *For the large school*
  - *For the school district*

We're ticked off!!!

After all the effort we put into writing this book, the computer industry has gone and changed things on us. Don't they know how hard it is to keep up? This brings to mind the problems faced by one of the major atlas publishers. As you probably know, the borders of the countries of the world have been changing a lot recently. Faced with the prospect of printing an atlas that was out of date before it even reached the bookstores, the managing editor issued a directive that countries would no longer be able to change their borders or their names after a certain date. Following this enormously logical precedent for handling the difficulties of publishing any book that deals with a constantly changing topic, we are issuing the following directive to the computer industry—we will not allow any further changes to computer networking hardware or software after we have written this revision until we are darned good and ready to write a new one!

The concept of the never-ending chapter

Of course we recognize the futility of the directive we just made to the computer industry. It is impossible to hold back technological change. Indeed, we welcome the changes being made in networking because it usually means greater power and utility for the computers we use both personally and at school. But that leads us to a difficult predicament when it comes to a book such as this. How can you keep it current? Our answer has been to create the never-ending chapter. Rather than rewrite the entire book, we have decided to comment on late-breaking developments here. As we mentioned earlier in the book, we recommend that you look at Section Two: An Introduction to the Technical Side of Networking after you read this chapter. We will help you to see the impact new developments have had on what we wrote in this section for the earlier edition. And yes, we know that we will eventually have to rewrite the entire second section of the book when it appears that the never-ending chapter is getting bigger than the rest of the book. But that task ranks right up there with World War III and cleaning out our garages in our priority schemes, and so we will try to get away with this chapter for as long as possible. So let's begin.

Our unabashed predictions

First, a change in name for computers running the Microsoft Windows operating system. Back in the 1980s, it was IBM that first developed the microcomputers that ran operating systems created by Microsoft. When other companies began to make clones of those machines, they were dubbed "IBM compatibles." Today, however, the water has been muddied by the fact that IBM is now working more closely with Apple than with Microsoft. IBM has now licenced the Mac OS, and its new computers based on the PowerPC chip will give users that option. Thus, we are giving up on the moniker "IBM compatible" and replacing it with one that reflects the fact that computers that run the Windows operating system use processor chips from Intel. We now refer to these computers as "WinTel" machines.

Changes in network operating systems

Network operating systems (NOSs) are developing far beyond the file and print services they were originally designed to provide. Two of the most important changes in the NOS are the strengthening of directory services and the ability to act as an application server. In addition, the increasing popularity of Windows NT Server has turned the NOS market for WinTel computers into a much closer race between the latest products from Microsoft and Novell.

Directory services

Directory services, also referred to as naming services, are critical to organizations that have many servers linked together. A complete directory service allows users to access resources on the entire network from any location. It also allows the network administrator to control those resources from any location, thus easing the management of a large enterprise-wide network. With a complete directory service, all the workstations and servers (file servers, printers, CD-ROM servers, communications servers, and application servers) as well as all the users appear in a single directory and can be managed from a single program. While this is most important for those who have many schools linked together in a wide area network, it may also be of value if all the servers in a large school are linked together.

The application server

The trend in programs designed for network operation is toward client server applications, in which a task is divided into two parts, one of which runs on a server while the other part runs on the workstation. The number of programs that are designed to run on a server is growing rapidly, making it desirable for the NOS to support easy program development and safe, efficient running of programs. When file and print services were the only tasks running on servers, a proprietary operating system designed for efficient disk and network access was good; now it is equally important to have a NOS for which it is easy to write programs.

Novell Netware 4.1

Novell made a major change between Netware version 3.12 and version 4.1. The newer version is designed around a new directory system called Netware Directory Services (NDS). NDS is very sophisticated, providing all the features required by a large enterprise. This allows a district to manage all network resources in a WAN from one location, but it adds complexity not required in a school with one server. Because of this complexity, Novell has had difficulty encouraging their user base to move from version 3.12 to 4.1. Novell offers both products for the same price, even though 4.1 has many additional features. It is likely that in the near future support for the older

version will be reduced or withdrawn. All certified netware engineers, technicians certified to support Netware networks, had to be trained with version 4.1 by June 1996 to keep their certification.

The changes in Netware make decisions more difficult for those planning school networks. Many schools continue to use version 3.12 because they do not need the added sophistication of NDS and do not want the added complexity. However, by choosing the older version, they miss out on additional features and may soon have a product that is not well supported. Eventually, they will have to move to version 4.1 and accept its added complexity. For large schools or school districts that have many servers and plan to connect them together into a wide area network, this may be a great benefit. However, smaller schools needing a file server for their one network may find the options promoted by Novell to be less than satisfactory.

Many schools will want to use their servers to run applications other than file and print sharing. The Netware operating system was designed to be very efficient at sharing files, not to run general purpose programs. Programs that run on a Netware server are called Netware Loadable Modules (NLM). They are difficult to write and tend to be hard to find and expensive. A poorly written NLM can easily crash the server, thus they are not appropriate for dabblers, such as the authors. Many schools that choose Netware as the operating system for their file servers will choose a different operating system for their application servers, including communications servers, World Wide Web servers, and database servers. Supporting multiple NOSs means additional costs and headaches, and it would not be the first choice of one who will have to maintain the network.

Choosing Netware is especially difficult for schools because Novell's normal educational pricing is unattractive. Determining the best educational price is difficult if not impossible. Many phone calls to a seemingly endless list of phone numbers reveals a wide variety of price schedules, many of which require a very large purchase to qualify for discounts. After much investigation, it is possible to find the best educational price is higher than the price available to anyone on the street. However, occasional sale prices can be very attractive. Investigating Netware educational pricing could easily convince one that Novell is neither well organized nor particularly interested in the educational market.

Microsoft Windows NT Server

Windows NT is a general purpose operating system often used for client workstations. NT Server is an extension that makes it a powerful NOS as well. NT Server is rapidly growing in popularity and is a serious challenge to Netware. The authors expect that it will become the predominant server software.

Windows NT version 4.0 does not have a sophisticated directory service to rival Novell's NDS. Instead, an NT network consists of domains. Each domain can have multiple servers and can be managed as one entity. All users on the domain are given access to the network or "authorized" by one machine called a primary domain controller. The primary domain controller keeps a database of users and passwords. If the domain includes more than one building, there must be reliable communication links to the primary domain controller to allow users to log on to the network. The user database can be replicated on backup domain controllers in each school, but a link to the primary domain controller is still required for administrative tasks.

If a district does not have reliable communications between schools but still wants to have all servers managed from one location, it is possible to set up many domains and establish "trust" relationships between domains, allowing users of some domains

access to the resources of other domains. These trust relationships can be complicated and inflexible, making network administration more difficult. The situation for the directory services offered by NT Server 4.0 is the reverse of Netware 4.1—it fits well into a simple network but not so well into a district that has a large WAN and requires remote administration.

The feature that is probably driving the growth of NT Server is its capability as an application server. Because it is an extension of the general purpose operating system, Windows NT, it is simple to write programs to run on NT Server. There are many programming languages for Windows NT, including simple batch file interpreters that allow dabblers such as the authors to write small programs to assist in administrative tasks. A wide variety of programs for NT Server can be found on the Internet in the form of freeware and shareware, and most programs written for Windows 95 will also run on NT Server.

NT Server 4.0 includes a variety of application programs that can be useful to schools, including a World Wide Web server and other programs that make it simple to connect the network to the Internet. Other communications programs, including electronic mail servers and Internet news servers, are available free over the Internet. Even if a school or district chooses a different operating system for their file servers, it is likely that they will end up running NT Server on their application servers.

Microsoft is doing a very good job of marketing NT Server to the education community. The pricing is very low and becomes even lower if the program is adopted by an entire school district. In some cases the educational price may be as low as one-fourth the retail price and less than one-fourth the price of Netware. It includes many extras that are expensive additions to Netware, making it a very cost-effective solution for schools.

Network planners should realize that NT Server is relatively new and is not as mature a product as the various versions of Netware. While the administrative tools are adequate, they are not designed for a school setting and may not meet all needs. In particular, it may be difficult to find someone who has done a particular setup in a secure student situation. The authors have found it necessary to download software from the Internet and write batch files to simplify the tasks assigned to school network administrators. While all available documentation was purchased, it was still necessary to resort to the Internet for advice and explanations. As this NOS matures and as more schools select it, NT Server will become easier to implement.

MacJanet 4.2

The MacJanet networking system has been completely rewritten. According to WatNet Technologies, this was required because the underlying logic in all previous versions of MacJanet created some severe limitations in adding new features. With the rewritten code, they feel they have a solid foundation on which to build a networking system that will have a useful life well into the future. However, as with all such projects, there are bugs. Be sure that you use MacJanet version 4.2 or later.

The major features that have been added to MacJanet are:

- a complete new interface
- multiple MacJanet servers on a network
  - users can log on to any server, have personal network disks on any server, and have different privileges on different servers
- improved e-mail system
  - use is more intuitive and a wider range of mail options
  - Internet mail support

    - the MacJanet e-mail system can be linked with the Internet via the SMTP protocol
- improved Apple File Protocol support
- improved print spooling
- improved control over the creation and management of network resources
- improved support for WinTel workstations
- MacJanet is now scriptable
    - using AppleScript you can now automate many of the features for network creation and information retrieval
    - examples include scripts for the creation of users and network disks, find and change user passwords, and check and download e-mail

The new version of MacJanet represents a major addition to the feature set of this networking software. It remains substantially better than AppleShare for use in the school environment.

Add-ons to NOSs for school use

In earlier chapters, IBM's Classroom LAN Administration System was recommended as a way of helping adapt Netware version 2 or 3 for use in a school setting. IBM has a similar, although much more sophisticated product available for Netware 4.1 called SchoolVista. A similar product from a different company for NT Server is called WinClass. While the authors have not used these programs and thus cannot recommend them, it will be worthwhile investigating such programs. As more schools start using Netware 4.1 and NT Server 4.0, the tools required to make them fit comfortably into a school setting will become available.

Changes in client operating systems

Microsoft Windows 95

Windows 95 has quickly become the world's most popular operating system with home and small business users leading the way. Those supporting large networks have been slower to adopt the new operating system, but the trend is very clear. New software is written for Windows 95 and Windows NT, not for Windows 3.1. For large organizations with large networks, the choice is now between Windows 95 and Windows NT. The advantage of Windows 95 is its compatibility with older software. The disadvantage is that it is not everything that the marketing hype led users, and especially network administrators, to hope for.

While Windows 95 can operate on a diskless workstation, booting from the server, it is not designed to do so and performance suffers greatly. Thus most, if not all, workstations will have local hard drives. These hard drives are the bane of the network administrator's life, especially in a school setting. In addition, many of the features that make Windows 95 a convenient and productive operating system on a stand-alone computer in an office have the opposite effect in a school network.

Windows 95 addresses security on network workstations through system policies. The workstation can be set to not start Windows 95 until the user has logged onto a network file server running either Netware or NT Server. As the user logs on, the system policy is downloaded from the file server, determining what appears on the user's desktop and what applications the user has access to. System policies are a nice addition to the operating system, but they are not yet complete.

Unfortunately, the system policy does not restrict the user's access to the local hard drive. So while users are restricted from running certain applications, they are not

stopped from deleting the programs on the local drive. This is not acceptable in a classroom network, and while the same was true for Windows 3.1, many network administrators hoped for more from Windows 95. Once again, it will be necessary to use additional software or mysterious tricks to secure the local hard drive.

Microsoft Windows NT

Windows NT does provide security for the local hard drive and is designed from the start for security and for use on a network. This is one of the reasons it is becoming so popular on large corporate networks. Windows NT also uses system policies, but they can be extended to limit access to the local drive.

The drawback to Windows NT as a workstation operating system is its incompatibility with many programs written for Windows 3.1. In particular, any program that directly manipulates the computer hardware rather than passing a request through the operating system will be stopped by the security of Windows NT. Luckily, Microsoft has pushed software developers into writing programs for both Windows NT and Windows 95, so the number of programs available for Windows NT is growing. However, many programs are still designed for home and school that will not run on Windows NT. While Windows NT may be a good choice for a high-school lab designed for specific purposes and running a limited set of software, many school networks will require a wider choice of software.

Macintosh System 8

Apple has been working on a major update on its operating system for Macintosh computers for some time. It will likely be called System 8 when it is released, but for now it has been code named Copland. Because of the delay of the new system, Apple seems to be adopting a strategy of a series of small releases as parts of the new operating system become available. This was certainly the case with the release of System 7.5.3, which offers some nice additions and enhancements while not representing a major departure from the existing system. However, System 8 will surely be released during the life span of this revision of our book, and it bears mentioning that this new operating system from Apple will likely contain a significantly enhanced feature set than what exists today in System 7. So keep your eyes and ears open for the release of System 8 and be sure to give it a good look before making any long-term decisions regarding Macintosh computers on your networks.

Changes in networking technologies

Larger networks and multimedia applications are creating a need for faster networks. While users need faster computers to meet their needs for multimedia applications, network administrators want to run these applications from file servers to ease the burden of network administrations. Two changes in network technology are contributing to faster networks: Ethernet switches and the 100-Mb (Megabits) Ethernet.

Ethernet switches

All computer networks allow a number of computers to share a single set of wires. A communications protocol, such as Ethernet, allows the computers to share the wire by sending packets of information. If two computers send data at the same time and their packets collide, the protocol specifies how they should re-send data to avoid further collisions. As the number of workstations on the network grows, the number of collisions will increase, and the performance of all the computers will decrease. Thus in a standard 10 Megabits per second (Mbps) Ethernet network, the combined bandwidth available to all the workstations on the network is not more than 10 Megabits

per second. This sharing of the network bandwidth is often the performance limit on networks with relatively fast computers. One way to stretch the limit is to replace the standard Ethernet hubs with switches.

An Ethernet switch acts like a standard hub, except it has the effect of breaking up the network into segments, each of which has a full 10-Mb bandwidth. If all workstations are connected to switches, each will have a full 10-Mb bandwidth, a great improvement over a shared network. While this would be the best solution, less expensive standard hubs are often connected to each port on a switching hub. The computers connected to one hub, and thus to one port on the switch, share a 10-Mb bandwidth. This compromise between a standard Ethernet network and a fully switched network is less expensive, but it also has lower performance than a fully switched network.

While a switched Ethernet network has very high performance, it moves the performance bottleneck to the next weakest link—the connection to the file server. The file server has an available bandwidth of 10 Mb, just like all the other computers on the network. Most of the traffic, however, is between the file server and one of the workstations. Thus each of the workstations may be waiting in turn to communicate with the file server. What is needed is a way of increasing the speed of the connection to the file server.

One way of increasing the file server's bandwidth is a duplex connection. A duplex Ethernet connection gives the equivalent of a 20-Mb bandwidth divided between two channels—one moving data each direction. Switches often allow duplex connections to the file server, but many go one step further by providing one even-faster, 100-Mb connection for the file server.

100-Mb Ethernet

The successor to the standard Ethernet network is known as Fast Ethernet and will transfer data at 100 Mb—ten times as fast as a normal 10BaseT Ethernet network. The two competing standards for the 100-Mb Ethernet are: 100BaseT and 100BaseVG. The 100BaseT standard is an extension of the Ethernet standard, using the same method of collision avoidance. It includes three separate wiring schemes: 100BaseTX uses two pairs of category 5 cable, 100BaseT4 uses four pairs of category 3 (voice grade) cable, and 100BaseFX uses fiber optic cable. A wide variety of companies produces 100BaseT components, and the price is falling quickly.

It is not always possible to change a network from 10BaseT to 100BaseT. While the allowable length of any one cable is still 100 m, the total distance across the network is much more limited. The number of hubs that can be connected together is also reduced. Making this change in a large school with multiple wiring racks may be complex and may require a careful consideration of the limitations. If the distances are too great, a fiber optic backbone connecting the hubs may be required. However, some network cards and hubs can support both 10- and 100-Mb operation, making it possible to make the changes in phases.

The second standard for the 100-Mb Ethernet, 100BaseVG, is strongly supported by Hewlett Packard Corporation but by few other suppliers. It uses a method of collision avoidance that is better than 100BaseT for large, heavily loaded networks. It uses four pairs of category 3 (voice grade) cable and has distance limits at least as far as standard Ethernet. Thus it may be easier to upgrade a standard Ethernet network to 100BaseVG than to 100BaseT. Because there no cards or hubs that support both 10BaseT and 100BaseVG, the upgrade must be done all at once. Unfortunately, 100BaseVG does not seem to be gaining much support in the networking market-

place. While many agree that 100BaseVG has technical advantages over 100BaseT, its limited support in the marketplace may make it difficult to find support in the future.

ATM

The advances in switching are not limited to Ethernet networks. New communications protocols have also been designed to take advantage of switching technology. The most important of these is Asynchronous Transfer Mode (ATM). The packets of data sent through these networks are small and simple, enabling switches to route them to the appropriate destination using hardware rather than software methods. Hardware works much faster than software, allowing ATM networks to move data at very high speeds. The standard speed is 155 Megabits per second (Mbps), although IBM is promoting a more affordable standard at 25 Mb, other high performance versions have been tested at hundreds of megabits per second.

ATM was developed for telecommunications and is probably used by your telephone company to send data between cities. It is only now making the transition from the WAN to the LAN. As it becomes more widely used, it should allow easier integration of the networks in schools with district-wide and worldwide networks. The costs of ATM today are too high to use it for every desktop, so it is used for the backbones of large networks, including the Internet. As the costs come down, it will become feasible to use it within a school network; so keep your eye on this technology.

Changes in telecommunications

While this book cannot hope to cover all the telecommunications issues that must be dealt with by a school system, some of these issues must be considered by school network designers. Most school networks will eventually be connected to the Internet, and some forethought can make such connections much easier.

The Internet

The two services on the Internet that appeal most to educators are e-mail and the World Wide Web. School networks should be designed with the assumption that they will at some point be connected to the Internet for these services. The information available on the Web makes it invaluable for education. The authors tend to write with a pile of books on one side and a Web browser on the other, and the Web is rapidly becoming the more useful resource.

Allowing many users on a network to browse the Web requires a high-speed, permanent connection. The cost of such a connection can be high—hundreds of dollars a month. However, the cost of the same number of individual users connecting to the Internet with modems on separate phone lines would be even higher. Each school or district will have to decide when it makes sense to start using their networks to allow users to access Internet resources, and budget accordingly.

E-mail alone can be done with a less-expensive, dial-up connection to the Internet. There are many e-mail systems available, some of which are proprietary and do not follow the standards of the Internet. While such systems may be less expensive and may offer attractive features, including gateways to allow mail to be exchanged with the Internet, they will likely be temporary solutions. When the decision is made to move to a true Internet connection, converting from such systems can be difficult and disruptive.

The networking protocol used throughout the Internet is Transmission Control Protocol/Internet Protocol (TCP/IP), and connecting directly to the Internet usually means using it on your network. The protocol used natively on the network may be something different, such as IPX for Netware or NT Server networks or AppleTalk for AppleShare or MacJanet networks. Thus adding the TCP/IP protocol may be necessary, requiring additional software and possibly additional workstation memory.

Networks using TCP/IP have a unique Internet Protocol address for each workstation. These addresses, which look like 204.239.100.120, can be difficult to maintain on each workstation. Automated methods of maintaining these addresses make the tasks much easier. These methods, with mysterious names like BOOTP or DHCP, are beyond the scope of this book but are well worth investigating.

If a network is connected to the Internet, it becomes possible to offer a wide variety of services. Of course, once these services become possible, they will quickly become a necessity; and a variety of application servers will appear on the network. These may include a Web server, a mail server, a mailing list server, a USENET news server, a domain name service server, and possibly a security server or firewall. Whether each of these services runs on the same machine or each has its own machine, maintaining them can be a daunting task. The design of a network that will be connected to the Internet should consider the management of any servers. Whenever another service is added to the network, its value will increase, and so will its costs—both capital and operating.

Intranets

Many corporations have private networks, often called Intranets, that use the same technologies as the Internet but for internal communications. These networks include all of the servers listed in the previous paragraph, as well as products designed to ease collaboration within the company. Examples of such collaborative software, or groupware, are Lotus Notes from IBM, Collabra Share from Netscape, Exchange from Microsoft, Groupwise from Novell, and a variety of worthy competitors from smaller companies. Because these Intranets are often connected to the worldwide Internet, close integration is needed between the internal communications (the Intranet) and the external (the Internet). It is especially important that the two appear seamless to the user, with a common user interface. Because the above products are proprietary and were not designed for use on the Internet, this can be a challenge. Some companies are moving away from such proprietary solutions to the standard Internet services, such as World Wide Web servers.

Our unabashed recommendations

In light of the changes to technology and education since the publication of the first edition of this book, the recommendations of the authors for your school network have changed. This is due to the emergence of new products as well as the increasing demands of educators for network services.

The workstations in your networks will probably come with the operating system: Windows 95 or Macintosh System 7.5.3 or later. Either can meet the needs of the school. For those using WinTel computers, we recommend you keep an eye on Windows NT as a workstation operating system. As soon as it is compatible with the application programs you need, it may become your best choice, largely because of the security it provides for the workstation's local hard drive.

The choice for network hardware is clear—it remains Ethernet using Category 5 unshielded twisted pair cabling. If budgets allow, or as costs come down, switching hubs and/or fast Ethernet should be considered. As technology progresses, ATM may also become feasible for schools with limited budgets.

For the small school

The factors considered in this decision should include the skill sets of school staff, district staff, and trusted vendors. If you have people who can confidently support one of these NOSs, that consideration should outweigh the small differences in features.

WinTel only network

The choice of NOS is between Microsoft NT Server and Novell Netware. While either of these two would be a good choice, NT Server will probably be the least amount of work to set up and maintain over the long term, while still giving maximum flexibility at low cost.

Mixed WinTel and Macintosh network

The choice of NOS now expands to include MacJanet and AppleShare along with Novell Netware and NT Server because all four will support both types of computers as workstations. This decision is less clear cut than any of the others. It will depend on the numbers of each type of computer in the mix and the software requirements of both types of machines. Generally, if there are more WinTel computers in the mix, then we recommend NT Server for ease of network creation and maintenance. If there are more Macs in the mix, then we recommend MacJanet.

Macintosh only network

The choice of NOS is between MacJanet and AppleShare. We continue to recommend MacJanet as the best networking solution for small schools.

For the Large School

WinTel only network

A large school with multiple servers requires more management and coordination, especially if all the servers are connected to a WAN or the Internet. We suggest the school seriously consider both NT Server and Netware. At the time of writing, our preference for complex networks is Netware.

As NT Server becomes more popular and third-party software is developed, it will become a worthy competitor. In fact, if your network will be connected to the Internet (and it will!), you may end up using both Netware and NT Server—Netware as the main NOS and NT Server running Internet applications.

Mixed WinTel and Macintosh network

Again, this decision is less clear cut than the others. The choice of NOS once again expands to include MacJanet and AppleShare along with Netware and NT Server because all four will support both types of computers as workstations. It will depend on the numbers of each type of computer in the mix and the software requirements of both types of machines. Generally, if there are more WinTel computers in the mix, then we recommend Novell Netware to handle larger, more complex networks. If there are more Macs in the mix, then we recommend either Netware or MacJanet.

Macintosh only network

The choice of NOS is between MacJanet, AppleShare, Novell Netware, and NT Server. We continue to recommend MacJanet as the best networking solution. However, larger, more complex networks may require the advanced features of Netware. To make matters even less clear, further development of NT Server promises to make this more of a horse race than ever before. Only a comprehensive evaluation of the needs of users and the requirements of the software to be run on the network can give you the final answer.

For the school district

The district creating a WAN must carefully consider both its needs and its resources when choosing a NOS. Connections to the Internet should be included in the network design. The skills of district and school staff and the availability of support from vendors should play a large role in the decisions. It is likely that Netware will provide a better foundation today for your WAN, but it is unlikely to meet all your needs. NT Server will likely be required for some applications and the strong medicine called UNIX may also be required if your needs are extreme.

Districts should keep a close eye on the capabilities, price, and support for Netware and NT Server. Netware is likely the best choice today, but both the technologies and the marketplace are fluid, making tomorrow's choice a difficult one indeed.

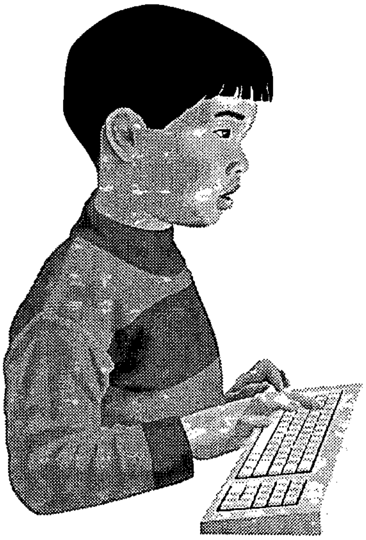

Glossary

Glossary

10Base2	The IEEE standard for Ethernet running on thin coaxial cable.
10Base5	The IEEE standard for Ethernet running on thick coaxial cable.
10BaseT	The IEEE standard for Ethernet running on unshielded twisted pair cable.
access method	A protocol which allows computers to send messages over the network. (also called communication strategy)
access rights	The rights that are given to a user to access specified files stored on a file server.
administrator	A special user ID which has maximum rights on a file server; also the person responsible for maintenance of the network.
AppleShare	Network server software from Apple Computer Inc. which manages a Macintosh network.
AppleTalk	A set of protocols defining how Macintosh computers will communicate over a network. AppleTalk is hardware-independent.
AppleTalk packet	A packet of information sent over an AppleTalk network.
AppleTalk zone	A single AppleTalk network which is connected to others by routers.
Arcnet	A type of network hardware that uses a token passing communication strategy, usually for IBM computers.
AUI	Attachment Unit Interface - a connector on an Ethernet adapter which is used to attach to a transceiver, normally for thick coaxial cable, but also for thin coaxial cable and fiber optic cable.
backbone	A network which connects other networks together, normally through a router.
Baseband	Although this word has other meanings, in this document Baseband refers to a specific type of networking hardware, also known as IBM PCNet Baseband. It communicates at 2 megabits per second over unshielded twisted pair cable.
BNC	A type of connector used with thin coaxial cable.
bus	A network topology in which all the computers are connected to a single wire. Also called trunkline.
cable ring or tray	A metal or plastic ring or tray system used to hold electrical cables in place in a floor or ceiling. An alternate to conduit.
CAI	Computer Assisted Instruction. Software which presents curricular material to students.
CD-ROM	Compact Disc - Read Only Memory. Optical disk storage with capacities to 680 megabytes.
cheapernet	A slang term for Ethernet using thin coaxial cable.
CMI	Computer Managed Instruction. Software which controls what exercises a student does, usually based on success in previous exercises.
coaxial cable	Network cable with a copper wire surrounded by a plastic insulator and a tube of metal shielding. Used with Ethernet and Arcnet networks.

collision	An attempt by two computers to send a message over the network at the same time.
collision avoidance	A communication strategy in which a computer wishing to send a message checks to see if the network is busy before sending. If a collision occurs the computers wait before trying again.
communication server	A computer on the network dedicated to sending messages from the network to a distant computer. Also called gateway.
concentrator	A fancy wiring hub, often with connections to multiple networks.
conduit	Aluminum pipe which provides support and protection for electrical cables.
dedicated server	A file server (or communications server) which does nothing else – it cannot also be used as a workstation.
distributed star	A network topology which has small stars connected to a central star. Used by Arcnet.
DOS	Disk Operating System. The operating system used by most IBM computers. Sold by IBM as PC-DOS and by Microsoft as MS-DOS.
driver	Software which allows a computer to communicate with another device, such as a network adapter or printer.
e-mail	Electronic mail. Messages sent through a computer network. May also include graphics and files.
electronic conference	Software which allows participants to add comments on specific topics. Can run on one computer or can be a distributed conference, running on a number of connected machines.
Ethernet	The most common type of networking hardware, used with IBM computers, Macintosh, UNIX computers and minicomputers.
EtherTalk	A high-speed AppleTalk network that uses Ethernet cabling.
FDDI	Fiber Distributed Data Interface. A 100 megabit-per-second network, originally designed for fiber optic cable.
fiber optic cable	Cable which uses light pulses sent over glass fibers. Not sensitive to electrical interference and able to carry a great deal of information at very high speed.
file server	A computer which contains files that can be used by everyone connected to the network. A file server may also act as a print server, gateway or router.
GEOS	A graphical user interface for DOS, similar to Windows.
gigabytes	A unit of data storage equal to one billion bytes - able to store one billion characters.
GUI	Graphical User Interface. A user interface based on graphics rather than characters, i.e. Macintosh, Windows.
hub	The center of a star topology network. Provides connections to all computers attached to the hub. Also called concentrator or MAU.
ICLAS	IBM Classroom LAN Administrator System. Software which is added to Novell Netware to make it better fit educational use.
ID	An identification code given to each person using a network - often some form of the person's name.
Init	Software which adds features to the operating system upon startup - usually used with Macintosh. See also driver.
Internet	A worldwide collection of networks, mainly educational and research institutions. Uses TCP/IP protocol.

interrupt (IRQ) .... A signal that suspends a program temporarily, transferring control to the operating system when input or output is required.

LAN .... Local Area Network. A network within one building.

LocalTalk .... The plug and play cabling system sold by Apple Computer Inc. This daisy-chain system is inexpensive but slow (.234 Mbps).

MacJanet .... Network server software for Macintosh computers on an AppleTalk network. Sold by WatNet Technologies in Ontario, Canada.

MAU .... Multi-station Access Unit. IBM's name for a Token Ring wiring hub.

Mbps .... MegaBits Per Second. A unit showing the speed of a network.

megabits .... One million bits. Can store one hundred twenty five thousand characters.

multitasking .... Running more than one program at the same time.

Netware .... Network server software from Novell Corp.

Network Interface Card (NIC) .. An adapter which connects a workstation to a network. Usually a card that fits into one of the expansion slots inside a computer.

network operating system .... A popular term for network server software. The type of operating system which runs on the file server and manages communication with the work stations.

Novell .... The company which makes Netware, the most commonly used network server software for IBM computers.

operating system .... The software of a computer which controls the execution of programs.

OS/2 .... Operating System/2. An operating system developed by IBM to replace DOS.

PBX .... Private Branch Exchange. A telephone system which automatically chooses the next available line.

peer-to-peer .... A network which lets any workstation contribute resources to the network while still running local application programs.

peripheral .... A device attached to a computer i.e. printer, modem, CD-ROM etc.

PhoneNet .... A cabling system from Farallon which runs AppleTalk on unshielded twisted pair wire.

print server .... A computer connected to a printer which spools (stores) printed data for the network. May be dedicated to this task or may also act as a workstation.

print spooler .... The software that holds print jobs sent to a network printer when the printer is busy. Each file is saved on a hard disk and then printed when the printer is available.

remote server .... A file server which is not on the same network, but which is on a network connected through a router.

remote print server .... A print server which is on a workstation, not the file server.

retransmission .... A packet of information which must be sent again after a collision on a network using a collision avoidance communication strategy.

RJ11 .... A normal modular telephone connector using four pins. Used with some unshielded twisted pair networks, such as AppleTalk, Arcnet, IBM Baseband..

RJ12 .... A modular connector for unshielded twisted pair cable with six connectors.

RJ45 .... A modular connector for unshielded twisted pair cable with eight connectors. Used by Ethernet 10BaseT networks.

router .... A device which connects two networks and permits packets of information to move between networks only if necessary.

server .... A computer providing a service to network users.

shielded twisted pair cable ...... Network cable consisting of two wires surrounded by a shield made of metal foil.

stand-alone .... A computer which is not connected to a network.

star .... A network topology in which wires run from a central hub to each workstation.

STP .... See shielded twisted pair cable.

structured cabling system .... A network cabling system in which every computer is connected to the central hub with its own wire. i.e. Token Ring, Ethernet 10BaseT.

subdirectory .... A section of a hard disk drive. Also called folder.

TCP/IP .... Transmission Control Protocol/Internet Protocol. A specification for software that manages communications. Widely used in research and in government applications.

telebridge .... A device which connects remote users to a network via telephone lines.

terminator .... A resistor placed at each end of a bus or daisy-chain network to ensure that signals do not reflect back and cause errors.

thinnet .... Ethernet using thin coaxial cable. Also called 10Base2.

Token Ring .... A token passing network pioneered by IBM which runs at 4 or 16 Mbps.

TokenTalk .... A high-speed AppleTalk network operating on a Token Ring.

topology .... The map or description of the physical connections of a network, describing how the cables are laid out.

trunkline .... A network topology in which all the computers are connected to a single wire. Also called bus.

twisted pair cable .... Network cable made up of two conductors twisted around each other, similar to telephone.

UNIX .... A powerful operating system developed by AT&T. Popular with a variety of computer types, especially scientific and engineering workstations.

unshielded twisted pair cable .. See also twisted pair cable, above. Has no electrical shielding. Comes in a variety of categories for different network speeds.

UTP .... See unshielded twisted pair cable.

WAN .... Wide Area Network. A communications network which connects together a number of local networks.

WatNet .... The company that makes the MacJanet network software.

Windows .... Workstation software designed by Microsoft which is added to MS-DOS. It provides a graphical user interface (GUI), multitasking, and improved use of memory.

wiring closet .... A room which contains network hubs, patch panels and other network wiring.

wiring rack .... A metal rack which holds network hubs, patch panels, and other connections. The standard rack is 19" wide. May have a locking door.

workstation .... A computer connected to a network which is used by a person to run application programs.

NOTICE

REPRODUCTION BASIS

- [x] This document is covered by a signed "Reproduction Release (Blanket)" form (on file within the ERIC system), encompassing all or classes of documents from its source organization and, therefore, does not require a "Specific Document" Release form.

- [ ] This document is Federally-funded, or carries its own permission to reproduce, or is otherwise in the public domain and, therefore, may be reproduced by ERIC without a signed Reproduction Release form (either "Specific Document" or "Blanket").